

**ENTERPRISE COGNITIVE COMPUTING APPLICATIONS:
OPPORTUNITIES AND CHALLENGES**

Monideepa Tarafdar

Professor of Information Systems, Lancaster University, Management School, United Kingdom
Research Affiliate, Center for Information Systems Research, MIT

Cynthia M. Beath

Professor Emerita at the McCombs School, University of Texas at Austin

Jeanne W. Ross

Principal Research Scientist, Center for Information Systems Research, MIT

Forthcoming in IEEE Computer Society IT Professional

ENTERPRISE COGNITIVE COMPUTING APPLICATIONS: OPPORTUNITIES AND CHALLENGES

Abstract

Enterprise cognitive computing applications are generating a great deal of excitement for organizations. However their business impact is yet to emerge on a large scale. An important reason for this is a lack of understanding of how such applications can contribute to a company's business objectives, and of the challenges associated with implementing them. In this article we provide an overview of cognitive computing applications for the enterprise. In particular we provide a classification of opportunities for developing enterprise cognitive computing applications and describe challenges in implementing them. Our findings are based on a study of fifty-one initiatives of enterprise cognitive computing applications across a broad range of industries in North America, Europe and Asia-Pacific. Given the lack of systematic description regarding what is possible from enterprise cognitive computing, we believe this article will be valuable to researchers and practitioners in unpacking the black box of cognitive computing.

Keywords: enterprise cognitive computing, cognitive computing tools, enterprise cognitive computing benefits, enterprise cognitive computing challenges

ENTERPRISE COGNITIVE COMPUTING APPLICATIONS: OPPORTUNITIES AND CHALLENGES

Cognitive computing applications are those that utilize tools such as natural language processing, image recognition, intelligent search, and decision analysis to adapt their underlying computational and modeling algorithms or processing based on exposure to new data. Cognitive computing is commonly and colloquially known by many other descriptions, such as machine learning, artificial intelligence and deep learning. We define Enterprise Cognitive Computing (ECC) applications as those that introduce cognitive computing into the software enabling an organization's business processes. The goal of ECC is to make business processes more efficient, accurate, relevant and reliable.

ECC applications are generating a great deal of excitement for organizations. However large-scale business impacts remain elusive. Many companies have invested in ECC applications, but most cannot point to significant benefits. An important reason for this is a lack of understanding of how ECC applications can contribute to the company's business objectives. Our research suggests that many companies don't fully appreciate the challenges associated with implementing ECC applications. The Center for Information Systems Research (CISR) at the MIT Sloan School of Management conducted a study on ECC initiatives in forty-one organizations, thirty-three of which are users of ECC applications, and eight of which are developers and vendors of ECC applications.

In the thirty-three user organizations, which represented a broad range of industries and were distributed across North America, Europe and Asia-Pacific, we studied a total of fifty-one initiatives/use-cases of ECC applications, about seventy percent of which were either in production/use or had a working proof of concept. (See Figure 1) In this article, we share our findings to provide an overview of cognitive computing applications for the enterprise. In particular we – (1) provide a classification of opportunities for developing ECC applications; and (2) describe challenges in implementing those applications.




Figure 1: Enterprise cognitive computing initiatives studied in this research – by region and stage of development

Opportunities for Developing Enterprise Cognitive Computing Applications

Two unique capabilities characterize ECC applications. The first is the processing and making sense of increasingly large and growing volumes of data. ECC applications can handle data that, although available to humans, can be overwhelming because of the sheer size of the corpus, such as the entire medical literature on a given topic and the case law of a certain country or state. Moreover ECC applications can adjust and adapt their models based on new data, to return results with ever-increasing consistency with the data. The second capability is the automation of tasks that formerly required human interpretation, such as queries within a call center that can now be interpreted by natural language processing or image recognition that can learn individuals' identities

Thirty five of the ECC applications we studied were in production or in development. They sought to leverage these two capabilities to pursue one of three distinct business objectives - driving operational excellence, delighting customers, or creating a superior experience for employees. See Figure 2. We describe these below and provide examples from among our applications in production.

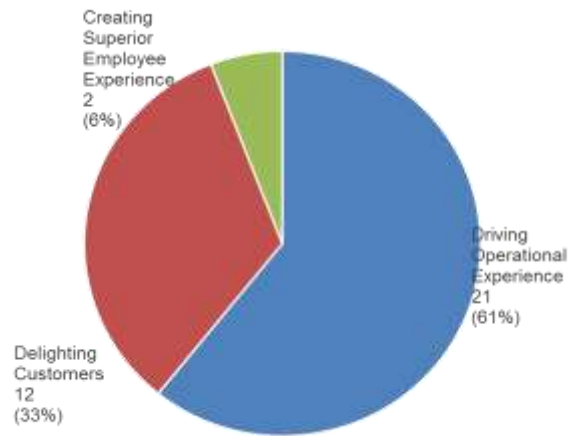


Figure 2: Opportunities for implementing enterprise cognitive computing applications

Driving operational excellence

Operational excellence is key to the organization’s ability to execute its core business processes in a reliable and repeatable way. Organizations that deliver superior operational excellence have stable core and routine processes (Ross et al. 2016). The largest number of use cases we found were ECC applications that enhanced an organization’s operational excellence. These included sophisticated search and retrieval from very large corpuses of technical information such as legal/accounting/financial laws and regulations and medical literature, predictive maintenance of machines, product classification, fraud detection, etc. In all of these examples, ECC applications were used to enhance the speed, accuracy and reliability of various core processes of the organization, and therefore drive its operational excellence.

In terms of speed, two sites (an auditing/consulting firm and a government regulatory body) used text mining, pattern matching, search, and natural language processing in ECC applications to parse large volumes of information and identified case law pertaining to their cases, helping to drastically speed up the grunt work of their attorneys. A global bank used Bayesian statistics and anomaly recognition for semi-automated fraud detection; the application was able to narrow down the list of probable cases for further consideration by the bank’s risk managers, thus speeding up the fraud detection process.

In terms of accuracy, an online retailer and a consumer analytics firm used pattern detection and data mining to quickly and reliably classify and categorize new products for correct presentation. A pharmaceutical company used genetic algorithm applications to develop models that could more accurately identify compounds that had the greatest likelihood of viability and success.

In terms of reliability, we found a number of use cases involving predictive machine maintenance in data centers. These applications used machine-learning models that learned from past data to predict future breakdowns, helping increase the reliability of the data centers.

Delighting customers

ECC applications intended to delight customers by either offering superior products and services or fostering customer loyalty and engagement formed the second largest number of use cases. These included loan processing and delivery, call center support, claims processing, citizen tax services, delivery of insurance products, and analysis of customer feedback. Several ECC applications allowed organizations to provide innovative, personalized and superior products or services to customers.

An online bank used Bayesian statistical modeling tools to understand customer behavior and demographics and offer loans in under seven minutes to retail customers. The bank's leaders believed that providing loans to customers so easily and quickly gave them a significant edge over their competitors. A leading insurance company used text analytics and pattern matching to speed up their claims analysis process. Similarly, a financial services company provided insurance agents with machine learning tools that enhanced their ability to quickly answer questions about insurance policies and claims.

Some ECC applications strengthen customer loyalty and trust. A pharmaceutical manufacturer used text mining and natural language processing to analyze customer feedback from multiple text and voice sources, which helped identify key product defects. Another bank harvested signals of customers' intent from their interactions on the bank's website. This enabled the bank to identify customers who were at risk of leaving and to proactively offer appropriate products and solutions. This application has reduced customer churn. A government tax

department used text analytics to mine laws for taxation implications so they could provide citizens accurate information about their tax liabilities.

Creating a superior employee experience

High performing organizations empower their employees to seamlessly, quickly and efficiently execute their key day-to-day tasks (Dery et al. 2015). ECC applications can help create a superior employee experience. For example, ECC applications in internal help desks help agents tackle a range of technical problems experienced by employees. An IT services organization used natural language processing, text mining and pattern matching to provide employees solutions to problems with their devices and applications. Machines generated answers to standard problems, while handing off non-routine and complex problems to human experts, who were also supported by the ECC application. A life insurance company in Europe developed a similar application for their computer help desk. In another use case, employees had access to an application that used text mining and pattern matching to provide answers to questions about work related travel such as passport and visa applications, currency, and airline tickets. In all of these organizations, ECC applications helped create a digital environment that delivered a superior employee experience by making routine and often time-consuming processes more reliable and efficient.

Challenges in Implementing Enterprise Cognitive Computing Applications

Although we sensed a lot of enthusiasm for the potential of ECC applications to deliver business benefits, there was a sobering recognition that achieving these benefits requires the organization to address multiple challenges. We observed that failure to acknowledge, understand and tackle these challenges resulted in glacially paced implementations, failed implementations, and in the worst cases, a strong disenchantment with, and deep skepticism of, ECC applications.

Four challenges are particularly important. To take advantage of ECC's ability to process massive amounts of data, business leaders must (a) choose the right tools, and (b) make sure needed data is available to those tools. To take advantage of increased automation, business

leaders must (c) consistently supervise the applications and (d) appropriately allocate responsibilities between humans and machines.

Choosing the right tools

While a large number of cognitive computing tools are available for ECC applications, not all tools are appropriate for all types of tasks. A tool that can predict loan worthiness of potential customers may not be appropriate for answering call center queries, because the nature of computations and data can be different for each. There are, in general, four types of tools (see Table 1), depending on– (1) whether they are point solutions that execute focused and relatively narrow tasks or broad-based applications that offer a suite of solutions; and (2) if they are open sourced or provided by specific vendors.

Point solution tools can tackle specific tasks such as natural language processing, text mining and searching through unstructured data. These tools embed existing and relatively uniform ontologies -- rules of vocabulary, syntax and grammar, or a structured corpus of domain information, such as case law. In principle, organizations can develop this type of tool in-house, procure a vendor-supplied tool, or adapt an open source tool. Among the ECC applications we studied, where firms could benefit from existing ontologies, they adopted open source tools or licensed vendor supplied tools that incorporated these ontologies. For example, an IT services company developed an ECC application for an internal technical help desk applying open source tools that provided ontology for natural language processing. In contrast, the vendor Luminoso developed an ontology for its sentiment analysis tool. Client companies build on this tool's 'base' ontology with data about their own products and customers to predict customer churn or anticipate demands for new product features.

Broad-based tools offer a set of general-purpose cognitive computing components that can be adapted to myriad specific ECC applications. For example, IBM's Watson suite of applications has a variety of components such as natural language processing, pattern detection, and classification. An organization licensing this tool can develop company-specific ontologies by training the general-purpose tool with company-specific data. Broad-based cognitive tools need massive amounts of upfront data and training sets, to develop algorithms. A leading bank in

Europe used a general-purpose cognitive computing tool to develop an ECC application that captured a raft of crucial information about the bank's highest-value customers. Much of the effort in the development of this application involved gathering this data from myriad sources and then teaching the tool the relationships among different data entities. The bank succeeded in adapting the general-purpose tool to its specific ECC application. It was an exception. Because of the significant resources and costs involved in doing this, multiple companies abandoned applications of broad-based tools. For example, a heavy machine manufacturing company tried to implement a trouble-shooting ECC application using a broad-based tool. They did not achieve the desired level of accuracy even after significant upfront effort. Eventually, the company abandoned the initiative while noting that a much simpler search engine might have been a more appropriate choice.

Open source and vendor-based tools require different types of resources and skills. Open source tools call for technical expertise to program and contextualize them to the specific requirements of a particular ECC application. We observed that organizations that successfully drew on open source tools to build their ECC applications had a strong base of technical expertise (e.g., a healthcare claims audit firm, and two pharmaceutical companies). On the other hand, purchasers of focused tools can, in principle, benefit from the continuous improvements that the vendor makes to the base ontology.

One tool vendor in our study mentioned that the high quality of their sentiment analysis tool was the direct result of continual monitoring and analysis by a product team of linguists, data scientists, and neural network scientists that worked to improve the underlying algorithm of the tool. It thus seems more efficient to license vendor tools and add additional company-specific ontologies on top of them to reduce requirements for validation and input from experts within the client organization.

Finally, very technically capable organizations can extend and enhance combinations of open source tools, such as natural language processing, speech recognition and intelligent search, to develop their own broad based tools. An organization in the IT services industry developed a broad based ECC platform largely from different open source components, to provide employees providing help desk support and other customer services.

	Open Sourced Tool	Vendor Provided Tool
<i>Cognitive Computing Tool Type</i>		
Point solutions that execute focused and relatively narrow tasks	<ul style="list-style-type: none"> • Tensor Flow: Helping systems transition from prototype to production • Spark: Anti malware • Apache OpenNLP: Natural Language Processing • Emdros: Text mining • WordNet::similarity : semantic pattern matching 	<ul style="list-style-type: none"> • Luminoso: Sentiment analysis • Naralogic: Data clustering and classification • Ross Intelligence: Legal discovery • Kabbage: Retail loan processing • ipSoft: Call center management
Broad-based tools that offer a suite of solutions	Internally developed ECC platform, leveraging many open source tools	IBM Watson

Table 1: Cognitive Computing Tools

Getting the right data, and getting the data right

ECC applications critically and completely depend on massive and continually growing amounts of data to return outputs that are relevant and accurate. There are three challenges with this data.

The first challenge is identifying the right data. A particular ECC application might require external data, such as legal documents, libraries, journals, consumer reports, or real time information feeds, or the organization’s internal data, such as product manuals, troubleshooting databases, operational databases, and subject matter expert knowledge. For example, an audit firm procured large external databases of accounting law manuals and financial statements for its accounting law analysis ECC application, while a heavy equipment manufacturer required

product manuals and troubleshooting databases to be accessed by its field service engineers in its product trouble shooting ECC application.

The second is making sure that needed data are available. For instance, search and query ECC applications require significant amounts of test data to train the model.

The third challenge is getting the data into the right format. ECC applications require data that is appropriately cleaned, formatted and structured (Early, 2016). Format problems may exist even with an organization's internal data. A bank found much of its data on high value customers in the locally stored files or memories of relationship managers. All of this had to be digitized, tagged and structured in order for it to be accessible to its ECC application. Another bank found that PDF documents could not be used as is because its ECC application could not read that format.

External data, which is not under the control of the organization, may present even more formatting problems. One way to acquire external data, especially that available on webpages, is to look for vendors that offer applications that can 'scrape' such data. A bank deployed an ECC loan processing application from a vendor that used another vendor's program to scrape financial statement data from the websites of applicants' banks. This was done in real time and saved the bank the cost and effort of acquiring the data. Another way to acquire external data is to deploy ECC applications that can interface with high traffic websites such as Twitter and TripAdvisor, to analyze data using techniques like text mining. A number of vendors of sentiment analysis ECC tools, such as Luminoso, provide such functionality.

One way to address data-related challenges is to link ECC initiatives to existing data analytics resources and capabilities. A government labor/employment agency developed a proof of concept for an ECC application that provided citizens guidance on jobs available and skills required. It leveraged curated data from a successful predictive analytics application for job-skill matching. In another example, a retail bank's ECC application for approving loan applications was implemented by the team that was responsible for the organization's data analytics applications.

Supervising on an ongoing basis

As ECC applications handle more and more of a company's routine tasks, it is tempting to assume they need little attention. However, a lack of supervision can cause the applications to drift, as the organization's customers and products change over time and new policies and business rules are needed. For example, a call center ECC application would need to take into account new products, customers, services, complaints and solutions in an ongoing way. It thus needs to be supervised by technical and subject matter experts so that the answers it returns are validated for relevance and accuracy, even in the face of these changes. Without supervision the application's models will lose their relevance over time and the application will be able to handle a smaller and smaller percentage of its assigned tasks. Call center agents will therefore rely less on it and eventually it may cease to have any value.

Organizations thus face the challenge of ongoing supervision of the ECC application, that is, of monitoring the performance of the ECC application and regularly re-calibrating it against desired or correct outputs. An online retailer mentioned that it revised its ECC recommendation engine model on a regular basis, with some product categories requiring more frequent revisions than others. An insurance company assigned product managers as subject matter experts to oversee the outputs of their call center ECC application.

Supervision is necessary for maintaining and managing the quality of the knowledge base residing in the ECC application. As noted earlier, vendors provide some of the supervision of their tools. A vendor of a case law analysis tool mentioned that they collected both system-generated and customer-generated feedback, which served as a basis for further improvements to their product's algorithms. Their system recorded the number of times a client-user actually clicked on the output links of a particular search query, and they also collected qualitative feedback from their clients about the tool's performance. However, even with vendor provided tools, it is important to monitor the output of the ECC application to ensure fit with changing organizational needs. Supervision is thus necessary to ensure that the tool retains its accuracy and relevance over time.

Clarifying the division of labor between the human and the ECC application

Clarity regarding allocation of responsibilities between the human user and the ECC application is essential to successful deployment of ECC applications. In general, this means that the ECC applications will take on the more routine tasks and leave complex or discretionary tasks to humans. In addition to retaining some prior responsibilities, ECC applications create new human tasks, such as training and sustaining the application. If executives have a good grasp of how cognitive computing works, they are more likely to make the necessary investments for data preparation, integration with other enterprise applications and development of ECC governance structures.

Understanding which aspects of a business process are routine enough to be handed over to the machine and which aspects need to be retained by human users is often, as we found, a matter of iterative use and learning. Two banks we studied had ECC applications for detecting fraud. They experimented with the application to understand which of the tasks could be performed by the application, which ones needed validation by human users, and how exceptions should be handled. Lack of clarity in articulating the respective roles of the human user and the ECC application can result in expecting too much of the application or, conversely, using human talent inefficiently.

Companies will gradually learn which tasks are best assigned to people and which are best assigned to the ECC application. Initially, ECC applications intended to help lawyers identify appropriate case law were limited to keyword searches that simply helped lawyers get their hands on potentially useful reference cases sooner. As cognitive computing has been applied to such applications, the application is better able to predict which cases are most useful and how they can be applied. The shift of tasks from legal talent to machines has been gradual. Machines can now do much of the work that law clerks once did. On the other hand, lawyers still provide the reasoning and creativity that convert the case material to the needs of a specific engagement.

Conclusion

The deployment of intelligent algorithms in business processes calls for caution because, as Kurzweil (2006) observed, such algorithms, if not properly understood, can lead to ‘runaway’ and ‘rogue’ information processing logic that is largely disconnected from underlying business realities. To add to Kurzweil’s concern, our research found that unrealistic expectations of mystical powers are also a potential problem. If executives regard cognitive computing as nearly magical, projects are more likely to fail. Executives leading ECC initiatives must become knowledgeable about the abilities and inabilities of cognitive computing tools. Enthusiasm is not enough -- a clear-eyed view of the possible is crucial.

Much has been written about high profile applications of cognitive computing, for example, those that can win Go, Jeopardy or chess. These are complex but focused applications with a fixed set of objectives and logic. In contrast, business processes can have competing objectives and dynamic business rules. We are still very early in our ability to apply cognitive computing applications in the business enterprise.

In this paper, we have described two key aspects of ECC applications. The first describes how they can help improve business processes – by strengthening the organization’s operational excellence, delighting customers and enhancing employee experience. The second provides guidance regarding challenges that organizations need to navigate to successfully implement ECC applications - selecting appropriate tools, ensuring the availability of appropriate data, supervising the ECC application’s performance, and clarifying division of labor between the human user and the ECC application. Given the relative newness of ECC applications, the associated hype around them, and a lack of systematic description regarding what is possible from them, we hope this article will help researchers, practitioners and those with a general interest in how technology benefits organizations, to see past the hype.

References

Dery, K., Sebastian, I.M., and van del Meulen, N., "Building Digital Value from the Digital Workplace," *MIT Center for Information Systems Research Briefing*, Volume XVI, Number 9, September 2016.

Early, Seth, "There is no AI without IA," *IEEE IT Professional*, May-June 2016, pp. 2-8.

Kurzweil, R., *The Singularity Is Near: When Humans Transcend Biology*, Penguin Books, 2006.

Ross, J. W., Sebastian, I. M., and Beath, C. M., "How to Develop a Great Digital Strategy," *MIT Sloan Management Review*, Winter 2017, 7-9.

AUTHORS BIOS

Monideepa Tarafdar is Professor of Information Systems at Lancaster University, Management School, UK, and Research Affiliate at MIT's Center for Information Systems Research. Her current research focuses on human-machine partnership and its consequences. She received her PhD in Management from the Indian Institute of Management Calcutta. She is a member of the AIS. She can be reached at Lancaster University Management School, Management Science Department, Bailrigg Campus, Lancaster, United Kingdom, LA14YX or monideep@mit.edu

Cynthia M. Beath is a Professor Emerita at the McCombs School, University of Texas at Austin . Her current research is on the uptake of artificial intelligence and on organization design for the digital economy. She received her PhD in Management from UCLA. She is an AIS Fellow. She can be reached at PO Box 989, Wimberley, TX 78676 or cbeath@mail.utexas.edu.

Jeanne W. Ross is Principal Research Scientist at MIT's Center for Information Systems Research. Her current research examines digital transformations and how people and machines use data to achieve strategic objectives. She received her Ph.D. in Management Information Systems from the University of Wisconsin-Milwaukee. She can be reached at 245 First Street (E95-1551), Cambridge, MA 02139 or jross@mit.edu.