



Primer on Cognitive Computing

Analyst: Jeff Kelly

Date: January 21, 2015

Modern analytics technologies are quite adept at analyzing and visualizing historical, structured data. Want to know how many widgets your company sold last quarter by zip code and product line? No problem. There's no shortage of data warehouse and business intelligence tools on the market that support interactive dashboards and other visualizations for this type of job.



But what if you're not just interested in what your company sold, but what how your customers *feel* about your products and your company based on their own words? Not only that, but let's say you also want to predict what products your customers are likely to buy (or not buy) from you next or whether they are likely to churn to a competitor? And what if you also want recommended actions, steps your company can take, to capitalize on (or prevent) this behavior. This is where things start to get tricky.

The real dividing line between these two scenarios is that one requires a backwards-looking view of a company's own internal data (sales and store locations data, for

example) and a relatively straight-forward visualization tool to display the results. The other requires the continual processing and analysis of large volumes of unstructured data and content (social media streams, news reports, customer feedback forums) that results in both *predictive* (what is likely to happen) and *prescriptive* (what should I do about it) analytic output.

This is where cognitive computing enters the picture. While definitions vary depending on whom you talk to, in essence cognitive computing refers to a style of advanced analytics that attempts to mimic the way the human brain functions but at a scale that no single person could achieve. While cognitive computing systems are extremely complex, they all require several foundational building blocks to be effective. These include:

1. **A Large Corpus of Data.** While the type of data depends on the particulars of the use case, all cognitive computing systems require a large corpus of data, usually unstructured textual data, upon which to apply various analytics techniques.
2. **Data Storage and Processing.** All that data needs to live somewhere, which means cognitive computing systems must include a scalable, flexible storage layer.
3. **Machine Learning.** Rather than requiring human intervention, machine learning algorithms do just what it says in the name – they learn and adapt to the results of data analysis to fine-tune its models resulting in a closed-loop feedback system.
4. **Natural Language Processing.** Natural language processing (NLP) is a style of analytics that understands the nuances of human communication and can interpret spoken language and textual data in order to perform predictive and prescriptive analytics queries.

While the concept of cognitive computing is not new, a number of developments in the underlying technologies have resulted in significant advancements in cognitive computing in recent years. Namely, new approaches to large-scale data processing and storage such as Hadoop have made it easier and more affordable for cognitive computing practitioners to manage extremely large volumes of unstructured data (a.k.a. Big Data). At the same time, the sheer volume of unstructured data being created on a daily basis means machines are required to parse and make sense of the data.

Despite some claims to the contrary, most cognitive computing systems do not aim to outright replace humans in the decision-making process. Rather, cognitive computing systems complement human decision-making resulting in analytic prowess neither cognitive computing systems nor humans could achieve on their own.

Real-World Use Cases

While cognitive computing is still in the early stages of development, there have already been a handful of successful applications of the technology on display. These include the [Durkheim Project](#), a non-profit initiative that set out to build a cognitive computing system that could help mental health professionals better identify returning military veterans at risk for suicide.

The project, directed by Chris Poulin of cognitive computing company [Patterns and Predictions](#) and funded by DARPA, resulted in a system that applies machine learning algorithms to textual data mined from veterans' social media feeds. It interprets what veterans are saying and sharing via social media and scores the risk for self-harm. Put

into practice, the system could alert mental health professionals *before* a veteran takes the drastic step to end his or her life, allowing for some type of intervention.

Dr. Craig Bryan, Associate Director at the National Center for Veteran Studies at the University of Utah and an advisor to the Durkheim Project, said cognitive computing systems have the potential to significantly aid mental health professionals in identifying at risk patients. Specifically, cognitive computing systems allow mental health professionals to analyze larger volumes of data than would otherwise be possible in addition to mitigating the risk of human bias.

“Of course the limitation is there’s only so much you can do as a human and in advance we assume we know what the indicators of risk will be and we look only for those indicators,” said Dr. Bryan in a recent interview. “What we can do with [cognitive computing] technology, of course, is analyze a much, much larger data set much faster and you could potentially uncover indicators that wouldn’t necessarily be immediately intuitive to humans.”

The applicability of cognitive computing systems in other settings is significant. Consider:



Healthcare. In addition to mental health scenarios, cognitive computing systems can be applied more broadly to clinical and research healthcare settings. At a clinical level, cognitive computing systems can be used to aid physicians and nurses in analyzing large volumes of electronic medical records, which often include text-based clinical notes, to better diagnose and treat patients. From a research perspective, cognitive computing

systems are already in use aiding drug developers mine and understand text-based research literature at a scale impossible for human researchers to achieve on their own.



Security. Cognitive computing systems are in use at three-letter U.S. government agencies to analyzing social media chatter to identify “bad actors” and predict the likelihood of a terrorist attack. Large enterprises across verticals, including the energy and utilities industries, could potentially use similar approaches to likewise identify and prevent potential catastrophic events before they impact critical operations.



Retail. With razor thin margins, retailers of all types are constantly looking for an edge on the competition. Cognitive computing systems can aid retailers in better understanding their customers and potential customers through the analysis of social media feeds, blog posts, news reports and other text-based data sources. NLP technology can also be used to improve the customer experience, enabling customers to ask questions and receive answers about products and services in using natural language queries.

These are just three potential areas where cognitive computing could have a significant impact. Others include financial services, insurance, real estate and hospitality.

Vendor Landscape

The cognitive computing vendor ecosystem is small but active and growing as of January 2015. It includes both vendors that sell cognitive computing systems for use by

enterprises as well as vendors who leverage cognitive computing techniques under the covers to support their own customer-facing products and services. Vendors in the cognitive computing space include:

- **IBM.** [IBM Watson](#) is the most well-known cognitive computing system on the market today. It made a splash in the media when, in 2012, the system took on and beat two returning champions on Jeopardy!, the popular TV quiz show. Since then, IBM has made strides commercializing Watson, making it available to developers as a cloud-based service in late 2014. While a general purpose cognitive computing system, Watson's early successes are largely in the healthcare space.
- **Patterns and Predictions.** The aforementioned Patterns and Predictions offers a cognitive computing system called Predictus, which, like Watson, can be applied to a variety of use cases. Centiment, for example, is an application built on Predictus that analyzes text-based financial data sources, such as news reports from Bloomberg, to score the risk-profile of various securities. Investors use Centiment to mitigating risky positions and identify potentially profitable investment opportunities.
- **Google.** Google acquired cognitive computing start-up [DeepMind](#) in January of 2014. The company has since integrated the DeepMind's NLP technology into its core search product, enabling it to better understand queries written (or spoken) in natural language. Apple similarly leverages NLP as part of Siri, it's mobile "personal assistant" technology.

Conclusion

While cognitive computing holds significant promise, its adoption in the enterprise faces several challenges. These include privacy and regulatory considerations, resistance by executives and front-line workers threatened by new methods of analytic-driven decision-making, and a relative lack of maturity of front-end tools and applications that enable non-expert users to interact with cognitive computing systems.

However, Wikibon recommends that both IT and business practitioners begin to evaluate cognitive computing systems and identify potential high-impact use cases in their respective enterprises (though mindful of potential risks.) While the technology is still maturing, there are already a number of use cases in which cognitive computing systems can deliver significant value today and early adopters increase their likelihood of realizing significant competitive advantage from the use of cognitive computing systems over more risk-averse organizations.