

WHITE PAPER

Cognitive computing reshapes enterprise decision-making

AI-assisted decision-making helps enterprises gain value and insight from their data



Contents

Introduction	3
Key business applications of cognitive computing	4
The road to cognitive computing success	7
Magellan brings machine learning and BI to the masses	7
About OpenText	8

Introduction

Enterprise content, the Internet of Things, social media, real-time, and digital transactions all have one crucial thing in common: generating vast quantities of data at ever-faster rates. With terabytes, petabytes—and soon, exabytes—of data from sensor networks, customer activities, or internal business applications, enterprises have enormous opportunities to gain intelligence about nearly anything they want to track.

The problem is that there is too much data, and much of it is not important. Picking out relevant information from mountains of non-relevant or badly formatted data and assembling these points into meaningful patterns to guide decision makers is a tall order, even for the most powerful business intelligence and analytics software. Another limiting factor is that the software can't act on its own; dozens of expert analysts and data scientists may be required to set it up, and then spend weeks or months sifting through its findings.

Enter **cognitive computing** systems such as OpenText Magellan, an artificial intelligence- powered analytics platform that blends enterprise-grade business intelligence with cognitive computing to enable machine-assisted decision-making, which in turn helps enterprises automate repetitive tasks, gain extra value and insight from their data stores, and operate more effectively. We'll discuss Magellan later in this paper. First, we will provide some helpful background.

Cognitive computing: Always learning how to do better

For most of computing's history, the basic principle has been that software programs can only do the specific tasks their programmers have coded them to do. By contrast, cognitive systems are self-learning systems that mimic human thinking, while processing large data workloads at blinding speed.

Cognitive computing, or **machine learning**, builds on a foundation of more than 50 years of research into **artificial intelligence (AI)**. Cognitive computing systems are not only interactive but self-improving, in the way they refresh or update themselves. As they consume more data, they can expand their insights or "world view," which helps them spot patterns and make concept-based associations for more helpful recommendations.

There are two major methods of machine learning: **supervised and unsupervised learning**. In "unsupervised learning," a cognitive system is "trained" on a set of data that a user wants the machine to find structures or relationships between different inputs— but doesn't tell it what to look for or how to classify the results.

Clustering is a typical unsupervised learning technique: The machine graphs outputs and then groups similar outputs into clusters. This is often used in market segmentation (e.g. figuring out how many different types of buyers there are, and the best ones to aim for) or fraud detection (spotting outliers among credit card users or mortgage applicants). Or it can turn up previously unsuspected language patterns that are clues to moods or actions. (For example, the word "just" in "Oh—hey! I'm just looking for a stapler on your desk...")

In "supervised learning," human users training a machine learning system give each input a corresponding target—that is, baskets to sort them into, or predicted values. **Classification** and **regression** are typical techniques here that help the machine make predictions. For example, it can sort through millions of corporate memos and emails to pick out and classify the ones on sensitive subjects that need to be preserved for legal reasons, or identify customer banking histories that make them likely prospects for a credit card offer.

The users will also periodically check in while training these systems and provide basic guidance (“you’re right; that is important,” or “sorry, that’s not what I meant”) that the system will use to continue learning what’s important, vs. what’s just noise in the data.

Over time, as the system trains, it gets better at understanding what in the data is important to the business and continues to refine that understanding.

Cognitive computing is increasingly recognized as the way forward as organizations embrace data-driven approaches to interaction, both at an internal level and for customer engagements. By 2020, [50 percent of all analytics vendors](#) will offer sophisticated predictive analytics based on cognitive computing, according to a recent prediction from market researcher IDC. Already, we’re seeing that emerging cognitive computing tools cover a lot of ground and can be applied to many situations.

Those benefiting from the power of cognitive computing include:

- **Professional or power users.** Professionals adept at manipulating data or even building their own solutions will find cognitive systems provide much greater speed and responsiveness—and less manual review—when they are seeking insights from corporate content and data. For example, a linguist combing through foreign-language texts could take advantage of cognitive computing capabilities to track difficult terms with unclear translations from one language to another, and come up with consistent terminology.
- **Data scientists and analysts.** Experts charged with deep analysis of data sets can employ cognitive computing to research and test high-level questions, more quickly and efficiently than analyzing large data sets manually.
- **Business users.** While cognitive computing is highly sophisticated, it can be delivered in easy-to-use formats. Business users without statistical training can consume analysis from pertinent sources through prebuilt dashboards and interfaces. If they want to dig a bit deeper, they will find cognitive systems helpful in leveraging pre-built queries and reporting to draw insights.

Key business applications of cognitive computing

The driving business case for cognitive computing is that it transforms content or data into not just data points on a graph but potential actions. It can spot patterns, identify the tone, sentiment, and specific emotions of written content, predict trends, and recommend the most productive next step.

In fact, many market analysts believe “machine-assisted decision making” is the sweet spot for cognitive solutions. Instead of science fiction clichés of out-of-control computers taking over the world, picture well-trained assistants who gradually develop the judgment to make suggestions to their boss, and then can be granted responsibility to make small, repetitive decisions in cases that would be too boring or time-consuming for a human to handle.

Temperature or pressure monitoring, such as in a factory or utility line, is one example. Sensors can deliver new readings dozens of times a second, but it wouldn’t be effective use of a worker’s time to have them just staring at gauges. Instead, a cognitive system could be trained to detect when a temperature is far enough from the norm to alert a human—and to recommend the best way to address the issue. (For example, either “there are no other signs of distress, and all other readings before and after are normal, so this is probably a random blip” or “this is problematic—maybe open the release valve.”)

Such machine-assisted decision making could also apply in many other industries that require frequent checks but only occasional human intervention, such as

detecting fraud, managing investment portfolios, or directing traffic flow through a downtown, stadium, or airport. And it can contribute to higher-level goals, such as improving customer satisfaction, reducing inefficiency, or identifying the best new products for given markets.

OpenText recognizes the current value and potential of cognitive computing across a wide range of industries and business functions. In fact, that's why we built Magellan and its machine learning functions on the framework of our industry-leading OpenText Analytics Suite, which offers sophisticated business intelligence functions, easy-to-use query and display functions, and the ability to process enormous volumes and a nearly unlimited range of data types.

The following are some typical Magellan use cases, reflecting experiences of our early adopters worldwide.

Increased customer satisfaction

Magellan can monitor all sources of enterprise data in real time, detecting and learning patterns, then make decisions based on the data and take appropriate action automatically, all within a split second. This can be applied to behind-the-scenes operations that directly affect customer satisfaction, and therefore, organizational revenues.

A large Asian airport was able to leverage cognitive analytics to increase operational efficiency and competitiveness. Like many global airports, this one suffered from inefficiencies associated with security line backups, foot-traffic bottlenecks, malfunctioning escalators, and broken washrooms. These inefficiencies cut into customer satisfaction, costing millions of dollars every year in lost passenger revenue.

The airport used Magellan to analyze data captured from around the facility and make faster decisions to alleviate problems. For example, by analyzing foot traffic data captured from hallway/doorway sensors and security cameras, and combining it with weather data and real-time flight arrival data from the control tower, Magellan helps the airport manage gates by automatically opening those closest to connecting gates, dispatching gate crews, and alerting facilities staff when and which washrooms to clean for disembarking passengers.

Integrated insights

Cognitive systems can find relationships between data from across the enterprise, even if it comes from disparate functions.

Using Magellan, a law firm was able to find relationships between staffing and billing from its data. This helps the firm monitor and analyze profitability of cases with real-time comparison with similar cases.

Magellan also helped improve the firm's discovery process by digging through its archive of legal contracts, tracking pertinent documents, and incorporating news and social media to support a case. Predictive analytics can anticipate the behavior of judges, juries, and venues based on past cases showing similar environments.

Enhanced sales and marketing

Cognitive systems help marketing campaigns succeed by observing real-time reactions of targeted customers most likely to take an interest in products or services, merged with campaign history data. These systems also help enhance brand and product management through social media analysis.

A consumer brands company uses Magellan to examine customer preferences based on data collected through a variety of interactions, including social media comments. The company also wants to deepen its understanding of its customers so that it can better nurture them through the purchasing funnel. Magellan can play a role in automating purchasing processes. It can also track data from orders, inventory, and shipments to predict and ease potential bottlenecks and assure rapid, event-free product delivery.

The system includes data on customer relationships and incorporates natural language processing plus a rules-based inference and calculation engine. Dimensions and metrics such as historical purchases, financial activities, mergers and acquisitions, regions, customers' profiles, and industries are used to pinpoint customer behaviors and tendencies, detect patterns, or help target specific markets or customers that could be hard to surface using classic methods.

Magellan's tracking and analytics capabilities can also monitor customers' attitudes about products, brands, or loyalty programs. Blended with data from finance, sales, engineering, legal, HR and IT, profitability evolution over time can be analyzed. Predictive analytics also support decisions such the best timing for upselling or cross-selling.

Improved quality of service

Cognitive systems can help assure higher-quality service delivery through knowledge management, and the blending of various data sources inside and outside organizations.

A governmental organization in Asia turned to Magellan to improve services provided to its citizens by monitoring key topics of citizen interest and gathering relevant information from a variety of diverse sources, including newspapers and social media postings. Magellan then calibrates this information against the operational data the government agency acquires from its healthcare system and other government service systems. Quality of service is based on monitoring the sentiment and emotion shown by data provided by citizens.

More accurate compliance

Cognitive systems can play a role in ensuring corporate compliance. For example, enterprises can offer user interfaces to manage compliance mandates. Such systems can digitally examine documents such as contracts, partnership agreements, or marketing materials to detect and flag areas that may be affected by laws, regulations, or existing corporate policies. As a result, compliance staff time is freed up for more pressing or strategic business problems, versus being tied down with manually scanning all corporate materials.

A U.S. defense agency that is going digital sought to better interpret and understand the content of its contracts. The project is massive, with hundreds of thousands of contracts being digitally encoded into the agency's contracts management system. The agency turned to Magellan to examine contracts for terms, key concepts, and numerical terms and provide deep insights.

The road to cognitive computing success

Organizations have only taken the first steps on their cognitive computing journeys with OpenText Magellan. This emerging way of managing and digesting information has no limits, and is capable of every critical function from managing contracts to enriching customer interactions.

With today's data-aware environments and infrastructures—such as OpenText Magellan, which is built on the cutting-edge platform Apache Spark™—cognitive algorithms can work within a single system to realize a broad range of business use cases. By processing and analyzing enormous amounts and varieties of data, cognitive systems will find patterns and deliver insights, in real time.

Magellan's machine learning-powered analytics process isn't a single function, but is based on a step-by-step process that takes data from its raw state at sources and transforms it into actionable information. Here's how to make it work:

1. Formulate a question. Determine a business question or thesis for which the answer is data-driven.
2. Access all the data needed to answer the question. This includes all types—structured or unstructured, one-time or continuous streams.
3. Process, merge, and prepare the data. Combine, federate, and store data from multiple sources and formats, in ever-larger quantities, employing a cognitive system to reconcile differences in formats and context.
4. Learn from the data. Model the data, apply machine learning, and train your system to classify information, recognize patterns, uncover trends, predict behavior, and identify correlations, disparities, or even sentiment and emotion.
5. Deliver insight. Let the system draw conclusions, predict outcomes, and make recommendations.
6. Take action. Act on the insights the system delivers. It should communicate smoothly with your systems of engagement so that they can automatically act, with no need for you to keep making decisions on the same questions.

Magellan brings machine learning and BI to the masses

Magellan is a flexible artificial intelligence (AI) and analytics platform that merges machine learning, advanced analytics, and enterprise-grade business intelligence (BI) with the ability to acquire, merge, manage, and analyze big data, both structured and unstructured. Magellan offers an easy-to-use, pre-integrated, and cost-effective cognitive computing platform to enable machine-assisted decision-making, automation, and business optimization for your organization.

As a unified platform, Magellan dramatically reduces the time, effort, and expertise necessary to implement the technologies required for an AI and analytics solution, relieving organizations of dealing with installation and integration headaches, so they can immediately focus on what's important: analyzing their valuable data. This means businesses of all types have a cost-effective and timely method of leveraging machine learning to drive their critical decisions.

Magellan discovers insights from big data and empowers IT teams, operational users, and business analysts to share findings, make more informed decisions, and take more effective action.

Magellan includes the following features:

A prebuilt, open foundation. Magellan leverages Apache Spark as the foundation for advanced analytics, machine learning, data modeling, and enterprise-grade BI. If businesses assembled these pieces themselves, it would take months of complex development—and that's before even analyzing any data. Magellan provides a simple way to implement an AI and analytics platform, which can then be further customized to an enterprise's individual needs, all at lower total costs of ownership and with far less complexity than other options.



AI-enabled insights. Gleaning insights from disparate forms of data used to be difficult. Now Magellan's AI core makes it not only possible, but also well within reach of both data scientists and business users. Magellan automatically scans existing data with natural language capabilities such as concept identification, categorization, entity extraction, and sentiment and emotion analysis.. In addition, Magellan acts as an enterprise's AI assistant—when given some basic guidance during training, Magellan learns the difference between noise and valuable data. As the system discovers new potential insights through regular and fully automated machine learning tasks, it alerts the right team for further analysis.

An emphasis on self-service. Magellan is built to empower IT and business analysts with self-service interfaces that let them apply sophisticated algorithms and dive deep into massive amounts of data without years of specialized training. IT teams can augment their applications with fast, self-serve analytics that anyone can use. By democratizing data, IT teams ensure that their implementations are valued across the enterprise.

Enterprises in any industry will discover immediate benefits from leveraging the AI-enabled analytics that Magellan offers. Those who are looking to bring AI and machine-assisted decision making to their existing enterprise data stores will find a tightly integrated, flexible, and empowering answer in Magellan. By enabling every user and relieving the burden on IT to develop and maintain a custom solution, Magellan allows enterprises to take stronger action from day one.

About OpenText

OpenText, The Information Company, enables organizations to gain insight through market leading information management solutions, on-premises or in the cloud. For more information about OpenText (NASDAQ: OTEX, TSX: OTEX) visit: opentext.com.

Connect with us:

- [OpenText CEO Mark Barrenechea's blog](#)
- [Twitter](#) | [LinkedIn](#)