

Deloitte.
University Press



Machine intelligence

Technology mimics human cognition
to create value

ARTIFICIAL INTELLIGENCE’S RAPID EVOLUTION HAS GIVEN RISE TO MYRIAD DISTINCT—yet often misunderstood—AI capabilities such as machine learning, deep learning, cognitive analytics, robotics process automation (RPA), and bots, among others. Collectively, these and other tools constitute machine intelligence: algorithmic capabilities that can augment employee performance, automate increasingly complex workloads, and develop “cognitive agents” that simulate both human thinking and engagement. Machine intelligence represents the next chapter in the advanced analytics journey.

DATA’S emergence as a critical business asset has been a persistent theme in every *Tech Trends* report, from the foundational capabilities needed to manage its exploding volumes and complexity to the increasingly sophisticated analytics tools techniques available to unearth business insights from data troves. By harnessing analytics to illuminate patterns, insights, and opportunities hidden within ever-growing data stores, companies have been able to develop new approaches to customer engagement; to amplify employee skills and intelligence; to cultivate new products, services, and offerings; and to explore new business models. Today, more and more CIOs are aggressively laying the foundations needed for their organisations to become more insight-driven.

Artificial intelligence (AI)—technologies capable of performing tasks normally requiring human intelligence—is becoming an important component of these analytics efforts. Yet AI is only one part of a larger, more compelling set of developments in the realm of cognitive computing. The bigger story is *machine intelligence* (MI), an umbrella term for a collection of advances representing a new

cognitive era. We are talking here about a number of cognitive tools that have evolved rapidly in recent years: machine learning, deep learning, advanced cognitive analytics, robotics process automation, and bots, to name a few.

We are already seeing early use cases for machine intelligence emerge in various sectors. For example, a leading hospital that runs one of the largest medical research programs in the United States is “training” its machine intelligence systems to analyse the 10 billion phenotypic and genetic images stored in the organisation’s database. In financial services, a cognitive sales agent uses machine intelligence to initiate contact with a promising sales lead and then qualify, follow up with, and sustain the lead. This cognitive assistant can parse natural language to understand customers’ conversational questions, handling up to 27,000 conversations simultaneously and in dozens of spoken languages.

In the coming months, expect to read about similar use cases as more companies tap into the power of machine. Spending on various aspects of MI is already increasing and is projected to reach

nearly \$31.3 billion in 2019.¹ It is also becoming a priority for CIOs. Deloitte's *2016 Global CIO Survey* asked 1,200 IT executives to identify the emerging technologies in which they plan to invest significantly in the next two years. Sixty-four percent included cognitive technologies.²

Data, now more than ever

What we think of today as cognitive computing actually debuted in the 1950s as a visionary effort to make technology simulate human intelligence. Though somewhat primitive AI technologies were commercially available by the 1980s, it wasn't until the 2000s that AI—and the cognitive computing capabilities that comprise the emerging machine intelligence trend—took off.³

A confluence of three powerful forces is driving the machine intelligence trend:

Exponential data growth: The digital universe—comprising the data we create and copy annually—is doubling in size every 12 months. Indeed, it is expected to reach 44 zettabytes in size by 2020.⁴ We also know that data will grow more rapidly as new signals from the Internet of Things, dark analytics, and other sources proliferate. From a business perspective, this explosive growth translates into a greater variety of potentially valuable data sources than ever before. Beyond the potential to unlock new insights using traditional analytics techniques, these volumes of structured and unstructured data, as well as vast troves of unstructured data residing in the deep web,⁵ are critical to the advancement of machine intelligence. The more data these systems consume, the “smarter” they become by discovering relationships, patterns, and potential implications.

Effectively managing rapidly growing data volumes requires advanced approaches to master data, storage, retention, access, context, and stewardship. From signals generated by connected devices to the line-level detail behind historical transactional data from systems across all businesses and functions, handling data assets becomes a crucial building block of machine intelligence ambitions.

Faster distributed systems: As data volumes have grown larger and analysis more sophisticated,

the distributed networks that make data accessible to individual users have become exponentially more powerful. Today, we can quickly process, search, and manipulate data in volumes that would have been impossible only a few years ago. The current generation of microprocessors delivers 4 million times the performance of the first single-chip microprocessor introduced in 1971.⁶ This power makes possible advanced system designs such as those supporting multi-core and parallel processing. Likewise, it enables advanced data storage techniques that support rapid retrieval and analysis of archived data. As we see with MapReduce, in-memory computing, and hardware optimised for MI techniques like Google's Tensor Processing Units, technology is advancing to optimise our ability to manage exponential data more effectively.

Beyond increases in sheer power and speed, distributed networks have grown in reach as well. They now interface seamlessly with infrastructure, platforms, and applications residing in the cloud and can digest and analyse ever-growing data volumes residing there. They also provide the power needed to analyse and actuate streamed data from “edge” capabilities such as the Internet of Things, sensors, and embedded intelligence devices.

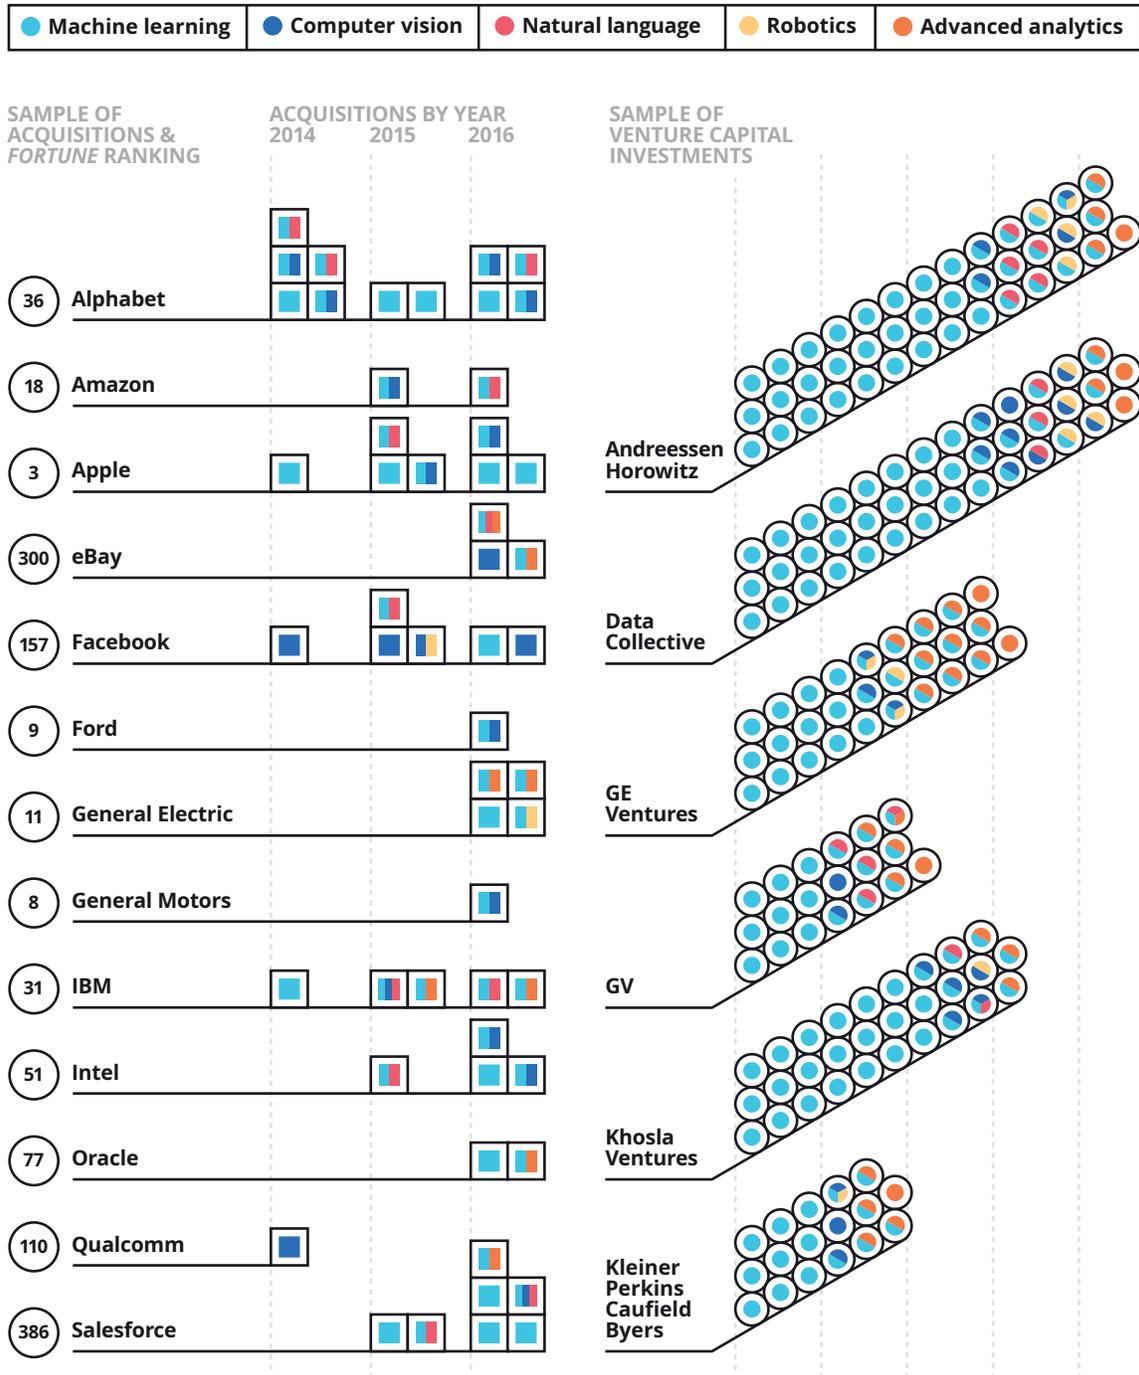
Smarter algorithms: In recent years, increasingly powerful MI algorithms have advanced steadily toward achieving cognitive computing's original goal of simulating human thought processes.

The following algorithmic capabilities will likely see broader adoption in the public and private sectors as machine intelligence use cases emerge over the next 18 to 24 months:⁷

- **Optimisation, planning, and scheduling:** Among the more mature cognitive algorithms, optimisation automates complex decisions and trade-offs about limited resources. Similarly, planning and scheduling algorithms devise a sequence of actions to meet processing goals and observe constraints.
- **Machine learning:** Computer systems are developing the ability to improve their performance by exposure to data without the need to follow explicitly programmed instructions. At its core, machine learning is the

Figure 1. Machine intelligence's impact: Sample acquisitions and investments, 2014–2016

A Deloitte analysis reveals that *Fortune* 500 companies and venture capital firms have recognised the potential of machine intelligence and are strategically investing to build new capabilities.



Sources: Publicly available information on all *Fortune* 500 companies and on technology companies in the *Fortune* 500; Erin Griffith, "Here are the 51 technology and telecommunications companies of the *Fortune* 500," *Fortune*, June 7, 2016, <http://for.tn/22o9uUO>; All investment information provided by cited companies, December 2016.

process of automatically discovering patterns in data. Once identified, a pattern can be used to make predictions.

- **Deep learning:** Developers are working on machine learning algorithms involving artificial neural networks that are inspired by the structure and function of the brain. Interconnected modules run mathematical models that are continuously tuned based on results from processing a large number of inputs. Deep learning can be supervised (requiring human intervention to train the evolution of the underlying models) or unsupervised (autonomously refining models based on self-evaluation).
- **Probabilistic inference:** New AI capabilities use graph analytics and Bayesian networks to identify the conditional dependencies of random variables.
- **Semantic computing:** This cognitive category includes computer vision (the ability to analyse images), voice recognition (the ability to analyse and interpret human speech), and various text analytics capabilities, among others, to understand naturally expressed intention and the semantics of computational content. It then uses this information to support data categorisation, mapping, and retrieval.
- **Natural language engines:** A natural language engine understands written text the way humans do, but it can manipulate that text in sophisticated ways, such as automatically identifying all of the people and places mentioned in a document; identifying the main topic of a document; or extracting and tabulating the terms and conditions in a stack of human-readable contracts. Two common categories are natural language processing for techniques focused on consuming human language and natural language generation for techniques focused on creating natural language outputs.
- **Robotic process automation (RPA):** Software robots, or “bots,” can perform routine business processes by mimicking the ways in which people interact with software applications. Enterprises are beginning to employ RPA in

tandem with cognitive technologies such as speech recognition, natural language processing, and machine learning to automate perceptual and judgment-based tasks once reserved for humans.⁸

How machine intelligence can create value

For CIOs, pivoting toward machine intelligence will require a new way of thinking about data analysis—not just as a means for creating a static report but as a way of leveraging a much larger, more varied data corpus to automate tasks and gain efficiencies.

Within machine intelligence, there is a spectrum of opportunities CIOs can consider:

Cognitive insights: Machine intelligence can provide deep, actionable visibility into not only what has already happened but what is happening now and what is likely to happen next. This can help business leaders develop prescribed actions to help workers augment their performances. For example, in call centres around the globe, service representatives use multifunction customer support programs to answer product questions, take orders, investigate billing problems, and address other customer concerns. In many such systems, workers must currently jump back and forth between screens to access the information they need to answer specific queries.

Cognitive engagement: At the next level on the machine intelligence value tree lie cognitive agents, systems that employ cognitive technology to engage with people. At present, the primary examples of this technology serve consumers rather than businesses. They respond to voice commands to lower the thermostat or turn the television channel. Yet there are business tasks and processes that could benefit from this kind of cognitive engagement, and a new field of applications is beginning to emerge. They will likely be able to provide access to complex information, perform digital tasks such as admitting patients to the hospital, or recommend products and services. They may offer even greater business potential in the area of customer service, where cognitive agents could potentially replace

some human agents by handling billing or account interactions, fielding tech support questions, and answering HR-related questions from employees.⁹

Cognitive automation: In the third—and potentially most disruptive—machine intelligence opportunity, machine learning, RPA, and other cognitive tools develop deep domain-specific expertise (for example, by industry, function, or region) and then automate related tasks. We’re already seeing devices designed with baked-in machine intelligence automate jobs that have, traditionally, been performed by highly trained

human workers. For example, one health care start-up is applying deep learning technology to analyse radiology images. In testing, its system has been up to 50 percent better than expert human radiologists at judging malignant tumors.

In the education field, machine intelligence capabilities embedded in online learning programs mimic the benefits of one-on-one tutoring by tracking the “mental steps” of the learner during problem-solving tasks to diagnose misconceptions. It then provides the learner with timely guidance, feedback, and explanations.¹¹



“Co-bots,” not robots

Facing cost pressures driven by prolonged low interest rates, increased competition, and evolving customer and market dynamics, global insurance provider American International Group Inc. (AIG) launched a strategic restructuring to simplify its organisation and boost operational efficiency. Part of this effort involved dealing with mounting technical debt and a distributed IT department struggling to maintain operational stability.

According to Mike Brady, AIG’s global chief technology officer, by restructuring IT into a single organisation reporting to the CEO, AIG laid the foundation for creating a new enterprise technology paradigm. The first step in this transformational effort involved building foundational capabilities, for which the team laid out a three-part approach:

Stabilise: Overall network performance needed improvement, since users experienced high-severity outages almost daily and the virtual network went down once a week.

Optimise: The strategy focused on self-service provisioning, automation, and cost-efficiency.

Accelerate: To move forward quickly, the team implemented a DevOps strategy to create a continuous integration/continuous deployment tool chain and process flow to deploy software in real time.

AIG turned to machine learning to help with these directives. The company developed an advanced collaborative robot program that utilises built-in algorithmic capabilities, machine learning, and robotic process automation. These virtual workers have been dubbed “co-bots”—a nod to the company’s desire for everyone on staff to treat the virtual workforce as an extension and assistant to employees.

In October 2015, AIG deployed “ARIES,” the company’s first machine learning virtual engineer, to resolve network incidents around the globe. During a 90-day pilot program, ARIES was trained in a “curate and supervise” mode in which the machine operated alongside, and learned from, its human counterparts. In this approach, ARIES learned through observation and experimentation how to assess outage sources and identify probable causes and responses. The co-bot was ready for full deployment on day 91. It’s not that these machines are dramatically faster—in fact, AIG has found that humans take an average of eight to 10 minutes to resolve a typical issue, while co-bots take an average of eight minutes. The benefit lies in its scale: Co-bots can work around the clock without breaks or sleep, and they can resolve incidents so rapidly that queues and backlogs never develop.

Within six months of ARIES’s deployment, automation identified and fixed more than 60 percent of outages. Within a year, ARIES’s machine intelligence, coupled with the expansion of sensors

monitoring the health of AIG's environment, was making it possible to programmatically resolve an increasing number of alerts before they become business-impacting events. The virtual engineer can automatically identify unhealthy devices, perform diagnostic tests to determine the cause, and log in to implement restorative repairs or escalate to a technician with "advice." Additionally, the co-bot correlates network issues, so if data patterns show one device caused 50 incidents in a month, for example, the IT team knows it needs to be replaced. Those efforts have reduced the number of severity 1 and 2 problems by 50 percent during the last year. They have also increased technician job satisfaction. Instead of having to perform mundane and repetitive tasks, technicians can now focus on more challenging, interesting tasks—and benefit from the co-bots' advice as they begin their diagnosis.

Four additional co-bots, each operating with a manager responsible for governance, workloads, training and learning, and even performance management, have been deployed with consistent successful adoptions.

Following the success of the co-bot program in IT, AIG is exploring opportunities to use machine learning in business operations. "We want business to use machine learning instead of requesting more resources," Brady says. "We need to leverage big data and machine learning as new resources instead of thinking of them as new costs." Internal trials are getting under way to determine if co-bots can review injury claims and immediately authorize payment checks so customers need not delay treatment. Other opportunities will likely emerge in the areas of cognitive-enhanced self-service, augmented agent-assisted channels, and perhaps even using cognitive agents as their own customer-facing channels.

"The co-bot approach takes work," Brady adds. "If it's really complex, you don't want inconsistencies in how the team does it. That's where design thinking comes in. Since we started doing this a little over a year ago, we have resolved 145,000 incidents. It's working so incredibly well; it just makes sense to move it over to business process and, eventually, to cognitive customer interaction."¹²

Patients, please

As health care moves toward an outcomes-based model, patients are looking to health insurers to provide the same level of highly personalized customer service that many retailers and banks deliver. To meet this expectation, Anthem, one of the nation's largest health benefits companies, is exploring ways to harness the power of cognitive computing to streamline and enhance its engagement with customers and to make customer support services more efficient, responsive, and intuitive. Anthem's end goal is to change the way the company interacts with its affiliated health plan companies' members over the life of a policy, not just when a claim is filed.

Anthem's strategy grows across three dimensions of machine intelligence: insight, automation, and engagement. In the first phase, the company is applying cognitive insights to the claims adjudication process to provide claims reviewers with greater insight into each case. According to Ashok Chennuru, Anthem's staff vice-president of Provider/Clinical Analytics and Population Health Management, "We are integrating internal payer data—claims, member eligibility, provider demographics—with external data that includes socioeconomic, clinical/EMR, lifestyle and other data, to build a longitudinal view of health plan members," he says.¹³

Currently, reviewers start with a process of document review, patient history discovery, and forensics gathering to determine next steps. With cognitive insight, the new system is continuously reviewing available records in the background to provide reviewers with the full picture right from the start, including supplemental information such as a patient's repeat hospital stays to inform possible care plans or targeted intervention, as well as applying intelligence to flag any potential problems with the claim. By the time the claims representative receives the case, she has the information necessary for a comprehensive assessment.¹⁴

In its next phase, Anthem will start to add cognitive automation to claims processing, freeing up time for adjudicators to dedicate their attention to

patients requiring added levels of support. “By deploying predictive and prescriptive analytics and machine learning algorithms, we will be able to process both structured and unstructured data in a more cost-effective, efficient way,” Chennuru says. At first, the system will identify any potential issues that need addressing and recommend a specific course of action. As the system matures, it can begin to resolve certain issues by itself, if its analysis reaches a certain threshold of certainty based on all signals and inputs. If the level of certainty falls below that threshold, then an adjudicator will still manually review and resolve the claim. As the system’s continuous learning capabilities monitor how adjudicators successfully resolve issues over time, the system will correlate specific issues with proper courses of action to continuously improve its automated resolution accuracy and efficiency.

In the third phase, as Anthem goes deeper into cognitive engagement, the company will more broadly utilise its neural networks and deep learning to engage one-on-one with health care providers recommending individualised care plans for patients. In a shift from simply reacting to claims to proactive involvement in a customer’s care, Anthem will be able to review a patient’s medical history and reach out to providers with recommendations for future care plans.

Anthem’s baseline of semi-supervised machine learning capabilities teach the system how to break down problems, organise them, and determine the best response. During test periods, observers will compare system behavior and performance to the traditional human-driven approach to gauge the system’s efficiency and accuracy.

The company is currently collecting and crunching data, training systems, and streamlining its solutions architecture and technology, and it is seeing positive outcomes across the board as a result of the claims management cognitive insights. A prototype of the automated adjudication system is scheduled to launch in 2017, followed by a minimum viable product version a few months later.

Anthem has built a broad cognitive competency, with multiple teams mapping out use cases to achieve results, evaluating proof of value, and optimising how teams prepare data, tune algorithms, and deliver program availability. “Eventually,” Chennuru says, “we will be able to leverage the platform in many areas such as value-based analytics, population health management, quality management, and to develop insights into gaps in care and cost of care.” Anthem wants to enable as many enterprise cognitive engagements as possible to train its models, optimise its program, and grow its cognitive intelligence to help the company better serve members.

MY TAKE

**MARIA RENZ, VICE PRESIDENT,
TECHNICAL ADVISER TO THE CEO**

**TONI REID, DIRECTOR,
AMAZON ALEXA**

AMAZON

With 2017 ushering in the most exciting time in artificial and machine intelligence history, the Amazon team is empowered to think big and chart new territory.

At Amazon, we believe voice will—and in many ways already has—fundamentally improved the way people interact with technology. While we're a long way from being able to do things the way humans do, we're at a tipping point for many elements of AI and voice technology. Solving unbelievably complex problems every day, voice makes the complex as simple as the most natural and convenient user interface.

The original inspiration for the Amazon Echo was the Star Trek computer. We wanted to create a computer in the cloud that's controlled entirely by voice—you can ask it things, ask it to do things for you, find things for you, and it's easy to converse with in a natural way. We're not quite there yet, but that was our vision.

One of the key capabilities of Alexa, the voice and brain behind Echo, is that she's a cloud-based service that is always getting smarter, in both features and natural language understanding and with improved accuracy. Because her brain is in the cloud, she continually learns and adds more functionality, every hour, every day, which only makes it easier to innovate and add features on behalf of customers.

Since launching Echo in November 2014, we have added more than 7,000 skills to Alexa. Her footprint is expanding across the Echo family of devices and is now embedded within other Amazon hardware (Fire TV and Fire tablets) and in third-party devices such as the Nucleus intercom system, Lenovo Smart Assistant speaker, and the LG Smart InstaView Refrigerator, plus embedding Alexa into cars from companies such as Ford and Volkswagen.

In terms of the surface area she covers and her accuracy within search material, Alexa understands users effectively. Even so, voice technology presents ongoing challenges. When we started working on this, the technology didn't even exist—we had to invent it. We're fortunate to have the power of the AWS cloud to put behind it, and we have teams of incredibly smart speech experts, including talented speech scientists, working on solving these problems.

“DON'T BE AFRAID TO INVENT
ON THE CUSTOMER'S BEHALF.”

We view the benefits to customers and opportunities with AI as nearly limitless. Right now, Alexa mainly operates through the Echo hardware, but in the future her brain will continue to expand through countless numbers of systems and applications. We've made the implementation process easier by making a series of free, self-service, public APIs available to developers with the Alexa Skills Kit (ASK), the Smart Home Skill API, and the Alexa Voice Service APIs.

Ultimately, our developments in machine intelligence, neural networks, and voice recognition advancements should offer our customers new capabilities that are helpful in meaningful ways.

At Amazon, we start any new product or service with a draft press release, imagining the core customer benefits that we would deliver when and if we launch the product. We focus on building the right experience first and solve the hard technical problems later.

With this in mind, we advise looking at your customer base, listening to them, and understanding their core needs and ways in which you can make their lives easier. From there, develop your product or service based off that feedback. That said, don't be afraid to invent on the customer's behalf—customers don't always know what to ask for. If you have the right focus on the customer experience, the rest should fall into place.

In the context of cybersecurity, machine intelligence (MI) offers both rewards and risks. On the rewards front, harnessing robotic process automation's speed and efficiency to automate certain aspects of risk management could make it possible to identify, ring-fence, and detonate (or, alternatively, scrub) potential threats more effectively and efficiently. Leveraging machine intelligence to support cyber systems could potentially help scale data analysis and processing and automate the means of acting in a deliberate manner on the risks these tools identify.

MI's efficacy in this area can be further enhanced by predictive risk and cyber models that extend its data mining net further into largely unexplored areas such as the deep web, and address nontraditional threats it may encounter.

Companies can also harness MI to drive channel activity, strategy, and product design. For example, using capabilities such as deep learning, sales teams can construct fairly detailed customer profiles based on information readily available on social media sites, in public records, and in other online sources. This information can help sales representatives identify promising leads as well as the specific products and services individual customers may want.

But there is a potential downside to MI's customer-profiling power: These same applications can create cyber vulnerabilities. MI might make inferences that introduce new risks, particularly if those inferences are flawed. By creating correlations, MI could also generate derived data that presents privacy concerns. Ultimately, companies should vet derived data based on inferences and correlations.

Indeed, as automation's full potential as a driver of efficiency and cost savings becomes clear, many are discussing broader ethical and moral issues. What impact will automating functions currently carried out by humans have on society, on the economy, and on the way individual organisations approach opportunity? How will your company manage the brand and reputation risk that could go hand-in-hand with aggressive automation initiatives? Likewise, will your organisation be able to thrive long-term in what some are already describing as "the post-work economy"?

Finally, risk discussions should address the "black box" reality of many MI techniques. At this juncture, it may not be possible to clearly explain how or why some decisions and recommendations were made. While there is an ongoing push for algorithmic transparency that could eventually drive development of new means for auditing and understanding assumptions, observing patterns, and explaining how conclusions are justified, those means do not currently exist. Until they do, try to determine where a lack of visibility could be an issue (legal, reputational, or institutional) and adjust your plans accordingly.

As we sail into these uncharted waters, CIOs, CEOs, and other leaders should carefully balance the drive for shareholder value with a host of potential risks to reputation, security, finance, and others that will likely emerge in the years to come.

“WILL YOUR ORGANISATION BE ABLE TO THRIVE LONG-TERM IN WHAT SOME ARE ALREADY DESCRIBING AS “THE POST-WORK ECONOMY”?”

Where do you start?

Few organisations have been able to declare victory in and around data. Even when data was largely structured and limited to information housed within the company's four walls, managing and analysing could prove challenging. Today, sophisticated algorithms and analysis techniques enable us to solve complex scenarios; we can move from passively describing what happened to actively automating business responses. Yet even with rapidly advancing capabilities, some organisations still struggle with data.

The good news is that machine intelligence offers new approaches and technologies that may help us finally overcome some longstanding data challenges:

- **Curate data:** MI techniques can be applied in a largely automated fashion to data taxonomies and ontologies to define, rationalise, and maintain master data. MI can analyse every piece of data, its relationships, and create a derived approximation of data's quality. Likewise, it can potentially provide a means for remedying content or context issues that arise.
- **Bounded and purposeful:** Focus on gaining insight into business issues that, if resolved, could deliver meaningful value. Let the scope of the problem statement inform the required data inputs, appropriate MI techniques, and surrounding architectural and data management needs. By resolving a few of these issues, you may acquire greater license to apply MI to more complex questions.
- **Sherpas welcome:** MI is enjoying its own age of enlightenment, with academia, start-ups, and established vendors bolstering capabilities and adding new techniques. Consider partnering with vendors willing to co-invest in your efforts. Likewise, collaborate with academics and thought leaders who can provide unbounded access to valuable expertise.
- **Industrialised analytics:** Data has become a critical strategic corporate asset. Yet too few organisations have invested in a deliberate, holistic commitment to cultivate, curate, and harness this asset across the enterprise. Industrialising analytics means driving consistent and repeatable approaches, platforms, tools, and talent for all dimensions of data across the enterprise—including machine intelligence. Tactically, this will likely lead to services for data ingestion, integration, archiving, access, entitlement, encryption, and management.

Bottom line

Artificial intelligence may capture more headlines, but the bigger story is machine intelligence, a term describing a collection of advances in cognitive computing that can help organisation move from the legacy world of retrospective data analysis to one in which systems make inferences and predictions. The ability to take these insights, put them into action, and then use them to automate tasks and responses represents the beginning of a new cognitive era.

CONTACTS



WARWICK GOODALL
Director, Technology Consulting
Deloitte MCS Limited
020 7007 9825
wgoodall@deloitte.co.uk



STEPHEN MERCER
Head of Technology
Consulting
Deloitte MCS Limited
0161 455 6836
stmercer@deloitte.co.uk



DAVID ADAMS
Associate Director, Risk Advisory
Deloitte Consulting LLP
0161 455 6813
davadams@deloitte.co.uk

AUTHORS

NITIN MITTAL

US Analytics & Information Management practice leader, Deloitte Consulting LLP

PETER LOWES

Principal, Deloitte Consulting LLP

SANDEEP KUMAR SHARMA, PH.D.

Deputy chief technology officer, Deloitte Consulting LLP

RAJEEV RONANKI

Cognitive Computing and Health Care Innovation practice leader, Deloitte Consulting LLP

JERRY WEN

Managing Director, Technology Strategy and Architecture practice, Deloitte Consulting LLP

ENDNOTES

1. IDC, "Worldwide spending on cognitive systems forecast to soar to more than \$31 billion in 2019, according to a new IDC spending guide," press release, March 8, 2016, www.idc.com/getdoc.jsp?containerId=prUS41072216.
2. Khalid Kark, Mark White, Bill Briggs, and Anjali Shaikh, *2016–2017 Global CIO Survey*, Deloitte University Press, November 10, 2016, <https://dupress.deloitte.com/dup-us-en/topics/leadership/global-cio-survey.html>.
3. David Schatsky, Craig Muraskin, and Ragu Gurumurthy, *Demystifying artificial intelligence*, Deloitte University Press, November 4, 2014, <https://dupress.deloitte.com/dup-us-en/focus/cognitive-technologies/what-is-cognitive-technology.html>.
4. Dan Vesset and David Schubmehl, "IDC FutureScape: Worldwide big data, business analytics, and cognitive software 2017 predictions," International Data Corp., December 2016.
5. Deloitte Consulting LLP, *Tech Trends 2017: Dark analytics*, 2017.
6. Andrew Danowitz et al., "CPU DB: Recording microprocessor history," *ACMQueue* 10(4), 2014, <http://queue.acm.org/detail.cfm?id=2181798>.
7. Schatsky, Muraskin, and Gurumurthy, *Demystifying artificial intelligence*.
8. David Schatsky, Craig Muraskin, and Kaushik Iyengar, *Robotic process automation*, Deloitte University Press, September, 14, 2016, <https://dupress.deloitte.com/dup-us-en/focus/signals-for-strategists/cognitive-enterprise-robotic-process-automation.html>.
9. Tom Davenport and Rajeev Ronanki, *The rise of cognitive agents*, Deloitte University Press, August 26, 2016, <https://dupress.deloitte.com/dup-us-en/focus/cognitive-technologies/rise-of-cognitive-agents-artificial-intelligence-applications.html>.
10. *Economist*, "Automation and anxiety," June 25, 2016, www.economist.com/news/special-report/21700758-will-smarter-machines-cause-mass-unemployment-automation-and-anxiety.
11. Barbara Kurshan, "The future of artificial intelligence in education," *Forbes*, March 10, 2016, www.forbes.com/sites/barbarakurshan/2016/03/10/the-future-of-artificial-intelligence-in-education/.
12. Interview with Mike Brady, global chief technology officer, American International Group Inc., November 22, 2016.
13. Ashok Chennuru, Anthem's staff vice president of Provider/Clinical Analytics and Population Health Management, email communication, January 24, 2017.
14. Deloitte Insights, "Health insurers embrace cognitive," *Wall Street Journal*, December 15, 2016, deloitte.wsj.com/cio/2016/12/15/health-insurers-embrace-cognitive-computing/.