Booz | Allen | Hamilton

THE MACHINE INTELLIGENCE PRIMER

Distinguishing Hype From Reality
In Our New Technological Era

CONSULTING | ANALYTICS | DIGITAL SOLUTIONS | ENGINEERING | CYBER

Machine intelligence (MI)—the ability of machines to perform tasks that would normally require human intelligence—is becoming an increasingly urgent topic of discussion in c-suites, newsrooms, and academic institutions around the world. Rapid gains in computing power, exponential increases in availability of digital data, and new research in computer and data science are enabling algorithms to meet or exceed human capability across diverse tasks. These developments have given rise to machine capabilities that at once inspire and unnerve: trucks that drive themselves, computer programs that develop drug therapies, software that writes news articles and composes music.

INTRODUCTION

Many executives find themselves alternately exhilarated, unsettled, and bemused by the predictions about machine intelligence's (MI) potential following its recent technical advancements.

Self-described "tech evangelists" predict MI will soon slash the cost of doing business, add billions or trillions of dollars to the economy, and liberate employees from drudge work. Cynics caution with equal vigor that we are on the verge of upending financial markets, seeding mass unemployment, and perhaps threatening humanity itself (much good cost savings will do us then). A growing cadre of skeptics disavow both positions, pointing to all the times so-called "artificial intelligence" has failed to deliver in the past and continues to now. Who hasn't cursed a virtual customer service agent?

So, it would be tempting for leaders to set MI exploration aside until there is greater consensus about its capabilities, benefits, and risks. But waiting would be a mistake. MI, to be sure, is a nascent technology with serious limitations. It is often difficult to implement and falls far short of our expansive human capabilities. But even in its current, limited form, MI is being used to substitute for human intelligence in a growing number of tasks. These tasks, taken together, account for meaningful portions of our professional and personal lives. Ceding them to machines—as we've just begun to do—will change our jobs, economy, relationships, and understanding of ourselves. The gravity of this impact calls us to act. We must create—not respond to—our future with MI. We will decide what MI means for our collective future.

The first step is cultivating our understanding of what MI is and what it implies. This Machine Intelligence Primer provides this foundational understanding. In it, we discuss where MI came from, how it got to where it is today, and where it's likely going. We explain and dispel common myths that surround it. While the Primer discusses technical topics, no technological expertise is required to understand it. Instead, it is meant to help executives, practitioners, and curious skeptics alike consider what MI will mean for them and their teams. We hope it will help you consider how MI might be used to create a world that is not simply more efficient, but more equitable, meaningful, and verdant.

In creating the Primer, we were inspired by the works of wide-ranging academics, entrepreneurs, and government organizations whose work came before ours. We hope this book will inspire you to watch, read, and listen to the insights of MI experts like Andrew Ng, Stuart Russell, Fei-Fei Li, and John Launchbury. We are honored to accompany you in your effort to understand and harness MI. We hope that you will share the ideas, aspirations, and worries the information in the Primer evoke for you. We are interested to hear these perspectives from you as we set out on our own journey to make the most of MI.

Please join the conversation at boozallen.com/MIPrimer and contact us at machineintelligence@bah.com

TABLE OF CONTENTS


ERP

1

HOW WE GOT HERE

Definitions and Origins of Machine Intelligence

DEFINING MACHINE INTELLIGENCE

The first challenge to understanding Machine Intelligence (MI) is defining it. A simple web search for a definition will return enough conflicting explanations and bad illustrations of robots to make you regret trying in the first place. In this section, we'll help you gain a realistic understanding of what MI is—and is not.

What is MI?

The best place to start is by understanding what academics, consultants, and internet savants can agree on: the "why" and the "how" of MI. First, the why. The purpose of MI is to produce machines able to perform tasks that would normally require human intelligence. Put simply, MI researchers want to build machines that can perform tasks we do—ideally better than we do them. Next, the how. MI works by enabling machines to perceive, learn from, abstract, and act on data (machine learning). MI and machine learning are often conflated, but machine learning *is a technique for achieving MI*—not the entire field. *Deep learning*, for its part, is a machine learning technique. (Fig. 1.1) These techniques are described in greater depth in the next section, *A Brief History of MI.*

MI and artificial intelligence (AI) are interchangeable terms. MI has grown popular because it lacks negative connotations sometimes associated with AI. Companies don't want you thinking about *The Terminator* when you're deciding whether to buy a virtual home assistant. More importantly, the "artificial" in AI suggests machines possess a fake form of human intelligence, the way a good reproduction of a painting could be confused with the original. In fact, intelligent machines' smarts are much different from human smarts. While our know-how comes from evolution, experiences, culture, and more, machines' intelligence comes mostly from math. Making this distinction in terminology helps remind us how machines differ from us.

What Isn't MI?

Today's MI is not anything like the killer robots we see in movies or on TV. In fact, the cardinal limitation of MI is that—at least for now—it can only perform narrowly defined tasks well. That's why today's MI is referred to as Machine *Narrow* Intelligence. Machine *General* Intelligence (MGI), in which machines would match humans' capacity to perform many tasks (think, C-3Po from *Star Wars*) does not yet exist. Despite recent acceleration in MI development, most researchers agree there is no clear path to creating MGI today. Therefore, claims that MI will soon "replace" us should be viewed with healthy skepticism. We may find our jobs boring at times, but few of us perform just a single task (and thank goodness for that).

Another frequent point of confusion is the relationship between MI and data science. Both fields are rooted in computer science and mathematics, and both are concerned with finding patterns in data to produce actionable insights. The critical difference between the two is that MI is concerned with enabling *machines* to autonomously perceive, reason, and act, while data science is the human tradecraft of producing insights from data that enable *us* to act. Online store managers might use data science to review click rates and shipping data, identifying opportunities to create efficiencies. Or they might use MI to parse other vendors' pricing data and automatically adjust their own store's pricing in real time. The magic of MI is in ceding some work to machines to avoid tedious tasks and allow us to focus on higher-order work.

MI is a field concerned with producing machines able to autonomously perform tasks that would normally require human intelligence by giving them the ability to perceive, learn from, abstract, and act using data.

WHAT IS MI?

- Interchangeable with the term artificial intelligence, or "AI"
- Able to perform certain narrowly defined tasks as well as or better than humans
- A substitute for human intelligence in *certain rote tasks* within jobs

WHAT ISN'T MI?

- Interchangeable with the term "data science"
- Able to perform wide-ranging tasks as well as or better than humans
- A substitute for any human's entire job

Figure 1.1 – What MI Is – What MI Isn't

A BRIEF HISTORY OF MI

MI has long captured our imaginations through science fiction, but it has also existed as an area of serious academic research for more than 60 years.

1950s: The Birth of Machine Learning

In the 1950s, advances in computer science were generating great excitement about the potential for machines to replicate or even surpass human capability at a wide range of tasks. Academics began testing the notion that, instead of providing computers explicit, step-by-step instructions for performing tasks (as is the case with computer programming), machines could be trained to perform tasks on their own.

In 1956, a small group of American researchers proposed an 8-week study of MI at Dartmouth University. The resulting "Summer Research Project on AI" is widely viewed as the birth of formal MI research. The leaders of the project set as their goal to give machines the ability to "use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."[1] Modern MI experts are still working to address each of these challenges. But those early researchers' efforts established the most promising technique we have for achieving MI today: giving machines the ability to learn from data without explicit programming, or "machine learning."

Defining Machine Learning

Machine learning *is a technique for achieving MI that involves training algorithms to parse, learn from, abstract, and act on data.*[2] To accomplish this, we feed algorithms, which are sets of rules used to help computers perform problem-solving operations, large volumes of data from which to learn. Generally, the more data a machine learning algorithm is provided the more accurate it becomes at making inferences.

Just as there is no one way to write a poem or compose a piece of music, there are many approaches to creating machine learning algorithms. Researchers have taken inspiration from a variety of academic fields to develop distinctive styles and structures for algorithms. The "symbolist" approach, for example, takes inspiration from logic and philosophy. It involves codifying human knowledge as sets of rules and programming these rules into computers. Other machine learning approaches draw inspiration from fields as diverse as neuroscience, psychology, evolutionary biology, and statistics.

Over time, various groups of researchers have grown to focus on or favor specific machine learning methods, coalescing into groups MI expert Peter Domingos dubbed the "five tribes of machine learning."[3] (Fig. 1.2) While they may at times disagree about how best to approach machine learning, the term "tribe" isn't meant to call to mind *Survivor*. Machine learning experts often do communicate and collaborate across tribal lines. Instead, the "tribal" umbrella is a helpful way to think about the spirit of discovery and competition among machine learning experts.

1960s–2000s: Challenges to Machine Learning's Advance

Since the 1950s, the tribes of machine learning have been working to bring MI out of the lab and into real-world applications. While MI often fell outside the realm of public consciousness during this time, it did achieve incredible feats. Many will remember, for example, Gary Kasparov's crushing defeat at the hands of IBM's "Deep Blue" computer in 1997. MI systems also beat the world champion at backgammon in 1979 and drove a car more than 1,000 kilometers on a Paris highway in 1994.[4,5]

"An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."

– Machine Intelligence Researchers J. McCarthy, M.L. Minsky, N. Rochester, and C.E. Shannon, 1956

Considering these successes, you may wonder why MI did not take off before the past 5 years.

Until recently, most approaches to producing MI were incredibly labor-intensive, and the resulting systems exhibited only very limited intelligence. This problem is most clearly represented in "expert systems," which grew popular in the 1970s and 1980s. Researchers created expert systems by carefully documenting as data the knowledge of human experts—like accountants and doctors—and representing it as logic-based rules. Once programmed with these rules, computers could answer questions and perform simple tasks from narrow problem domains with great accuracy and scale. A hypothetical doctor expert system input with the symptoms "runny nose" and "grogginess" could, for example, indicate a diagnosis of "cold."

The problem with expert systems was that they needed to be painstakingly reprogrammed each time the rules upon which they relied changed. If the medical community, for example, discovered a new raft of symptoms indicative of a disease, the machine would have no way of knowing without being explicitly reprogrammed with new rules. For this reason, expert systems, and other early forms of MI, were called "brittle"—a term MI researchers use to convey a system is easily broken because it does not evolve to reflect changing circumstances. By the 1980s, most expert systems had been abandoned, and the commercial sector had grown wary of claims about "artificial intelligence" solving business and research problems.[6]

Early Challenges to Advancing Machine Learning

Why couldn't MI experts simply find a way to make systems that were less brittle? The biggest challenge was the lack of raw materials needed to make machine learning work—particularly the dearth of computer-readable data. Remember: Creating a doctor expert system in the 1980s would have required carefully capturing a set of doctors' thoughts, ideas, and inclinations as data. A second, and related challenge, was the lack of computing power to process data for machine learning in problem spaces where computer-readable data was readily available. Running machine learning experiments in the lab often required days as computers pored through massive amounts of data. To understand better why a lack of data and computing power held MI back until recently, it's helpful to take a look at how machines learn [Fig 1.3].

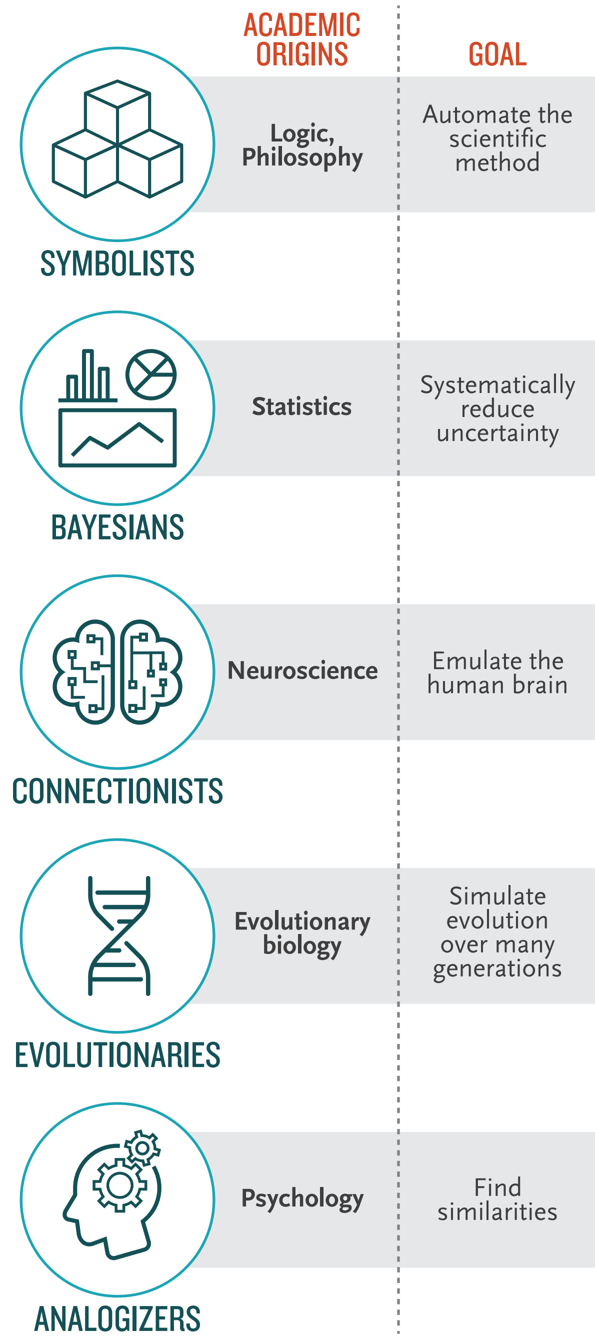

Figure 1.2 – Five Tribes of Machine Learning
Source: Pedro Domingos, "The Five Tribes of Machine Learning" Booz Allen analysis

A QUICK GUIDE TO HOW MACHINES LEARN

The Bottom Line Up Front

All forms of machine learning rely on the availability of large volumes of data to train algorithms. Most forms of commercially viable MI rely on supervised machine learning, which is generally more accurate and reliable than unsupervised learning (at least for now). Therefore, the limiting factor in organizations' ability to take advantage of MI is access to a large set of well-organized, labeled data. Improving techniques for unsupervised learning, which relies only on unstructured, unlabeled data, and semi-supervised learning, which requires less labeled data, will be critical to advancing the field of MI and increasing the number of organizations that have access to it. Until then, when considering an MI initiative, it's important to consider whether you have the data available to support it, and if not, how you might leverage partnerships to access such data.

Supervised Machine Learning

In supervised learning, machine learning algorithms are given training data categorized as input variables and output variables from which to learn patterns and make inferences on previously unseen data. The goal of supervised learning is for machines to replicate a mapping function we have identified for them (for example, "fuzzy ears and whiskers map to the label 'cat'"). Provided enough examples, machine learning algorithms can learn to recognize and respond to patterns in data without explicit instructions. Supervised machine learning is typically used for classification tasks, in which we segment data inputs into categories (e.g., for image classification), and regression tasks, in which the output variable is a real value, such as a price or a volume. The accuracy of supervised learning algorithms typically is easy to evaluate, because there is a known, "ground truth" (output variable) to which the algorithm is optimizing. For example, if we are using supervised machine learning to identify cats in photos, we can easily tell if the algorithm is successful in that task. Most commercial applications of MI rely on supervised machine learning.

Unsupervised Machine Learning

Unsupervised machine learning is an approach to training machine learning in which the algorithm is given only input data, from which it identifies patterns on its own. The goal of unsupervised learning is for algorithms to identify underlying patterns or structures in data to better understand it. Unsupervised learning is closer to how humans learn most things in life: through observation, experience, and analogy. One might, for example, conclude which neighborhood restaurants are popular by observing foot traffic, tidiness, and food quality—no "good" or "bad" label is needed. Unsupervised learning is best used for clustering problems—for example, grouping customers by purchasing behavior. It is also useful for "association," in which algorithms independently discover rules in data; for example, "people who like popsicles also tend to enjoy sorbet." The accuracy of unsupervised learning is harder to evaluate, as there is no predefined ground truth the algorithm is working toward.

And a few twists...

Sometimes researchers combine these approaches in a method called "semi-supervised learning." In this approach, machine learning algorithms are given a small amount of labeled training data and a much larger pool of unlabeled data from which to learn. This approach can combine the best of both worlds—improved accuracy associated with supervised machine learning and the ability to make use of unlabeled data, as in the case of unsupervised machine learning.

Figure 1.3 – A Quick Guide to How Machines Learn

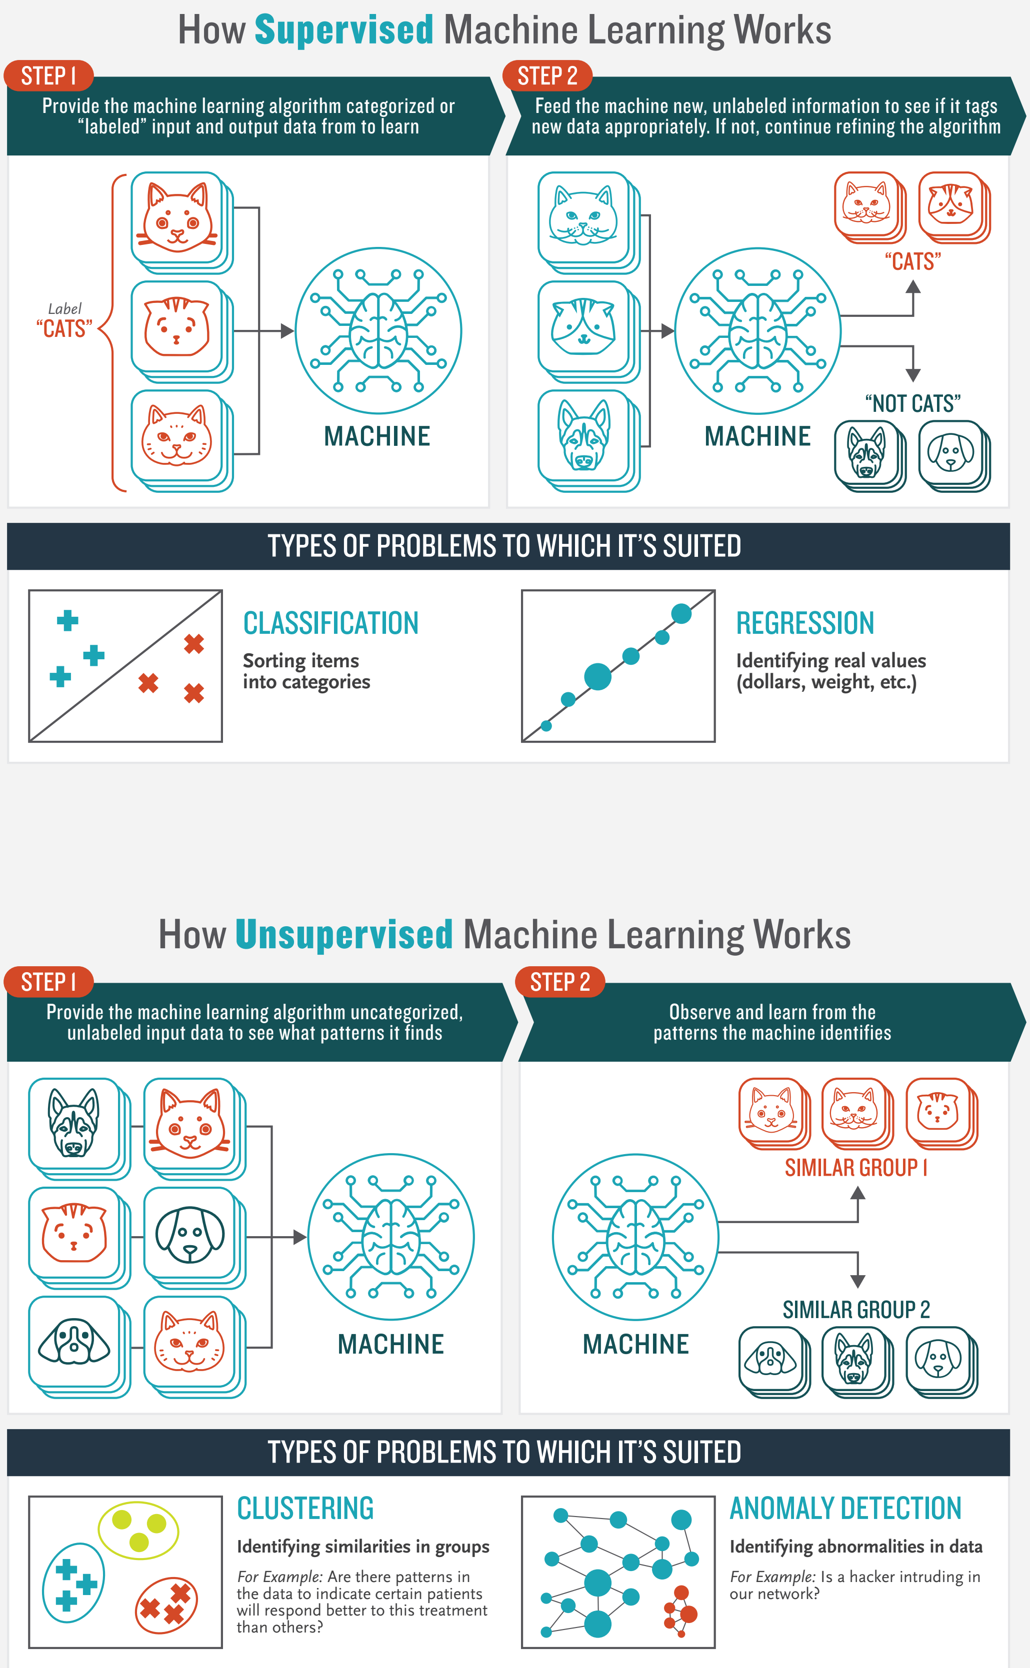

Figure 1.4 – Supervised vs. Unsupervised Machine Learning

"Children don't have adults telling them what each pixel represents in every image they see, or what are the objects present in every image, what is the grammatical structure and the fine sense of every word in every sentence they hear. We extract most of the information from simple observation, and that is what unsupervised learning in principle does."

– Yoshua Bengio, world-renowned MI expert, January 2016

2010–The Present: A Machine Intelligence Renaissance

The convergence of three interrelated technology trends [Fig. 1.5] are propelling MI past the hindrances that have encumbered its development over the past 60 years. Their effects are so highly anticipated that many experts forecast they will spark a new technological revolution akin to industrial revolutions of the past. While we like to avoid hyperbole, we concede the factors contributing to MI's advance today signal the start of a major technological transformation. [7,8,9,10]

The hallmark achievement these trends have enabled is a powerful machine learning technique known as **deep learning**.

Trends Contributing to Machine Intelligence's Ascent (2012–2017)

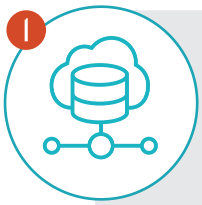

An Exponential Increase in Availability of Digital Data

Training machine learning algorithms requires massive amounts of data. In the past, researchers painstakingly codified pieces of information into digital format to make them useable for machines. Today, our Internet-connected devices produce massive quantities of machine-readable data without our explicit direction.

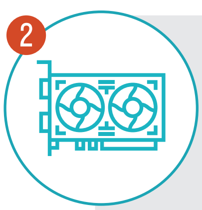

Better Hardware/High-Performance Computing

Not too long ago, running a single machine learning experiment could take days or weeks due to the sheer volume of data the algorithms needed to process to make MI. Today, new types of computing chips have enabled much faster experimentation. Graphics processing units (GPU), originally created for the world of computer gaming in which players move through visually rich digital worlds, have proven exquisitely useful for machine learning. By carrying out many computations in parallel—rather than sequentially, like traditional chips—GPUs work much more efficiently than their forebears.

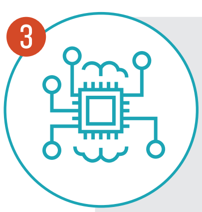

Breakthroughs in Machine Learning Research

Connectionist machine learning research has in recent years produced new algorithms that are remarkably efficient and accurate in their interpretation of massive data. Artificial Neural Networks, a type of algorithm that has existed since the 1950s, now have far more data to learn from and are far more complex than previous iterations. These changes have contributed to the creation of Deep Neural Networks and "deep learning."

Figure 1.5 – Factors Contributing to Machine Intelligence's Recent Advancement [11, 12, 13]
Sources: The New York Times, The Economist, MIT Technology Review, Booz Allen analysis

AN INTRODUCTION TO DEEP LEARNING

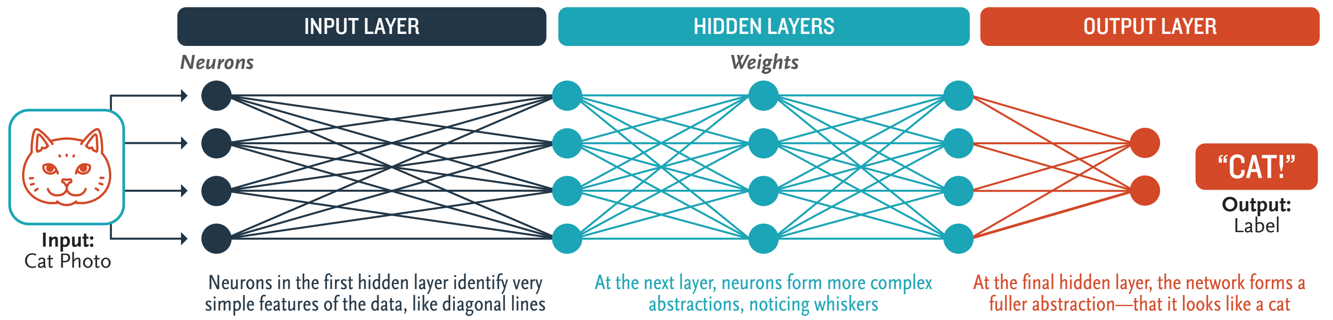

Figure 1.6 – How Neural Networks Interpret Information

Deep learning is a machine learning technique that relies on algorithms called Deep Neural Networks (DNN). DNNs are a form of Artificial Neural Network (ANN)—a *connectionist* algorithm loosely inspired by the circuitry of the human brain [Revisit Fig. 1.2].

ANNs, first invented in the 1950s [Fig. 1.7], comprise artificial neurons, which are connected by numbers called "weights" and arranged into series of layers. This structure resembles the arrangement of biological neurons in our brains. Data are passed through each neuron in a layer and assigned a weight before being passed on to the next. As data are passed through more neurons, the algorithm makes increasingly complex abstractions. At the output layer, the ANN produces a final inference about the data it ingested (e.g., "there is a cat in this photo," or "employee pay is correlated with attrition").

A simple way to think about ANNs is to use the metaphor of vision. When we perceive an object, say, a slice of cake, light waves are bouncing off it, hitting our retinas, and traveling along the optic nerves of our brains in the form of electrical impulses. The individual impulses that travel to our brains are nonsensical in isolation, but our brain pulls them together and forms an internal representation of the cake. Similarly, ANNs take in many singular data points that are nonsensical in isolation and reconstitute them into a meaningful pattern [Fig 1.6].

ANNs can have anywhere from dozens to thousands of layers of neurons. Advances in digital data availability and computing power in the past 5 years have made it possible to make ANNs perform better by massively increasing the number of layers in them. This is where the "deep" in DNN and deep learning comes from, and it's what makes deep learning so special and exciting.[14] DNNs have proven to be remarkably accurate and efficient at wide-ranging tasks—particularly those concerned with categorization, such as image identification and speech recognition.

If you'd like to visualize how a neural network works, we recommend you visit Google's TensorFlow playground, which allows you to see how neural networks are trained and provides more background on how these amazing algorithms work.

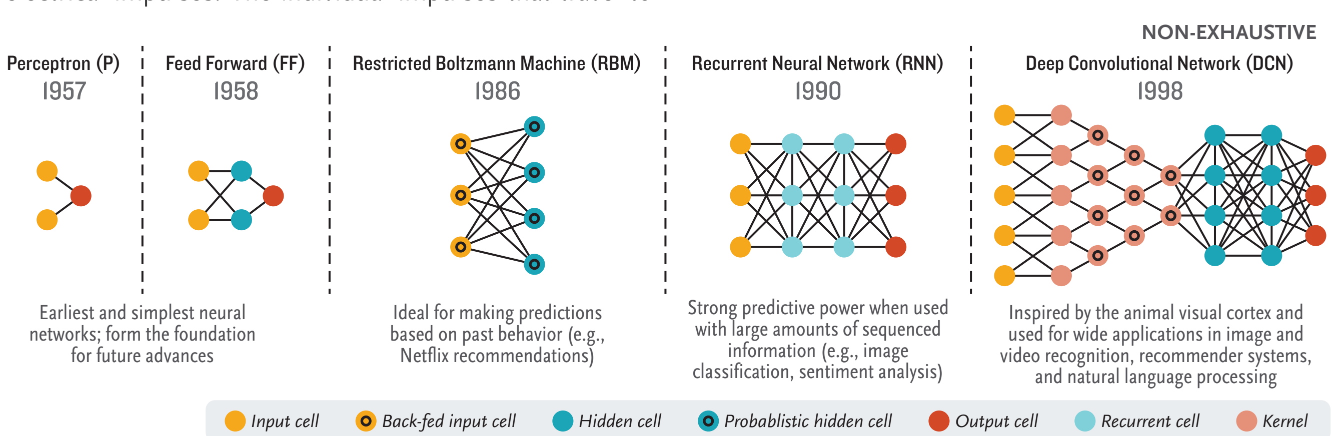

Figure 1.7 – Neural Networks Through the Ages
Sources: "A Mostly Complete Chart of Neural Networks" by the Asimov Institute, Booz Allen analysis

2

SEPARATING HYPE FROM REALITY

What Machine Intelligence Can and Can't Do

"This kind of euphoria of AI has taken over, and [the idea that] we've solved most of the problem is not true"

– Fei-Fei Li, Computer Vision Expert and Associate Professor at Stanford University, June 2017

COME SO FAR, GOT SO FAR TO GO

Recent advances in MI are awe-inspiring—particularly for the many researchers who've dedicated decades to seeing the technology to fruition and endured its many setbacks. New developments in MI research and enabling trends have made possible astounding technological feats: driverless cars, robotic surgeons, software that can compose music, write news articles, and translate languages in real-time. Amid a steady drumbeat of announcements over MI's latest achievements and applications, a sense that humans may soon become passé has begun to take hold in certain circles. MI experts would say, "Not so fast."

Andrew Ng, one of the researchers most responsible for MI's breakthroughs in deep learning, said recently, "The types [of MI] being deployed are still extremely limited...If a typical person can do a mental task with less than one second of thought, [MI] can probably automate it"[15] Stuart Russell, a professor at the University of California, Berkeley who has advanced MI research for decades, said this year, "I think we have at least half a dozen major breakthroughs to go before we get close to human-level [MI]."[16]

In this section, we'll give you an understanding of these researchers' pragmatism by explaining what MI can and can't do given the state of the technology. To start out, consider the following facts about what MI can and cannot do today—and notice which list is longer. (Fig. 2.1).

WHAT MI CAN DO WELL TODAY

- Perform some simple, well-defined tasks as well as or better than humans
- Find and act on patterns in data—including patterns invisible to humans
- Get better at performing a given task when given lots of labeled, well-organized data from which to learn

WHAT MI CANNOT DO WELL TODAY

- Perform any entire job better than humans can
- Explain its mechanism for finding patterns in information or what those patterns mean (beyond giving them a label)
- Understand the context that surrounds a given task
- Learn from unorganized, unlabeled, or small amounts of data
- Perform tasks that require creativity, empathy, or complex judgment

Figure 2.1 – What MI Can and Cannot Do

	1 SIMPLE TASK AUTOMATION	2 PATTERN RECOGNITION	3 CONTEXTUAL REASONING
MI ERA	**YESTERDAY**	**TODAY**	**TOMORROW**
MACHINES...	**Execute narrow, deterministic tasks in a static environment**	**Uncover patterns in data and use them to act in static environments**	**Derive context from multiple data sources and use it to act in dynamic environments**
TECHNOLOGY MATURITY	**Mature/Deployed**	**Maturing/Pilots**	**Emerging/In the lab**
APPLICATIONS (SEE FIG. 2.3)	• Robotic Process Automation • Expert Systems	• Core machine learning software & platforms • Computer vision • Natural language processing & engines • Cognitive robotics	• Contextual machine intelligence ("cognitive" or "semantic" computing)
TASKS SUITED TO THIS TYPE OF MI	• Low complexity, or "swivel chair" • Rules-based ("Deterministic") • High-volume/repetitive • Require no context/reasoning • Closed world environment	• Moderate complexity • No hard and fast rules • Semi-repetitive • Require minimal context/reasoning • Static, steady-state environment	• Moderate to high complexity, creative • No hard and fast rules • Non-repetitive, unpredictable • Require context/reasoning • Real-time, dynamic environment
STRENGTHS	• Speed • Scale • Precision • Risk profile (low) • Easy-to-implement, inexpensive	• Insight discovery • Scale • Precision in aggregate • Ability to self-improve with data	• Closer approximation of human intelligence • Ability to operate in dynamic environments • Ability to self-improve with data
LIMITATIONS	• Inability to manage ambiguity • Minimal self-improvement	• Requires massive data • Data must be tagged/labeled/organized • Resource-intensive to implement • Inability of MI to explain its insights • Inability to grasp context leads to failures, mistakes	• Risk profile (higher)

Figure 2.2 – What Machine Intelligence Can Do: Machine Intelligence Capabilities
Sources: DARPA, "Three Waves of AI, 2016; "Three Ways Work Can Be Automated," Harvard Business Review, Booz Allen analysis

WHAT MI CAN DO

Let's start with the good news: MI research over the past 60 years—and especially the past 5—has given us some truly incredible machine capabilities.

Automate Simple, Rules-Based Tasks

Since the early days of MI, researchers have developed techniques for enabling machines to perform simple, rules-based tasks, usually by codifying human knowledge as sets of logic-based rules and programming rules into computers ("symbolist machine learning"). Expert systems are the quintessential example of this early MI technology. Using the knowledge of accountants, doctors, and other types of human experts, researchers built machines that could answer straightforward questions with precision and scale. Many commercial applications of these systems went out of fashion after the 1980s, due in large part to their rigidity and labor intensiveness. Collecting and formalizing the knowledge of experts takes a great deal of time, and each time the experts' knowledge evolves, the machines require updates. Useful applications of expert system-like products do, however, persist. For example, tax preparation and accounting software can be viewed as an expert system.[17]

New applications of early machine learning research that rely on the same, basic concepts are also still emerging. Robotic Process Automation (RPA), which has grown increasingly popular over the past 3 years, is one such example. In RPA, virtual software "bots" are programmed onto machines to "observe" clerical staff performing simple, high-volume, rules-based tasks. Through observation of their human counterparts' digital behavior, the bots develop rules they can use to repeat the steps themselves, autonomously. RPA is preferable to more classic computer programming-based automation because it adapts to specific organizations' unique processes. RPA is also easier to update than computer programming, as bots can be reconfigured to retrain when an organization's processes change. Organizations are adopting RPA quickly, seizing an opportunity to perform rote tasks more accurately and quickly and enabling staff to spend time on more complex, value-added activities.

Find and Act on Patterns in Data

While simple task automation has great potential to save organizations time and money, the current excitement surounding MI is largely the result of significant improvements in MI's ability to find and act on patterns in data. In the past 5 years, enabling technology trends have made available massive amounts of data and computing power with which to train machine learning algorithms. The development of sophisticated deep learning algorithms inspired by the circuitry of the human brain has further boosted MI's ability to find structure in data.

The math and computer science behind the developments that enable machines to find patterns in data is complex, but the applications we have today are relatively straightforward. Machine learning algorithms ingest, learn from, abstract, and act on data in an "Input Data A generates some Simple Response B" fashion.[18] Most of the tasks pattern recognition MI executes could be completed by a human with a few seconds of thought.

Given the relative simplicity of pattern recognition MI, it may be surprising that it generates such technological euphoria. But, a larger portion of tasks that we perform than one might expect take just a few seconds of thought. Strung together, these moments-long actions make up a meaningful portion of our personal and working lives. A doctor reviewing an x-ray of an arm takes just a moment to determine if there's a break. Driving down the highway, we know almost instantaneously whether we have time to merge lanes. A fluent translator takes just a moment to convey what was said by one party in English to a colleague in French. The applications of pattern recognition data are expansive (Fig. 2.3), and for many industries, transformative.

An accounting of the many applications of pattern recognition MI makes clear why it is inspiring such enthusiasm and investment on the part of government and private sector organizations. It also elucidates growing concerns over how MI will impact our jobs.

It's important to note, on this front, that while many jobs—from those in radiology to trucking—will be significantly impacted by MI, relatively few will be eliminated entirely. Simple, straightforward tasks are building blocks of many of our jobs, but they don't form the full picture of our working lives. They fit into a broader context of more complex, higher-order cognitive tasks and human interactions—and today's machines aren't very good at context.

	Goal	Use Cases	Example
"CORE" MACHINE LEARNING SOFTWARE	Perceive, learn from, abstract, and act on data to perform pattern recognition-related enterprise tasks	• Cognitive or "Intelligent" automation • Machine learning analytic platforms	A software program ingests information about an online store's competitors, adjusting prices in real time to reflect demand
COMPUTER VISION	Perceive, learn from, abstract, and act on data in images and video to perform visual tasks (e.g., navigating an environment, labeling a photo)	• Object identification (i.e., in photos, video) • Image and video tagging • Biometrics • Sentiment analysis • Facial recognition • Scene analysis	Security cameras monitor video footage in real time and identify anomalous behavior, alerting human personnel to risks
NATURAL LANGUAGE PROCESSING & GENERATION (Including Text Analytics)	Perceive, learn from, abstract, and act on data in text and audio in natural language (i.e., conversational speech that does not adhere to rigid rules of syntax, grammar, etc.) to perform language-related tasks (e.g., generate text, translate audio, etc.)	• Machine translation • Speech recognition • Language detection • Sentiment analysis • Document analysis • Virtual assistants • Report generation • News article writing • Insight summarization	An app translates the audio of diplomats speaking in real time, enabling foreign dignitaries to communicate in real time without a human translator
COGNITIVE ROBOTICS	Perceive, learn from, abstract, and act on data to take action in a physical environment	• Autonomous vehicles • Manufacturing "co-bots" • Retail service robots • Home companion robots	An office robot roams the hallways of a university, guiding students to professors' office hours and answering questions about classes

Figure 2.3 – Pattern Recognition Machine Intelligence Applications

WHAT MI CANNOT DO (YET)

Despite at-times astounding advances in MI, it is still quite far from approaching parity with human intelligence.

Understand Context

Intelligent machines suffer from a fundamental inability to perceive or make use of context—the circumstances or information that surround and give meaning to the tasks they perform. While humans rely on prior knowledge and situational awareness gleaned from multiple sources of information to take the right actions at the right times, machines see only the task directly in front of them.

Why is it that MI can perform certain tasks with great competence while lacking basic comprehension? Today's MI is generated using machine learning algorithms designed to address specific problem spaces and trained to recognize patterns in those problem spaces with thousands upon thousands of carefully labeled example data relevant to that specific problem space (supervised machine learning). A given machine learning algorithm's understanding of the world is therefore limited to whatever use case it was designed to address and whatever data was used to train it. An algorithm trained to recognize cats in pictures cannot recognize zebras unless it's been shown thousands of tagged photos from the plains of the Serengeti. Moreover, an algorithm trained to recognize cats in pictures cannot also be trained to translate the word cat from English to Spanish—different use case, new algorithm needed. With each machine learning algorithm operating in relative isolation, it is quite difficult for machines to stitch together context. And, without context, it can be hard for machines to make good decisions in ambiguous or uncertain environments.

Driverless cars are an excellent example of an emerging MI technology that attempts to address this problem. Autonomous vehicles rely on dozens or hundreds of sensors and cameras, each with their own carefully designed algorithms configured to work together to perceive and act on stimuli in the environment: a red light, a car in the next lane, a deer in the road. This works well enough for highway driving, where the roadway scenarios a car encounters are relatively predictable. Driving on city roads, however, is an exponentially more difficult and varied problem. While we can perceive whether a pedestrian is preparing to jaywalk, a driverless car would need dozens of sophisticated algorithms trained on data concerning people's body language, gestures, facial expressions, speed of movement, and even the region the person is in (there are more jaywalkers in New York City than Charleston). MI's issues grasping context is the reason driverless cars for highways are already here, while car companies forecast fully autonomous driving in cities to be 5, 10, or even 20 years away. [19, 20]

"If you compare what computers can do now to what they could do 10 years ago, they've certainly made a lot of progress, but if you compare them to what a 4-year-old can do, there's still a pretty enormous gap."

– Alison Gopnik, Developmental Psychologist, University of California at Berkeley

Explain Itself

The complexity and volume of the math that underlies machine learning algorithms makes them largely inscrutable—even to the researchers who create them. This is particularly true of deep learning—currently the most popular and commercially ubiquitous form of machine learning—where each algorithm comprises many interdependent abstractions that cannot be disentangled or analyzed in isolation. When a deep learning algorithm produces an output incongruent with what its creator, or the organization that commissioned it, expected or intended, there is often no way to know why or how it happened. This has led many to refer to machine learning algorithms as "algorithmic black boxes."

Some MI researchers question why this matters: Humans make decisions all the time that are opaque, even to themselves. And humans, some argue, are inherently more fallible and prone to bias than cold, unfeeling computers. How do we know, for example, that a hiring manager's decisions aren't based on prejudice? Or that an executive's "gut instinct" isn't misguided? We can't split our brains down the middle to understand our impulses and inclinations, so why should it matter that the same be true of the intelligent machines we create?

There are a few important reasons to care about—and work to overcome—the inscrutability of MI. First, MI enables tremendous scale: Machines can process information and make decisions on a much larger scale than humans. Their decisions and actions, therefore, can have much more far-reaching consequences than those of one or a few humans. A hiring manager might in isolation discount a few qualified candidates, but a recruiting algorithm used across a company could erroneously shut off many hundreds of candidates from job interviews. This ties into the second problem with algorithmic black boxes. If MI is perceived to have made a mistake—treated a certain group unfairly or led a driverless vehicle to take a wrong and dangerous turn—it becomes quite difficult to assign blame.

Learn Without Lots of Labeled, Organized Data

MI's reliance on supervised machine learning—and the tens to hundreds of thousands of examples or more of tagged data it requires—is one of its chief limitations today. While a child can learn to understand her mother's tone of voice almost instinctively at a young age, machines require massive reams of labeled data to extract the same meaning from sound. This requires time and resources many organizations and individuals don't have.

Demand for labeled data is so high that some have suggested labeling data might become the new blue collar job of the 21st century (not a job most of us would want!). Organizations have turned to interns, junior staff, and even crowd-sourced labor platforms, like Amazon's Mechanical Turk and Crowdflower, just to prepare their data for machine learning algorithms. This is an expensive solution that precludes smaller, less well-resourced companies from taking advantage of emerging MI applications. [21, 22]

Advancing unsupervised learning will unlock and make more accessible wide applications of MI.

WHERE MI IS GOING NEXT

Researchers are hard at work addressing the limitations of today's MI.

Around the world, companies, academic institutions, and even entire nations are making massive investments in pushing MI technology forward. The level of investment is so great that some of the limitations of MI that we've described here could be overcome in just a few months or years—we can't be sure. But early developments in these areas show promise [Fig. 2.4].

1

Enabling Machines to Understand Context

Today's MI systems lack context: They can perform tasks with great competence, but only in narrowly defined contexts with little ambiguity. Their rigidity stems from their inability to piece together multiple pieces of information from different sources—individual algorithms address specific problems and operate in relative isolation. Researchers are working to make machine learning algorithms more interoperable so that they can perform more complex tasks in dynamic environments. One potential solution researchers are investigating is the use of progressive neural networks: separate deep learning systems connected to share pieces of information.

2

Making Algorithms Explain Themselves

MI systems today are largely opaque—there is no way, even for their creators, to understand how they arrive at the inferences they make. MI researchers are working to address the problem of "algorithmic black boxes" in a few ways. One method bypasses the problem of opening up algorithms completely by instead testing their outputs to detect bias. In 2016, researchers at Google and the Toyota Research Institute co-authored a paper detailing this approach, saying, "Our criteria does not look at the innards of the learning algorithm. It just looks at the predictions it makes."

3

Enabling Machines to Do More with Less Data

Researchers are working hard to reduce MI's massive requirements for data. In 2016, the International Conference on Machine Learning (ICML), a prestigious academic conference first held in 1980, hosted a workshop on data-efficient machine learning to tackle the problem of enabling machines to "learn in complex domains without large quantities of data." Participating researchers explored a variety of techniques for achieving this goal. One promising method is generalized learning, in which machines take knowledge from one domain space and apply it to another rather than receiving an entirely new data set to train on.

Figure 2.4 – Three Hard Problems MI Researchers are Focused on Today [23, 24, 25]

If MI researchers successfully enable machines to understand context, explain their reasoning, and learn from less data, it will usher in a new era of MI in which the technology will be a much closer approximation of human intelligence. MI will then be able to robustly operate in dynamic environments, easily picking up new information as it experiences the world and explaining its mistakes along the way. This era of contextual reasoning is what many MI researchers aspire to today.

3

GETTING STARTED

Advice for Using Machine Intelligence

Given all that you now know about MI, how can you get started on using it to address problems you see at your organization, or even in your day-to-day life? First, it helps to consider what you've learned about the kinds of problems MI is and is not useful for today.

As an emerging technology, the vendor, research, and expert ecosystem within which MI is developing becomes more complex by the day. It is difficult to discern which MI technologies and use cases will bring value to a problem space, and even more complicated to disentangle what MI implies for an organization's future when it comes to its clients, teams, and long-term trajectory. As with any strategic initiative, the best place to begin is by developing a clear map of your goals, vision, risk tolerance, and readiness to deliver.

TASKS WELL SUITED TO MI

- Simple, requires little or no context
- Involves finding patterns in data
- Characterized by large amounts of data—preferably labeled
- Occurs in a static environment or one with little uncertainty

TASKS POORLY SUITED TO MI

- Complex, requires contextual understanding
- Requires explaining patterns in data
- Little or no data to characterize the problem
- Occurs in a dynamic environment with lots of uncertainty

Figure 3.1 – A Recap of What We Have Learned

FIVE QUESTIONS TO ASK WHEN GETTING STARTED WITH MI

What are our goals?

The first question to consider when evaluating MI investment is the value you hope to capture through the investment. Your aspirations may be as simple as creating efficiencies in internal operations or as audacious as transforming your organization's brand or the industry within which it operates. Often, MI aspirations go beyond the immediate use cases under consideration to a broader vision for how the future will unfold and what your organization's role will be in that future state. With high expectations on the horizon, it's important to clearly articulate the return on investment expected from MI in the near and long term.

An idea to consider: Develop a 1-, 5-, and 10-year vision. When developing a multiyear plan, executives should start by evaluating how MI is already impacting their industry, and how this disruption might evolve in coming years. What value could MI enable in the industry that other organizations are not yet capturing? What capabilities are competitors delivering that the organization will need to respond to? What parts of the organization are most vulnerable? Next, using this picture, executives should think about their organization's role in these developments: Will the organization be a leader or a fast follower? What are the weaknesses and benefits of different approaches? Using this vision, executives can develop shorter-term, 1- to 3-year goals that will set them on the path to realizing goals for their organization's long-term trajectory.

How much risk are we willing to tolerate?

When considering investment in MI, it's important to remember that commercial applications of MI are still nascent and no MI product or platform is truly "off-the-shelf" (yet). All but the most basic applications of MI—for example, Robotic Process Automation (RPA)—come with a certain level of risk. An organization may, for example, set out to apply MI to a particular problem, only to discover the data the effort requires are harder to access and structure than originally estimated, lengthening timelines and complicating the validity of outcomes. Or, an organization may develop a client-facing MI application that makes some mistakes, frustrating or even embarrassing the organization. The good news is that mistakes and unexpected hurdles in MI are not insurmountable, and the return on investment for MI applications is often worth the risk. Still, it's important that executives develop a clear articulation of how much risk the organization is willing to tolerate when it comes to MI so that they can set clear parameters with partners and avoid frustrating organizational leadership.

An idea to consider: Start small, then scale. Executives at risk-averse organizations should consider starting small. RPA, for example, can be applied relatively easily and quickly to many clerical tasks—often in just a few weeks and with minimal disruption to staff activities. Relative to other MI engagements, the cost of RPA is small. Demonstrating the value of MI in this use case can provide an on-ramp to more experimental MI efforts. Alternatively, executives might consider a narrow MI application within a small part of the organization where staff are open to and savvy about new technology. Many organizations apply MI first in their IT organizations, particularly to address threat monitoring and response in cyber security. Those internal experts that go first in using MI can later become ambassadors and communicators to other staff as MI is scaled across the organization.

What is the state of our data assets?

MI experts are working hard to make machine learning less data intensive, but until then, MI requires a lot of labeled, well-organized data. This data intensiveness is a stumbling block for many organizations working to capture the value of MI. Even those organizations with huge data assets find data is usually harder to extract, structure, and organize than they expect. A 2016 survey found data scientists spend more than 80 percent of their time simply preparing data for use.[26] It is helpful to enter any MI engagement clear-eyed about the challenges to making use of data and consider strategies for addressing these challenges.

An idea to consider: Use partnerships. MI is an inherently collaborative field: It is evolving so quickly that no one organization can afford to operate in isolation, and even the largest organizations are looking to partners to make progress. Organizations lacking data assets can consider looking to startups, public institutions, and government organizations to harness others' data assets. The U.S. Government, for example, is working to make open data the default standard for its data assets, yielding rich information with wide business applications. Organizations may also consider funding academic research at institutions pursuing industry-relevant use cases.

An idea to consider: Structure infrastructure to capture the data you will need. Any executive who's attempted a data science or analytics effort knows how difficult it can be to structure, organize, and label data. The same is true for MI. Rather than trying to back into data that's disorganized or difficult to extract from organizational silos, consider thoughtfully collecting new data to power your MI efforts.

What MI talent assets do we have?

MI talent is scarce, and the battle for the top experts is fierce. Even the most prominent tech organizations—from Google and Amazon to Tesla and Uber—can rarely hold on to top talent for more than a year or two. Peter Lee, a VP in Microsoft's MI research department, recently compared the cost of acquiring a top MI researcher to that of signing a quarterback in the NFL.[27] This can be quite disheartening for many organizations that can't offer NFL salaries to talent who will only be with them for 12–24 months. But the scarcity of MI talent needn't be a barrier to entry for MI investment. Through partnerships with academics and investment in existing talent, most organizations can build the bench needed to put MI to work.

An idea to consider: Balance between borrowing, buying, and building MI talent. Partnering with research institutions and academic organizations enables organizations to "borrow" world-class MI talent. Firms should consider the assets they can offer to assist these organizations: investment dollars, data sets, and access to executives' industry-relevant insights. "Buying" MI talent is the right choice when an organization already has an MI application in mind that a hired expert can be put to work on right away. Executives should expect that it will be difficult to retain an MI expert long term and focus their efforts accordingly on interesting, time-constrained projects. Finally, organizations should invest in building MI talent in house. Most organizations will not have large cadres of in-house technical MI experts, but can upskill staff who have technical competencies. Those staff who don't should be familiar with what MI is, how it operates, and how to interact with it.

What are our values?

MI's portrayal in the media can conjure up fears of wide-scale automation, job loss, and even robot violence. While many of these fears aren't warranted (at least for now), MI does have the potential to significantly negatively impact our lives. From automating parts of our jobs to amplifying bias (if bad data are used) and disrupting our privacy (when our data are used in ways we don't expect), MI is not all good news. It is therefore critical that executives consider how MI will impact their clients, shareholders, employees, and the public from an ethical, equity, and privacy standpoint when beginning any MI initiative. Executives should preempt concerns by transparently and authoritatively communicating their organizational goals and values as they relate to MI. It is equally important to clearly explain any public-facing MI efforts to clients and investors. Clients may worry about their privacy, how their data will be used, and whether they are being treated fairly. It is important to convey your system of checks and balances for ensuring their concerns are addressed.

An idea to consider: Develop a list of questions you imagine you may receive about your use of MI, and agree on your answers. Solicit input from a range of colleagues to ensure a diversity of perspectives are incorporated into these questions. Speak with executives, departmental leads, legal experts, and human capital resources about the questions and concerns you expect your MI efforts may generate. For example, "Will this technology replace my job?"; "How are my data being used?"; or, "What are you doing to protect my privacy?" Once you are clear on the questions that may arise, develop answers that you and fellow leaders can agree on. This is not simply an exercise in risk avoidance—it will also help your organization define the ethical principles and values that will guide your use of MI.

When considering where to invest in MI, it's helpful to consider the different types of MI you might want to implement, and compare them against your goals, risk tolerance, data assets, and talent to identify gaps and opportunities.

	SIMPLE TASK EXECUTION	PATTERN RECOGNITION	CONTEXTUAL REASONING
GOALS	• Automate simple processes • Free staff from rote work • Create efficiencies/save money • Transform internal operations	• Augment staff in low- to medium-complexity tasks • Generate new business/ operational insights • Create a new client-facing capability • Enhance corporate brand	• Create an industry-leading technology/tool • Create an enduring competitive advantage • Enhance corporate brand
RISK TOLERANCE	Low	Moderate–High	High
TIME INVESTMENT	SHORT TERM 2–6 weeks	MID TERM 6–12 months	LONG TERM 12–36 months
DATA ASSET REQUIREMENTS	• Few data assets • Data assets siloed/unorganized	• Massive data assets or access to data • Data labeled, organized	• Variable data assets—generally large
TALENT REQUIREMENTS	• New in-house MI talent generally not required—implementation typically through vendors, consultants	• MI talent required, either in-house or accessed via partnerships (e.g., with academic institutions)	• Leading MI talent required, typically accessed through partnerships (e.g., in engagements with large tech vendors)
SOLUTIONS	• Robotic process automation • Expert systems • Simple automation	• Core machine learning software • Computer vision • Natural language processing • Cognitive robotics	• Contextual machine learning ("semantic" or "cognitive computing")
EXAMPLES	• Automate expense reporting • Automate call center data entry	• Create a virtual agent to interact with clients in call centers	• Develop autonomous vehicles

Figure 3.2 – A Framework for Deciding How to Invest in MI

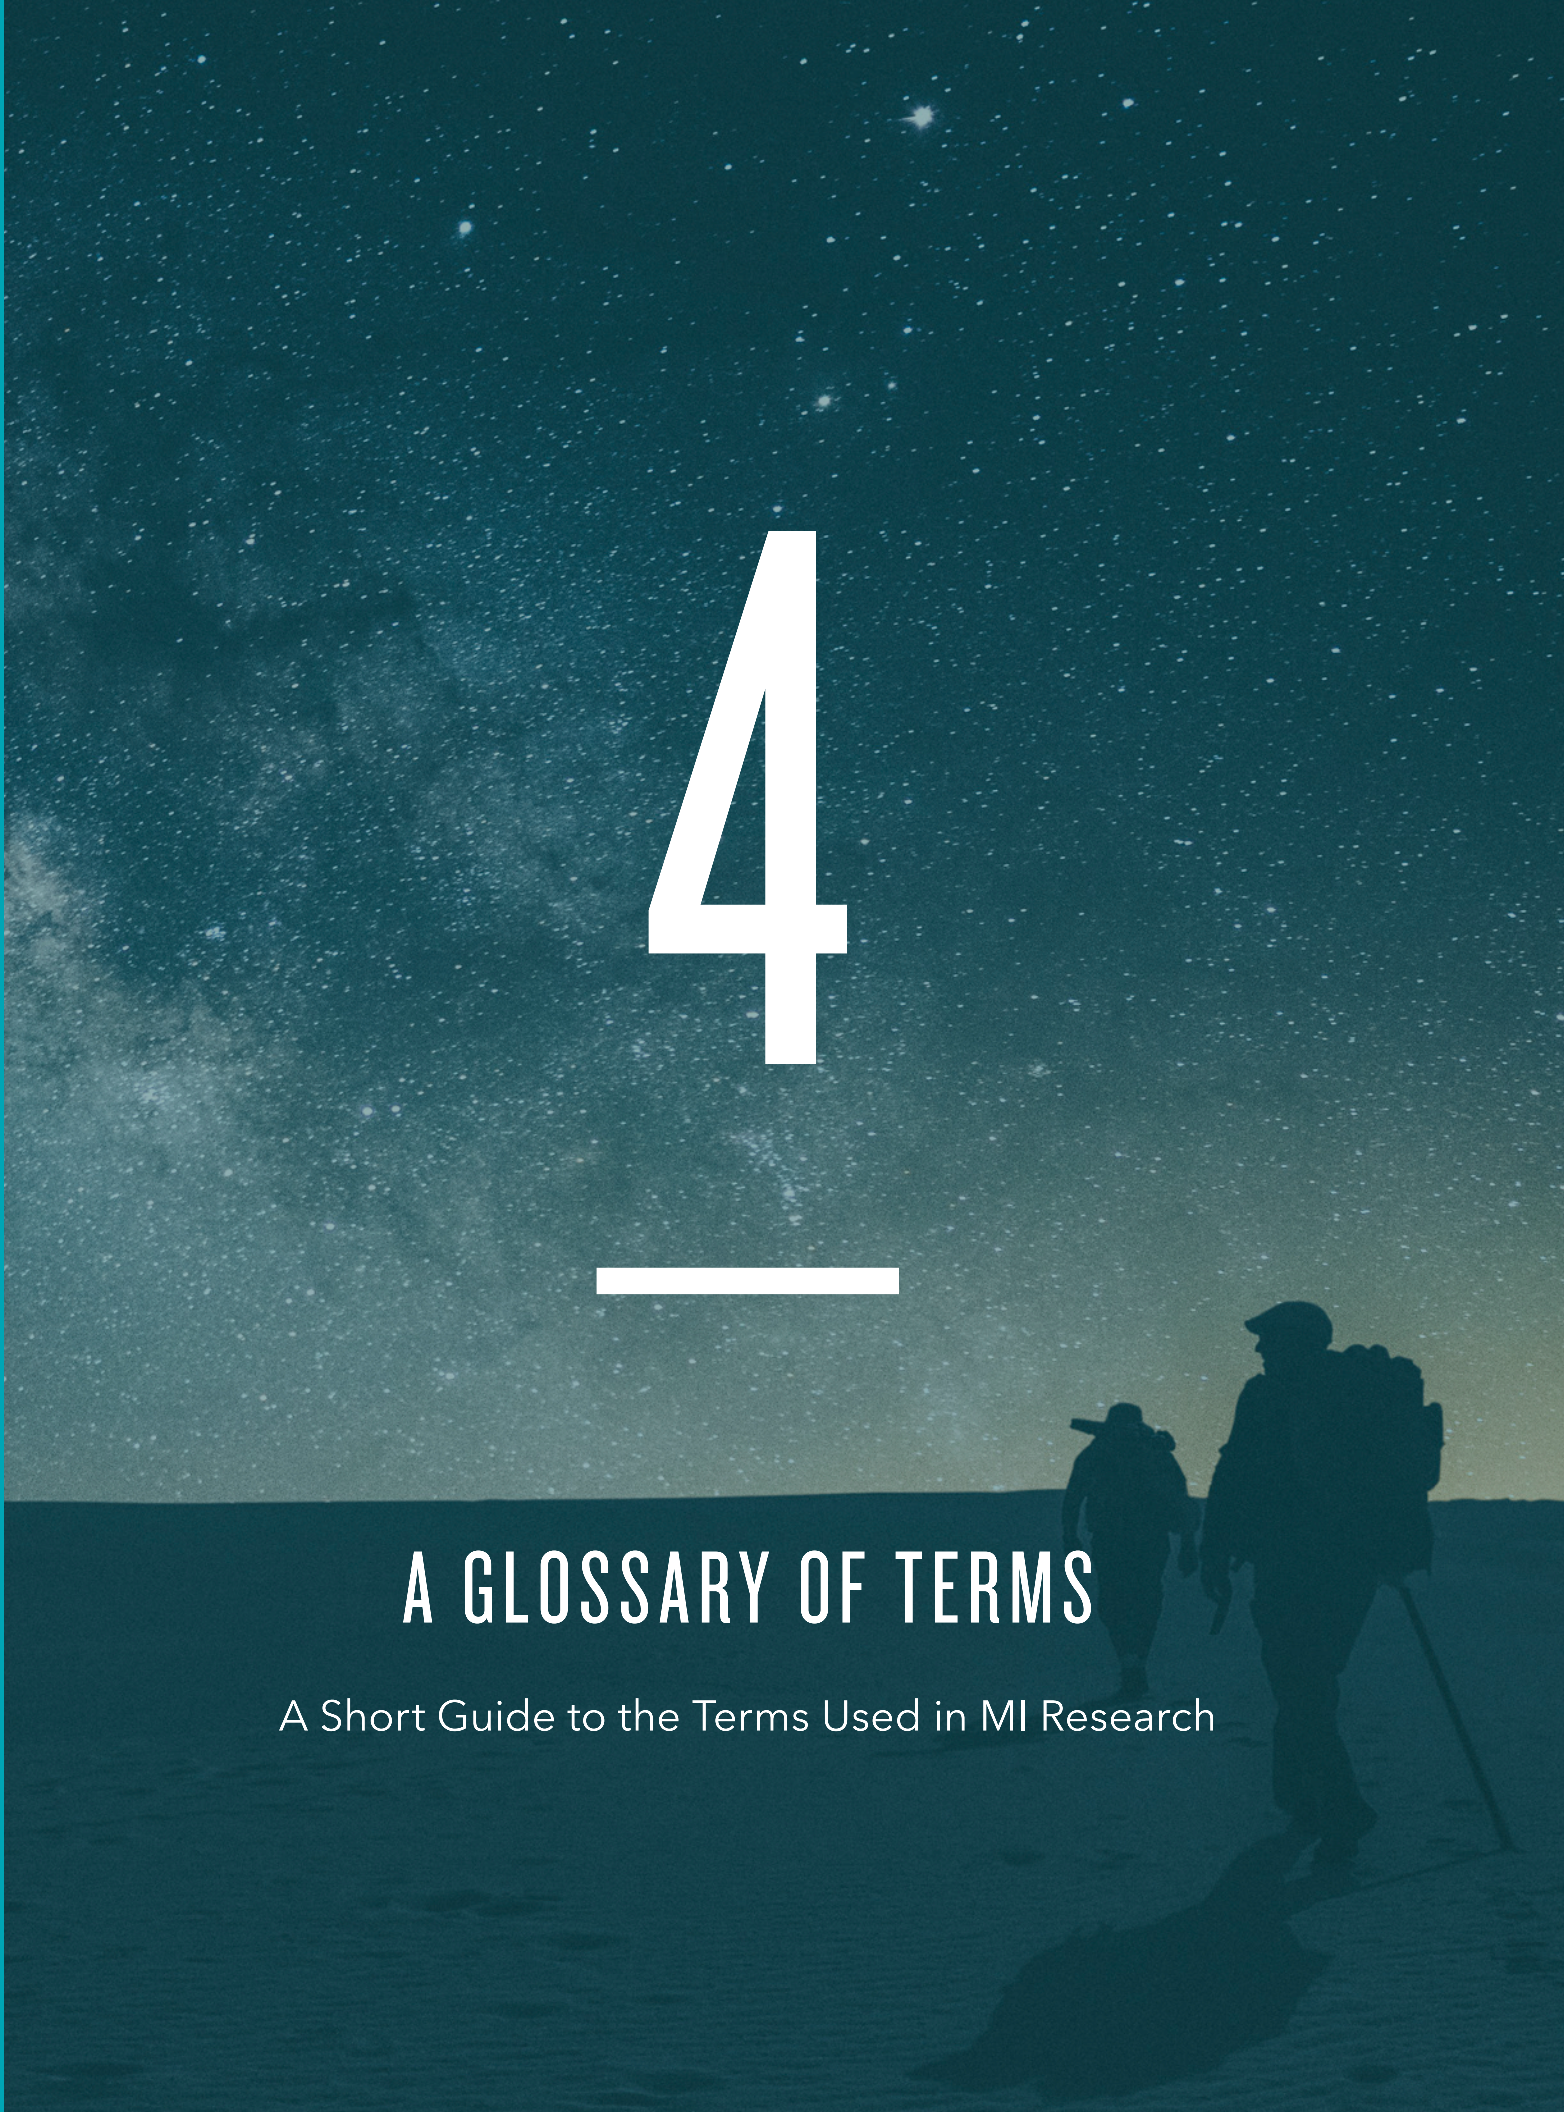

4

A GLOSSARY OF TERMS

A Short Guide to the Terms Used in MI Research

Artificial Intelligence: A term interchangeable with "MI." AI was the original term often used to classify the technology now often dubbed MI. Some organizations prefer the term MI because it is less likely to evoke ties to science fiction.

Computer Vision: MI systems that can perceive, learn from, abstract, and act on data in images and video to perform visual tasks (e.g., navigating an environment, labeling a photo). Computer vision systems typically consist of enabling technologies and techniques—sensors, cameras, and image processors—and machine learning algorithms. Applications of computer vision include object identification (i.e., in photos and videos), image and video tagging, biometrics for security and other purposes, sentiment analysis, facial recognition, and scene analysis.

Expert Systems: An early form of MI in which researchers codified human experts' knowledge as sets of logic-based rules and programmed these rules into computers, enabling them to perform straightforward tasks with precision and scale. These systems lost popularity during the 1980s due to their rigidity and poor handling of changing or ambiguous circumstances.

Five Tribes of Machine Learning: A term coined by contemporary MI expert Pedro Domingos for disparate groups of MI researchers who favor competing approaches to designing machine learning algorithms.

Labeled or "Tagged" Data: Data which have been given a label, (e.g., "there is a cat in this photo") typically by humans, that corresponds to an outcome variable. Labels make these data useful for training supervised machine learning algorithms. Labeling data can be time consuming, and data that come pre-labeled are often expensive.

Machine Intelligence: A field concerned with producing machines able to autonomously perform tasks that would normally require human intelligence by giving them the ability to perceive, learn from, abstract, and act using data.

Machine Learning: A technique to achieve MI in which computers are given the ability to learn without being explicitly programmed using algorithms that can parse, learn from, and make determinations based on data.

Natural Language Processing and Generation (NLP, NLG): MI systems that can perceive, learn from, abstract, and act on data in text and audio in natural language (i.e., conversational speech that does not adhere to rigid rules of syntax, grammar, etc.) to perform language-related tasks (e.g., generate text, translate audio, etc.). Applications include machine translation, speech recognition, language detection, sentiment analysis, document analysis, virtual assistants, report generation, news article writing, and insight summarization.

Robotic Process Automation: A recently developed form of MI that enables machines to perform simple, rules-based, high-volume tasks—often clerical or "back office." Software "bots" observe clerical staff performing these tasks and develop rules for repeating them, enabling the execution of simple tasks at scale and with precision.

Supervised Machine Learning: An approach to training machine learning algorithms in which the algorithms are given input variables and output (or "tagged" variables) from which to learn patterns and make inferences. Supervised machine learning is useful for classification (assigning items to categories) and regression (finding values, such as price or weight) problems.

Unsupervised Machine Learning: An approach to training machine learning algorithms in which the algorithms are given only input variables from which they uncover patterns and make inferences on their own. Unsupervised machine learning is good for clustering problems (grouping people, behaviors, types, etc.) and classification problems (finding rules in data).

Unlabeled or "Untagged" Data: Data which have not been given a label to make them useful for training supervised machine learning algorithms. The vast majority of data in the world are unlabeled, making them relatively less expensive to use or obtain than tagged or labeled data.

CITATIONS

1. McCarthy, Minsky, Rochester, and Shannon, "A Proposal for the Dartmouth Summer Research Project on AI," 1955
2. Arthur Samuels, "Some Studies in Machine Learning Using the Game of Checkers," 1956
3. Peter Domingos, "The Master Algorithm," 2015
4. Hans. J. Berliner, "Backgammon Computer Program Beats World Champion," http://www.bkgm.com/articles/Berliner/BackgammonProgramBeatsWorldChamp/, 1980
5. IDC, "Autonomous Robot and Autonomous Vehicle," http://www.idc-online.com/technical_references/pdfs/electronic_engineering/Autonomous_Robot_and_Autonomous_Vehicle.pdf
6. Management Information Systems Research Center at the University of Minnesota, "Early Expert Systems: Where are they Now?," https://www.jstor.org/stable/249711?seq=18#page_scan_tab_contents, March 1995
7. The World Economic Forum, "The Fourth Industrial Revolution," https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/, January 2016
8. The Economist, "Why Artificial Intelligence is Enjoying a Renaissance," http://www.economist.com/blogs/economist-explains/2016/07/economist-explains-11, July 2016
9. The New York Times, "Artificial Intelligence Software is Booming- But Why Now?," https://www.nytimes.com/2016/09/19/technology/artificial-intelligence-software-is-booming-but-why-now.html, September 2016
10. MIT Technology Review, "MI Winter Isn't Coming," https://www.technologyreview.com/s/603062/ai-winter-isnt-coming/, December 2016
11. Pew Research Center, "The Evolution of Technology Adoption," http://www.pewresearch.org/fact-tank/2017/01/12/evolution-of-technology/, January 2017
12. Venture Beat, "Google's Kaggle Now has 1 Million Developers Signed up for AI Competitions," https://venturebeat.com/2017/06/07/googles-kaggle-now-has-1-million-developers-signed-up-for-ai-competitions/, June 2017
13. Nvidia, Nvidia Announces Financial Results for First Quarter 2018," http://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2018, May 2017
14. WIRED, "The Man Behind the Google Brain: Andrew Ng and the Quest for the New AI," https://www.wired.com/2013/05/neuro-artificial-intelligence/, May 2013
15. Harvard Business Review, Andrew Ng, "What AI Can and Can't Do," https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now, December 2016
16. Venture Beat, "AI is Still Several Breakthroughs Away from Reality," https://venturebeat.com/2017/06/23/ai-is-still-several-breakthroughs-away-from-reality/, June 2017
17. DARPA; "Three Waves of AI," http://www.darpa.mil/attachments/AIFull.pdf, February 2017
18. Harvard Business Review, Andrew Ng, "What AI Can and Can't Do," https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now, December 2016
19. Driverless Future, "Autonomous Car Forecasts," http://www.driverless-future.com/?page_id=384, June 2017
20. PRI, "Driving in Some Cities is a Blood Sport. Can Autonomous Cars Compete?," https://www.pri.org/stories/2017-05-12/driving-some-cities-blood-sport-can-autonomous-cars-compete, May 2017
21. Tech Republic, "Is Data Labeling the Blue Collar Job of the AI Era?", http://www.techrepublic.com/article/is-data-labeling-the-new-blue-collar-job-of-the-ai-era/, March 2016
22. International Business Times, Samantha Olson, http://www.ibtimes.com/artificial-intelligence-robots-why-human-baby-brains-are-smarter-ai-2446588, November 2016
23. The Verge, "These Are Three of the Biggest Problems Facing Today's AI," https://www.theverge.com/2016/10/10/13224930/ai-deep-learning-limitations-drawbacks, October 2016
24. The Guardian, "Discrimination by Algorithm: Scientists Devise Test to Detect AI Bias," https://www.theguardian.com/technology/2016/dec/19/discrimination-by-algorithm-scientists-devise-test-to-detect-ai-bias, December 2016
25. International Conference on Machine Learning, https://sites.google.com/site/dataefficientml/, May 2016
26. Forbes: Cleaning Big Data: The most Time-Consuming, ELast Enjoyable Data Science Task, https://www.forbes.com/sites/gilpress/2016/03/23/data-preparation-most-time-consuming-least-enjoyable-data-science-task-survey-says/#704b401f6f63, March 2016
27. WIRED, "The Battle for AI Talent Only Gets Tougher from Here," https://www.wired.com/2017/03/intel-just-jumped-fierce-competition-ai-talent/, March 2017

MEET THE AUTHORS

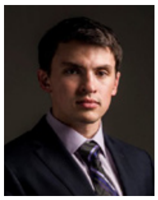

J.D. DULNY

Dulny_joseph@bah.com

301-489-3357

EMMA KINNUCAN

Kinnucan_emma@bah.com

617-512-2846

JOSH ELLIOT

Elliot_josh@bah.com

202-406-1949

STEVE MILLS

Mills_steven@bah.com

240-456-7942

DREW FARRIS

Farris_andrew@bah.com

410-865-3790

JOSHUA SULLIVAN

Sullivan_joshua@bah.com

301-543-4611

The developments occurring in Machine Intelligence (MI) are undeniably among the most exciting technological advancements of our era. MI will, over time, have broad impacts on how we engage with our work, communities, and one another. To make the most out of these changes, we must remain pragmatic about MI's capabilities and limitations. We hope that the MI Primer has enabled you to better understand these considerations, and to imagine how you will use MI to create a more equitable, verdant, and meaningful future. We encourage you to share these ideas with us. Join us in the conversation at machineintelligence@bah.com.

ABOUT BOOZ ALLEN

Booz Allen Hamilton has been at the forefront of strategy and technology for more than 100 years. Today, the firm provides management and technology consulting and engineering services to leading *Fortune 500* corporations, governments, and not-for-profits across the globe. Booz Allen partners with public and private sector clients to solve their most difficult challenges through a combination of consulting, analytics, mission operations, technology, systems delivery, cybersecurity, engineering, and innovation expertise.

With international headquarters in McLean, Virginia, the firm employs more than 22,600 people globally and had revenue of $5.41 billion for the 12 months ended March 31, 2016. To learn more, visit BoozAllen.com.
(NYSE: BAH)

© 2017 Booz Allen Hamilton, Inc.