Article from

Predictive Analytics and Futurism

December 2016
Issue 14

Five Myths and Facts about Artificial Intelligence

By Dr. Anand S. Rao

Artificial intelligence (AI) is splashed throughout the headlines these days. AI recently beat the human Go champion, autonomous cars can drive themselves, social bots can mimic human interactions and converse with others in social networks, and AI totally manages some hedge funds.[1]

However, countering all of this exciting news is fear about the potential for AI to take the place of and even pose an existential threat to humans. In a seminal paper published in 2013, Frey and Osborne[2] estimate that nearly 47 percent of U.S. employment is susceptible to computerization and automation. Moreover, insurance agents, sales, underwriters, and claims adjusters are some of the jobs that have a high potential for automation.

AI will play—and arguably is already playing—a prominent role in insurance.[3] In personal insurance,[4] commercial insurance,[5] and life insurance,[6] AI has already had an impact on how insurers 1) underwrite, price, market, and manage coverage, 2) target and manage their customers, and 3) distribute products. And, while AI will automate within 10 years some of the routine tasks insurance agents, underwriters, actuaries,[7] and claims adjusters perform today, we believe that it offers a great opportunity for insurers to transform these roles. Far from just automating certain tasks, AI will augment human expertise with more sophisticated data, tools and algorithms. This will enable insurers to make faster, better and cheaper strategic and operational decisions. At the same time, human experts embody a wealth of knowledge of and experience with the economy, markets, customers' needs and products and can play a critical role in training, testing and refining AI.

However, to better explain what AI may mean for the future of the industry, we need to debunk some of the myths about AI and explain realities.

MYTH #1: ARTIFICIAL INTELLIGENCE IS A DISTINCT AND MONOLITHIC AREA OF STUDY.

The term "artificial intelligence" was coined only 60 years ago. It has been subject to periods of excitement (often called "AI spring") followed by ones of despair (often called "AI winter"). We are currently in a new AI spring, in which hype about the "triumph of AI" and the "fall of yet another human activity to AI" is incessant. While some of the latter claims are realistic and do reflect significant advancements in AI, most are hyperbole. The media, in its quest to simplify things for the layman, tends to equate AI with a specific subfield of AI to the complete exclusion of everything else.

For example, in light of the important advances taking place in machine learning and deep learning, many articles that address advances in these areas equate AI only with those two areas of AI. Similarly, an article on robotics will claim that AI is simply robotics or robotic process automation or conversational agents. This gives the reader—including business executives—a very limited perspective and typically leaves him wondering just what AI is and what it is not.

FACT #1: ARTIFICIAL INTELLIGENCE IS AN INTER-DISCIPLINARY AREA WITH MANY DISTINCT SUBFIELDS.

In the words of John McCarthy, one of the founding fathers of artificial intelligence, "AI is the science and engineering of making computers intelligent." Given the lack of consensus on what is "intelligent," AI researchers have pursued a number of areas of human intelligence, including pattern recognition, understanding, learning, problem solving, reasoning, and decision making. With a view towards unifying these disparate areas, Russell and Norvig,[8] in their classic text define AI as "the study and design of intelligent agents where an intelligent agent is a system that perceives its environment and takes actions which maximizes its chances of success."

FIGURE 1: ILLUSTRATIVE SUBFIELDS OF ARTIFICIAL INTELLIGENCE

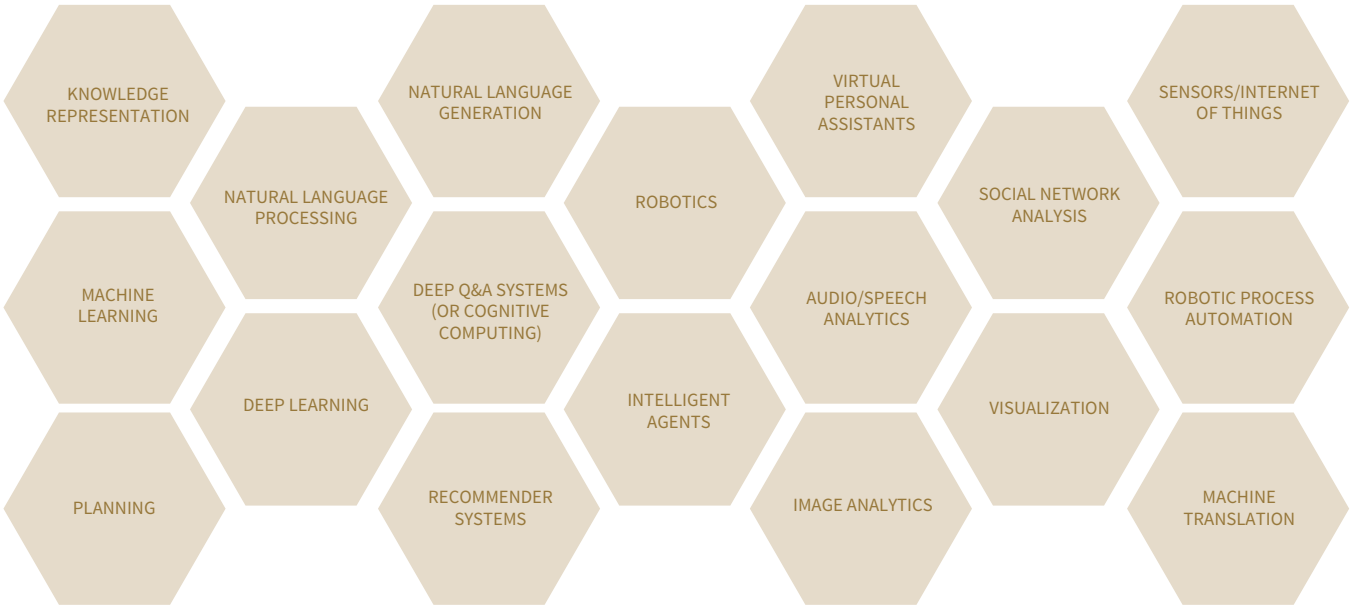

As of yet, there is no single unifying theory or practical solution to implement intelligent agents that can perceive and act in all environments as well as humans. AI researchers have typically adopted a "divide-and-conquer" approach that incorporates techniques from a variety of scientific disciplines, including computer science, statistics, mathematics, physics, biology, philosophy and logic. In addition, whenever any so called "intelligent" problem is solved by AI, it ceases to be called AI. For example, handwriting recognition, speech synthesis, or voice recognition were all considered AI in the 1980s, but as they became part of everyday solutions (e.g., handwriting recognition is now available on tablets, speech synthesis and voice recognition are now available on smartphones), they moved out of the AI realm.

Figure 1 shows a number of subfields that come under the banner of AI. Note that these topics are not mutually exclusive and collectively exhaustive (MECE), but are researchers and businesses are actively pursuing them; moreover, they are mature enough to offer viable solutions to some practical problems.

MYTH #2: ALL TYPES OF PROBLEMS CAN BE SOLVED BY A SINGLE AI SOLUTION.

A more serious consequence of Myth #1 is that the media, business executives, AI solution providers, and in some cases even AI researchers truly believe that a single AI subfield or a solution based on that subfield can solve ***any*** business problem. Part of the myth is due to a genuine ignorance of all AI's different subfields and what they can or cannot accomplish. Also, given the amount of venture capital and corporate funding that is going into AI, there is an increasing tendency to make claims that a particular AI solution is relevant to "all" business problems. This could have negative consequences for insurance companies and other organizations that are trying to pilot AI solutions. High hopes that AI solutions can be panaceas may well lead to disillusionment when producers and users don't change the world in one fell swoop. This could have ripple effects, such as failing start-ups and decreasing funds for AI (which, in turn, could lead to the next AI winter).

FACT #2: DIFFERENT TYPES OF PROBLEMS REQUIRE DIFFERENT TYPES OF AI TECHNIQUES.

The repertoire of human activity and intelligence is so vast that claiming a single AI solution—or even a combination of solutions—can match, much less surpass them, seriously underestimates the complexity of human cognition. Figure 2 outlines the key capabilities necessary for solving business problems. In spite of all the recent advances in AI, the world is still nowhere close to developing a unified theory or a single solution that incorporates all of the capabilities we outline. While the particular business problem a given researcher or company is trying to solve may not require all of those capabilities, it is very likely to require at least some of them. Even if a business problem is narrow enough to require limited capabilities, an AI solution may not be able to fully address it. For example, a natural language question answering system may be suitable for holding a conversation with the customer in order to solve a service issue, but may not be able to simulate multiple strategic options and thereby help management identify an optimum business strategy.

FIGURE 2: KEY BUSINESS CAPABILITIES

SENSING, QUERYING & CONVERSING

Ability to search or query structured databases, text, audio, image, and video data. Ability to communicate and converse with the machine in natural language.

DESCRIBING, CLASSIFYING, UNDERSTANDING & VISUALIZING

Ability to describe, classify, understand and visualize structured data, natural language sentences or documents, audio files, image and video data.

DIAGNOSING, DECIDING & REASONING

Diagnosing root causes, deducing rules from existing data, and reasoning what happened and what is likely to happen.

TRENDING, FORECASTING, PROJECTING & PREDICTING

Ability to identify trends from past data, forecast the future from the past data, project into the future across multiple scenarios and make predictions.

SIMULATING, ACTING, LEARNING & ADAPTING

Ability to simulate scenarios, autonomously act, learn from data and over time adapt to changing environmental conditions.

RECOGNIZING, SENSING & RECOMMENDING

Recognizing features from datasets (structured, unstructured—text, audio, image, video), sensing the environment through sight and speech (smell, taste, touch in the future) and recommending actions to be taken.

MYTH #3: MACHINE LEARNING AUTOMATICALLY LEARNS FROM DATA WITHOUT ANY HUMAN INTERVENTION.

In their enthusiasm to promote machine learning (and deep learning in particular), as well as to contrast it with other techniques, researchers, companies, and media often portray machine learning as something magical; if one feeds data into the system, then it will learn all the patterns and be able to answer any future queries. This perception is not only inaccurate, but also can be counter-productive for companies adopting AI. We have seen companies suddenly shift focus from building robust data collection platforms to trying to build machine learning or deep learning systems. In our view, organizations should have a basic level of "data hygiene" and the foundational AI skills and culture before embarking on advanced analytics, machine learning and deep learning.

FACT #3: MACHINE LEARNING REQUIRES A LABORIOUS PROCESS OF ACQUIRING AND CLEANSING LARGE AMOUNTS OF DATA, AND SELECTING, TRAINING AND GUIDING THE ALGORITHM.

Herbert Simon, a leading AI researcher, defined machine learning as "computers that automatically improve their performance through experience." The three key phrases here are "automatically," "experience" and "performance." Unlike other types of AI that are programmed, machine learning learns patterns from the data it is given in order to improve its performance according to a specific criteria. Once the machine learning system has been "trained" to recognize certain patterns, it can be provided with new data that it should be able to recognize or classify.

Machine learning systems vary by their representation (e.g., neural networks, rules), how they navigate through the search space of all possible solutions, and optimization criteria. Most machine learning systems require large quantities of training data. In addition, based on the type of learning (e.g., supervised learning) many of them also require humans to provide labeled data (i.e., training data requires the "right" answers). Given the number of different machine learning techniques, human expertise and judgement is critical in selecting the right techniques and fine-tuning relevant parameters. Deep learning, a specific type of machine learning based on multi-layered neural networks, has performed well on a number of image, video, audio and natural language processing tasks. However, the success of deep learning algorithms has depended largely on human ingenuity, notably fine tuning and modularizing the algorithms and collecting and preparing the right labeled data to train the system. In one particular implementation,[8] humans spend more than 70,000 hours collecting, analyzing, and labelling more than 328,000 photos.

MYTH #4: IN THE WORLD OF AI, THERE IS NO ROLE FOR HUMAN EXPERTISE AND INTELLIGENCE.

Another common myth is that AI will replace humans and take away their jobs, leaving no role for human expertise and intelligence in the future. Some data science and machine learning advocates claim that all they need is data, not human domain

expertise or intelligence. While it is true data scientists who have no domain expertise in a given problem area have won data science competitions, humans have had to frame problems, state hypotheses, set criteria for success, and finally evaluate the solutions.

FACT #4: AUGMENTING HUMAN INTELLIGENCE WITH ARTIFICIAL INTELLIGENCE WILL LEAD TO BETTER RESULTS AND DECISIONS THAN EITHER ONE COULD ACHIEVE ON ITS OWN.

Artificial Intelligence has matured to the point that AI software can perform some common tasks and decisions (e.g., recognizing faces or landmarks in a photograph and tagging them appropriately). However, there are a number of more complex tasks and decisions in business and everyday life that still require human skills, expertise and ingenuity.

Accordingly, AI can play an important role in helping humans make better decisions. But, for example, evaluating a number of strategic options for an insurance company entering the "connected home" insurance market is not something that AI software can do autonomously today. However, we can build sophisticated AI models that capture the dynamics of consumer adoption, competitive dynamics and technology trends to postulate strategic options that AI software can evaluate. In such systems, based on criteria humans provide, AI software can evaluate and optimize a multitude of strategic scenarios—something that is impossible for human brains to do. As a result, augmented intelligence, where AI systems are initially based on human knowledge and then inform humans in an endless feedback loop, will become the norm. (In fact, there are a number of evolving AI systems that approach this level of sophistication today.)

MYTH #5: AI POSES A LOOMING, EXISTENTIAL THREAT TO HUMANITY.

This myth would have not been part of this list unless some well-known academics and business leaders publically announced that it's a potentially serious issue. In his book *On Superintelligence*, Nick Bostrom[9] argues that an AI system which is able to create better and better versions of itself could very quickly surpass current levels of human intelligence. If this were to occur, he claims that we will be unable to predict or guarantee the values of this superintelligence. Stephen Hawking, Elon Musk, and Bill Gates[10] also have recently echoed these concerns.

FACT #5: SUPERINTELLIGENCE IS NOT A TECHNICAL REALITY. AT LEAST NOT YET.

Superintelligence is a fascinating concept, but is decades away from potential realization. Almost all of AI today fails to approach even general intelligence and is good for addressing only very specific problems (see Myth #2). Even the field of Artificial General Intelligence, which aims to build general purpose intelligence, is still in its infancy. Andrew Ng,[11] a leading AI researcher, puts it succinctly: "*The reason I say that I don't worry about AI turning evil is the same reason I don't worry about overpopulation on Mars. Hundreds of years from now I hope we've colonized Mars. But we've never set foot on the planet so how can we productively worry about this problem now?*" Nevertheless, because of concerns that AI may go rogue at some point in the future, there are a number of researchers who are addressing how to build AI systems with sufficient safeguards and ethics.

CONCLUSION

As insurance companies evaluate the use of artificial intelligence in their organizations, they need to understand the true potential and limitations of today's AI technology. While AI has matured substantially over the past couple of decades, it is still a long way from being a "silver bullet" for solving most problems, and it is incapable of completely replacing human labor and decisions. As a result, any deployment of AI needs to start with an understanding of the types of business problems AI can actually address and the best AI subfield(s) and techniques for addressing them. ■

Anand S. Rao, PhD, is Partner, Innovation Lead, PwC Data and Analytics at PricewaterhouseCoopers in Boston, Ma. He can be reached at *anand.s.rao@pwc.com*

ENDNOTES

1 Artificial Intelligence: March of the Machines. The Economist, June 25, 2016.

2 The Future of Employment: How Susceptible are Jobs to Computerization? by Frey, and Osborne, September 17, 2013.

3 AI in Insurance: Hype or Reality. Top Issues, Insurance, PwC White Paper, March 2016.

4 Artificial Intelligence in Personal Insurance: Chatbot agents, RPA 'No Shoring' AdminBots and more by Anand Rao. Carrier Management, August 3, 2016.

5 Artificial Intelligence in Commercial Insurance by Anand Rao and Joseph Calandro. Carrier Management, August 21, 2016.

6 Why Artificial Intelligence will not replace Insurance Agents by Anand Rao. Independent Agent, July 20, 2016.

7 2036: An Actuarial Odyssey with AI by Dodzi Attimu and Bryon Robidoux. Predictive Analytics and Futurism, published by Society of Actuaries, Predictive Analytics and Futurism Section, Issue 13, July 2016,

8 Artificial Intelligence: A Modern Approach by Stuart Russell and Peter Norvig. Prentice Hall, 1994.

9 Microsoft COCO: Common Objects in Context. Lin et al., *http://arxiv.org/pdf/1405.0312v3.pdf*, February 21, 2015.

10 Superintelligence: Paths, Dangers, Strategies. Nick Bostrom, Oxford University Press, Oxford, July 3, 2014.

11 Stephen Hawking, Elon Musk, and Bill Gates Warn About Artificial Intelligence by Michael Sainato. Observer, August 19, 2015.

12 Andrew Ng: Why 'Deep Learning' is a mandate for Humans, not just machines. Caleb Garling, Wired.