



The Elusive Promise of AI—a Second Look

by Jeff Riley

Editor's Introduction

A 2006 Ubiquity article titled “[The Elusive Promise of AI](#)” contended that the field of artificial intelligence (AI) promised much but had not yet delivered on its promises. This follow-up article reviews some of the more significant events and progress in AI over the intervening decade-and-a-half since the original article, describes roughly where we are today, and speculates as to what might be ahead of us.

The Elusive Promise of AI—a Second Look

by Jeff Riley

In 2006 I wrote an article for Ubiquity titled “[The Elusive Promise of AI.](#)” The thrust of that article was the field of artificial intelligence (AI) had promised much from its beginnings in the middle-part of last century, but it had not yet delivered machines that think; machines that display human intelligence. In the article I observed learning “from scratch” is an incredibly difficult problem, and surmised that it would be a very long time before humans are able to construct machines that can reason and learn from scratch. We have come a long way in the past 14 years, but we are still a very long way away from constructing machines that can think, machines that can reason and learn from scratch. The current state of AI is still characterized by humans endowing machines with prior knowledge in one way or another—but that may not necessarily be an impediment to constructing machines that are actually intelligent.

In this new article I will review some of the more significant events and progress in AI over the past decade-and-a-half, describe roughly where we are today, and speculate as to what might be ahead of us.

What Has Happened in the Past Decade-and-a-half?

AI has experienced probably the most rapid growth in its history over the past 15 years or so, the two biggest drivers of which were continued growth in computing performance (e.g. supercomputer performance has increased by three orders of magnitude since 2005), and the explosion of data available to train AI algorithms, itself driven by the growth in internet usage and the Internet of Things (IoT). Current estimates put the amount of data stored by Google, Facebook, Microsoft and Amazon alone at approximately 1,200 petabytes—that’s 1.3 million million million bytes.

Following is a brief account of some significant events in AI since the appearance of my 2006 article.

2006 Fei-Fei Li, then at Princeton University, formulates an idea that leads to ImageNet, a database of some 14 million images annotated with labels. The existence of a large database of labeled images available as training examples for machine learning (ML) algorithms spurs progress in AI-aided computer vision and the development of ML image recognition algorithms. The ImageNet project runs an annual software contest, the “ImageNet Large Scale Visual Recognition Challenge,” for which competitors develop software to (try to) correctly classify and detect objects and scenes.



2009 Google develops self-driving automobile technology at its X lab, leading to an investment of billions of dollars and the creation of the Waymo division in 2016. Google began limited test-driving of driverless cars on public roads in late 2017, leading industry observers to believe that driverless vehicles capable of safely navigating any environment or road condition is a real possibility in the very near future.

2009 Deep learning gets a boost. Most deep learning models are based on artificial neural networks (ANN), specifically deep neural networks (DNN) and convolutional neural networks (CNN), and while ANNs, DNNs and CNNs have been around for decades, the 2009 NIPS Workshop on Deep Learning for Speech Recognition leads to the understanding that very large training data sets allow networks to learn more rapidly and with greater accuracy (i.e. with lower error rates). ImageNet is presented at the NIPS conference, and it is generally recognized that deep learning techniques are faster and more computationally efficient than many other ML techniques.

2011 Apple releases Siri on the iPhone. Siri is a virtual assistant capable of a surprisingly wide array of tasks on an IOS device. Not long after, in 2012, Google follows with Google Now (to become Google Assistant in 2016), and in 2014 Amazon releases Amazon Echo—a smart speaker with an integrated virtual assistant known as Alexa. With about 100 million smart speakers shipped worldwide in 2019, and the market growing at roughly 45 percent, affordable AI for the home is fast becoming ubiquitous.

2011 IBM’s Watson supercomputer challenges two all-time champions of the “Jeopardy!” television quiz show, and wins. Watson was constructed as part of IBM’s DeepQA project and incorporated a variety of question answering technologies including parsing, question classification, question decomposition, automatic source acquisition and evaluation, entity and relation detection, logical form generation, and knowledge representation and reasoning.

2012 Computing speed, particularly for GPUs, has increased significantly, and with the increase in computing performance the deep learning revolution gains momentum. A deep CNN for image classification achieves superhuman performance, and AlexNet (a CNN) wins the [ImageNet Large Scale Visual Recognition Challenge](#). AlexNet is considered one of the most influential papers published in computer vision.

2012 Google creates an ANN with 1 billion connections that learns to differentiate images of cats, human faces, and other human body parts—unsupervised and in a collection of unlabeled and unclassified images. The ANN used is an order of magnitude bigger (in terms of number of neurons and connections) than the state-of-the-art at the time. While learning what cats look like may sound frivolous, the real takeaways here are that AI can, without too much human direction (humans designed the network and the training algorithm, so there was some human direction), discover concepts previously unknown to it, and that the computer hardware to develop very large, complex ANNs is with us already.



2014 Ian Goodfellow and his colleagues develop generative adversarial networks (GAN). GANs use two neural networks in competition with each other to generate synthetic data that mimics real data. GANs are generative models where one model (the generative network) generates new data, and another (the discriminative network) evaluates them against true data. The goal of the generative network is to fool the discriminative network into evaluating synthesized data as true data. This adversarial game allows GANs to learn effectively unsupervised from the chosen training data. GANs are used widely in image generation, video generation and voice generation, and underpin DeepFakes.

2015 Skype integrates real-time voice translation (Skype Translator) into Skype for Windows desktop. People who don't speak a common language can now carry on normal conversations while Skype translates each to the other in real-time. While only seven languages were supported at launch, support for more languages is being added over time. Sci-Fi's Universal Translator has almost come of age.

2016 AlphaGo, an AI system developed by DeepMind to play the ancient Eastern game of Go, defeats the reigning European Go champion, Fan Hui (representing France, originally from China), 5-0 in a five-game competition, and world number two player, 18-time world champion Lee Sedol (Korea), 4-1. The ANN inside AlphaGo was trained to play Go using a large database of Go moves from expert human players then, by playing against slightly different versions of itself it was able to generate new moves that were later used to train the ANN with an even larger database of Go moves.

2018 An AI program developed by Chinese eCommerce company Alibaba becomes the first computer program to outperform humans on a Stanford University reading comprehension test. The Stanford Question Answering Dataset (SQuAD) was developed by Stanford University to test AI systems' ability to process and comprehend large amounts of information. We could soon have our newspapers, books, and business reports summarized for us by computers, presenting our daily doses of information to us in easily digestible chunks.

2018 Google releases as open source BERT (Bidirectional Encoder Representations from Transformers), a method of training natural language processors and an integral part of its own search technology. BERT provides a pre-trained starting point for deep neural networks. BERT's transformer network learns relationships between sentences, enabling it to better handle language nuances, ambiguity, and context.

2019 A year after Alimama, Alibaba's digital marketing subsidiary, claims that its AI-based copywriting tool passed the Turing Test, Springer publishes its first machine-generated book. There was human generated content in the book—only the introductions, tables of contents and references were machine-generated—but the release of the book served to demonstrate how far computer-generated content had come. Leading U.K. news agency Press Association, for example, produces 30,000 local news stories a month using AI.

2020, Where Are We Now?

AI, in conjunction with the general advancement in speed of computer hardware, has helped deliver machines that are faster, stronger, and better (subjectively to be debated) than humans at many things, and they are only going to get even faster, stronger, and “better.” Machines compute many times faster than humans (though we still outperform them in parallelism). Machines are capable of defeating human champions in many games and competitions, including chess, Jeopardy!, and Go. Machines can now match, or outperform, humans when it comes to reading, writing, listening and speech—machines even outperform us when it comes to sensitivity of senses such as smell and touch. These capabilities are already being used to underpin other activities, such as driving cars, trucks, buses and trains, operating complex machinery, assembling other complex machines, etc. And, importantly, machines can do what they do repeatedly, with fidelity, without complaint, and for as long as they have a power source.

At the same time AI has been advancing there have been similar advances in robotics. Robots are a physical manifestation of AI—containers that allow AI to be mobile, with access to an array of sensors. Over the past decade or so robots have become increasingly prevalent in the military, industry, and at home. Miniaturization of electronics, including, and importantly, batteries, in conjunction with the advances in AI have led to robots that are capable of a range of movement, dexterity, and finesse that has not previously been possible.

Industrial robots—autonomous machines that weld, stack, pack, cut, fold, clean, etc.—are if not ubiquitous, almost commonplace. Autonomous vehicles, though less common, are becoming a reality. We can now choose from several robotic vacuum cleaners available at reasonable prices. Robots that can do our laundry, even fold our clothes, are available and will eventually become as commonplace as industrial robots. Robots can serve drinks and clean spills. Robotic guides are in use at some museums. Robots are routinely used in surgeries of all types, including remote surgery where the doctor and the patient are separated by some distance.

The military routinely deploys drones of all types: on land, in the sea, and in the air. We have robots that can navigate terrain of all types, even open doors and not only avoid obstacles, but go around them and stay on track to their destination. Quadruped robots, in use in the military and useful in areas that are not navigable by wheeled vehicles, are able to carry weights and move faster than any live pack animal. Bipedal robots are not yet at the same level as the quadrupeds, and probably won't be for some time. Humanoid robots, especially humanoid robots that actually look and move like humans, are even further away—advances in the dexterity and finesse in robotic movement have not brought us there yet, nor are the materials to simulate the human exterior sufficiently advanced. Certainly Star Trek's Data is not even close to being on the horizon.

Sentient machines are not walking among us yet, and won't be for some time—many decades at least, if at all—but machines are certainly capable of performing tasks that have until now been the sole purview of humans. But are they (the machines) intelligent?

What intelligence actually looks like, and how to define it (and indeed, who should define it) is not the subject of this article—that discussion could, and does, fill entire books. For the purposes of this discussion, I'll simply define the intelligence that is the subject of AI as *human* intelligence—broadly, the ability to exhibit a range of behavior (not necessarily movement) similar to that exhibited by an “average” human. To be considered to be truly intelligent, a machine needs to be able to do the things that humans do, but not just one or two things—a whole range of them. We can build machines that can outperform humans at specific tasks, but no machine has outperformed a human across the whole gamut of behaviors and abilities routinely exhibited by humans.

We should be careful not to conflate the ability to follow instructions and complete tasks quickly, even complex instructions and tasks, with exhibition of human intelligence. It is true many machines today are capable of following very detailed and complex instructions in order to complete complex, and often difficult, tasks. It is even true, as discussed earlier in this article; many machines today are capable of outperforming humans in many complex tasks. But following instructions and completing complex tasks, even if done quickly, doesn't necessarily require human intelligence. The Jacquard machine, created by Joseph Marie Jacquard in 1804, and based on work done in 1725 by Basile Bouchon, was a weaving loom modified to allow it to follow instructions recorded on a series of punched cards, to automate the manufacture of textiles with complex patterns. The Jacquard machine was an incredibly important machine, to textile manufacturing, automation, and computer science—but it didn't possess human intelligence, and it didn't need to in order to do its job.

The world is awash with “smart” devices—smart phones, smart televisions, smart watches, etc. These devices are labeled “smart” because they are able to compute—they have embedded processors that earlier devices didn't have, and they are able to do things their predecessors could not. But smart? Smart connotes intelligence, and while these devices are able to perform tasks and exhibit functionality far superior to their (even recent) predecessors, by-and-large they simply follow instructions programmed into their embedded computers. Some may exhibit predictive or learned behavior, but even those are just doing what they've been instructed to do—they just have more complex and involved instructions.

Some machines are able to learn and predict, but only because they've been programmed at least with the strategies, if not the actual algorithms, to do that. The strategies and algorithms they use to learn and predict are all developed by humans, generally incrementally over a long time. No machine has yet learned, unguided by humans, the ability to learn. Moreover, machines that can learn are told, by humans, what to learn. If a machine is learning to solve a problem, humans first have to tell the machine that a problem exists and what the problem is—and generally what form the solution should take. No machine has yet learned, unguided by

humans, to recognize that a problem exists, that it needs to be solved, and how to solve it. The reason for that is as I have outlined: learning *from scratch* is an incredibly difficult problem. Carl Sagan described the problem most eloquently when he said, “If you wish to make apple pie from scratch, you must first create the universe.”

Of course, learning everything from scratch is not necessarily a requirement of an intelligent system. After all, humans don’t learn everything—even most things—from scratch. Mammalian brains are composed of three areas, often referred to as separate “brains within the brain,” that evolved over time and serve different purposes. The two oldest parts of the mammalian brain, the “R” (reptilian) complex and the limbic system, are responsible for what we might call our reflex actions: hunger, temperature control, fight-or-flight fear responses, keeping safe, defending territory, emotional responsiveness, memory formation, etc. Humans don’t learn these things; they’ve been refined over millennia by evolution, and we get them for free at birth. We do, though, learn from them, and even add to them as we experience life. Humans learn most things by using the past experience of others—by bootstrapping our knowledge and understanding from existing knowledge and understanding, then learning incrementally from there. Human children don’t learn everything they come to know by adulthood from scratch—some of it is inherent, endowed by evolution, and much of it is passed on from family and tribal elders. We can bootstrap the knowledge of our AI systems in a similar way. The AI systems we construct don’t need to learn from scratch; they can use the collected knowledge of humans and learn from there.

Acquiring knowledge and being capable of following instructions isn’t necessarily the same as becoming intelligent. Wikipedia [defines](#) Artificial General Intelligence (AGI) as “the hypothetical intelligence of a machine that has the capacity to understand or learn any intellectual task that a human being can,” and notes that “Today’s AI is speculated to be many years, if not decades, away from AGI.” Machines haven’t achieved AGI yet, but that doesn’t mean they haven’t acquired some semblance of intelligence—we have made great strides in AI since 2006.

Beyond 2020, Where Are We Heading?

Machines and virtual machines (i.e. software robots) have obviated, or are on track to obviate, the need for human involvement in many tasks, particularly, but—importantly—not limited to, tedious, repetitive, and dangerous tasks. Machines are already “better” than humans at many things. What does all this mean to us?

In a 2019 white paper, leading research firm IDC estimated the contribution of “digital workers” (software robots and AI) to the global workforce would grow by more than 50 percent by 2022. We won’t just see intelligent machines working in dangerous or inhospitable places, we will see digital workers in the office, trained to perform many and varied business and office related tasks.

Clippy may have been almost universally reviled, and unceremoniously dumped by Microsoft, but we will all have digital worker assistants, even co-workers, in the future. Many organizations have already implemented, or begun implementing, AI in the workforce—in the form of both physical and software robots. This trend will only continue, and at some stage physical robots and digital workers will outnumber human workers. A clear corollary is that we as a society will undoubtedly have more, and probably “enhanced” (e.g. implanted electronics to facilitate interfacing and communication), interactions with AI, both physical and software-based.

Autonomous vehicles, particularly self-driving cars, will bring great changes to our society, our way of life, and to business and commerce. As self-driving cars and trucks become ubiquitous our roads will eventually become safer and more efficient, leading the world to consume less fossil fuel and, as a consequence, pump less pollutant into the atmosphere. The cost of transportation, for both commerce and leisure, will fall due to increased efficiency on the roads and a decrease in accidents causing damage, time and product loss.

There are potentially huge benefits in having machines perform all sorts of tasks, especially those that are dangerous or tedious, to some people, just as there are potentially huge negatives to others. For every job where a machine replaces a human, there is a human who either needs to reskill or is out of a job. For all the gains in efficiencies and safety autonomous vehicles will bring, their advent raises issues that need to be addressed—one of the biggest questions that faces us is whether we are ready to hand over moral decisions to AI.

Current aircraft autopilots are relatively simple compared to autonomous road vehicles. There's not a lot of uncontrolled, close traffic up there, and even fewer children and animals running out in front of the aircraft, not to mention the fail-safe of pilots sitting at the controls. What does an autonomous road vehicle do if a small child runs out in front of it, and the only options are to run the child over or drive the vehicle off a bridge (or cliff, or into a brick wall, or another vehicle, or bystanders)? If the choice is (potentially) kill the child or (potentially) kill the occupants of the vehicle, what choice does it make, and on what basis? Are we yet willing to put this sort of decision in the hands of AI? Are we willing to indemnify the makers of the vehicle against legal action and liability in such circumstances? Who would decide, and how, whether the AI controlling the actions of the vehicle acted correctly, or whether the algorithm was defective? The technology is clearly not insurmountable, but are we ready to deal with it? Siri, Alexa, and other virtual assistants will no doubt be a great boon to many of us, helping us manage our busy lives. But these devices are always listening: How can we be certain our privacy and the security of our data is protected? Google, Facebook, Microsoft and Amazon already collect petabytes of information about us—there is no reason to believe that virtual assistants won't become another source of data and income for the digital giants.

We must be vigilant in ensuring the protection of private and personal data, and should look to regulating the harvesting, storage, and use of private and personal data. The European Union (EU) seems to be leading the way with the General Data Protection regulation (GDPR), which

became enforceable in the EU and the European Economic Area (EEA) in May 2018. The U.S. state of California adopted the California Consumer Privacy Act (CCPA) in June 2018. Some other countries either have introduced, or are planning to introduce, their own laws and regulations in this space. Uniform laws and regulations should be negotiated and implemented.

In Summary

There have been many advances in AI over the past decade-and-a-half, but I think its most remarkable achievement is moving from science fiction to mainstream in the minds of the general public. People are generally more willing to consider AI to be a good thing, and to contemplate it becoming a natural part of their lives than at any time in the past.

As AI becomes more ubiquitous in the everyday lives of humans, and as we cede responsibility for decision making in a wide variety of situations to AI, ethical dilemmas, and the conformance of AI to some underlying set of morals, and who decides what those morals should be, will be of paramount importance. Trust is a must. If we cannot rely on the AI in our lives to make reliable, repeatable decisions on our behalf that reflect us, our desires, and our morals, then AI is destined to be consigned to the factory floor and mineshafts.

Further advances in AI will happen. AI is already redefining and reshaping how we live, play, and work, and will continue to do so, probably at a faster pace and more dramatically than most people thought possible. We, as a society, need to think through the consequences carefully and prepare for their eventuality.

Finally, the path to AGI—to machines thinking like humans—will be a long one marked by many sidetracks, wrong turns, and dead ends. Recent advances in AI and robotics have largely been incremental technological advances, each one building a little on existing technology, rather than any radical new insights providing a leap forward. We are decades, probably tens of decades, from our goal of creating machines with human intelligence. To make real progress we may yet need to recast artificial intelligence as a science in its own right, rather than constrain it as a subfield of computer science.

The promise of AI is indeed elusive.

Biography

Jeff Riley is a 40-year veteran of the computing industry, and has worked in many roles in the field including Master Technologist and Scientist with Hewlett-Packard. Jeff holds a master's degree in IT, a Ph.D. in AI, and is a former Adjunct Principal Research Fellow of RMIT University in Melbourne, Australia. He is the founder of Praescientem, a company specializing in AI and IT consulting and education services. Jeff is currently pursuing a second Ph.D., in theoretical



astrophysics, and builds robots in his spare time. Jeff's major areas of interest in the AI space are machine learning and evolutionary computation. His publications are available at <https://www.praescientem.com.au/pages/research.html>.

DOI: 10.1145/ 3458742