Technica®

AI was coined as a term in 1956 at a Dartmouth College Computer Science conference. It refers to a line of research that seeks to replicate the characteristics of human intelligence. Over the years this broad and varied line of research has touched such disciplines as:

- Automated Planning and Scheduling
- Computer Vision
- Robotics
- Reasoning and Cognition
- Knowledge Representation
- Natural Language Processing
- Machine Learning

In the 1950s, "General AI"—equality in intelligence between humans and machines—was thought to be within reach. This sentiment proved to be too enthusiastic. While computers could be taught to beat the best chess players, this did not translate to intelligence on par with humans. In the last 50 years, AI has gone through two boom-bust periods. Will this iteration be any different? A key to answering this question is to focus more narrowly on important drivers of the current advancements in AI.

The Technica Innovation Platform White Paper Series presents advanced topics that will drive competitive advantage for next-generation IT over the next three-to-five years.

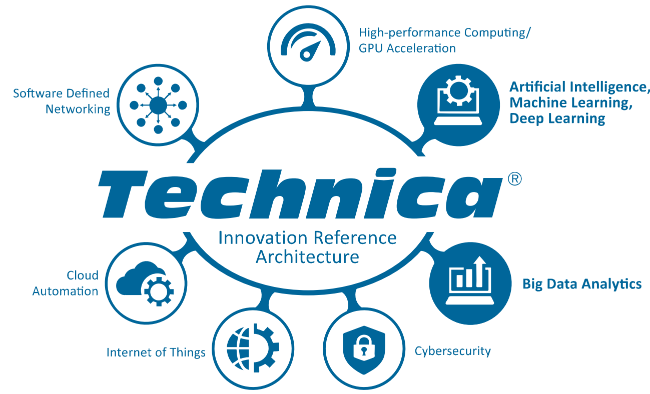

ARTIFICIAL INTELLIGENCE (AI): HYPE OR HOPE?

"Artificial Intelligence," (AI) a term with a long history, is once again an industry super-buzzword. Like previous buzzwords of "cloud computing," "Big Data," and "Web 2.0," AI is taking a life of its own. Startups highlight it in every funding pitch-deck to garner higher valuations. Established vendors re-examine their websites making sure AI shows up somewhere. Scientists and billionaires like Stephen Hawking, Bill Gates, and Elon Musk ring the warning bells of AI super-intelligence—the point in time at which computer-smarts overtake human brains.

Despite the hype, recent advances are truly remarkable. AI will indeed impact nearly every business in the developed world. This paper seeks to cut through the hype of AI and provide a balanced conceptual framework by which to understand the promises and pitfalls of this latest iteration of AI.

MYTH BUSTING DEFINITION

AI is a branch of computer science that deals with algorithms—digital recipes—inspired by various aspects of natural human intelligence.

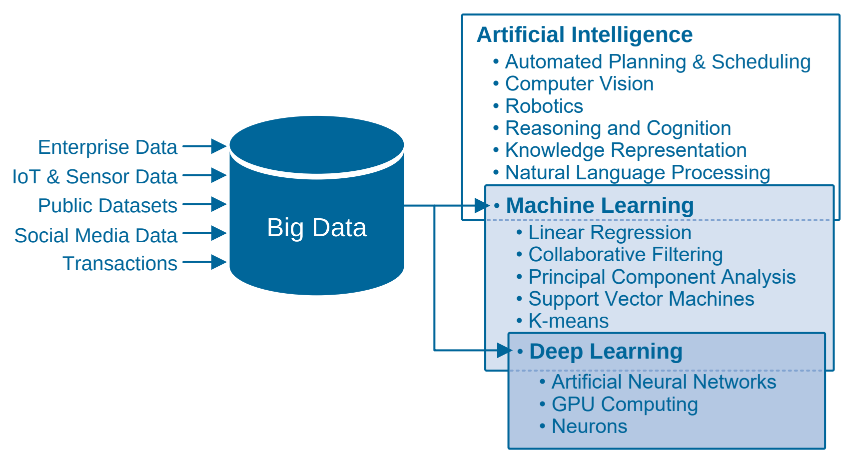

Figure 1 – Relation of Big Data to AI

Unpacking some of the terminology related and often used interchangeably can help clear some confusion. **Figure 1** presents some of these concepts.

BIG DATA

Big Data refers to data that is so large that standard processing mechanisms, like relational database management systems, can no longer be utilized. As highlighted in Figure 1, sources of Big Data are varied. They include existing enterprise data found in Customer Relationship Management (CRM) systems and other enterprise data stores and transactional data.

But the real growth in Big Data comes from external sources like that found in public data, like OpenGov.org, and social media data sets. The growth of the Internet has caused an exponential increase in data. Previously unimaginably large data sets are now being collected.

In the near future, data from sensors and the Internet of Things (IoT) will increase, by multiple factors, the current exponential growth of Big Data. The current and future trends in the growth of Big Data variety, velocity, and volume mean data analysis will only increase in importance. Machine learning has come to the rescue.

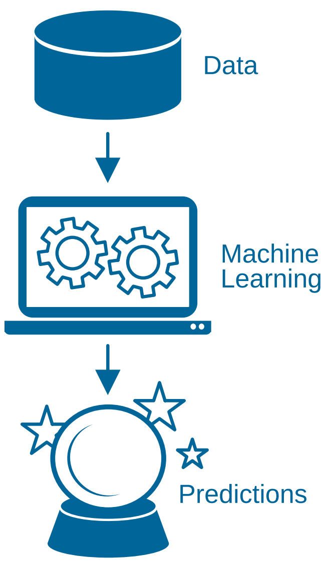

Figure 2 – Machine Learning Steps

MACHINE LEARNING

Machine Learning is a sub-discipline within AI that employs algorithms that can learn from data. Machine learning is widely used in spam detection, credit card fraud detection, and product recommendation systems from the likes of Amazon, Netflix, Spotify, and iTunes. In fact, much of the hype surrounding AI actually relates directly to machine learning.

The key aspect of machine learning algorithms is that they build a model based on inputs. Using this information the models then make predictions based on the inputs, as shown in **Figure 2**. Machine learning distinguishes itself from previous similar approaches, like expert systems, in that the algorithms make predictions without following explicitly programmed instructions.

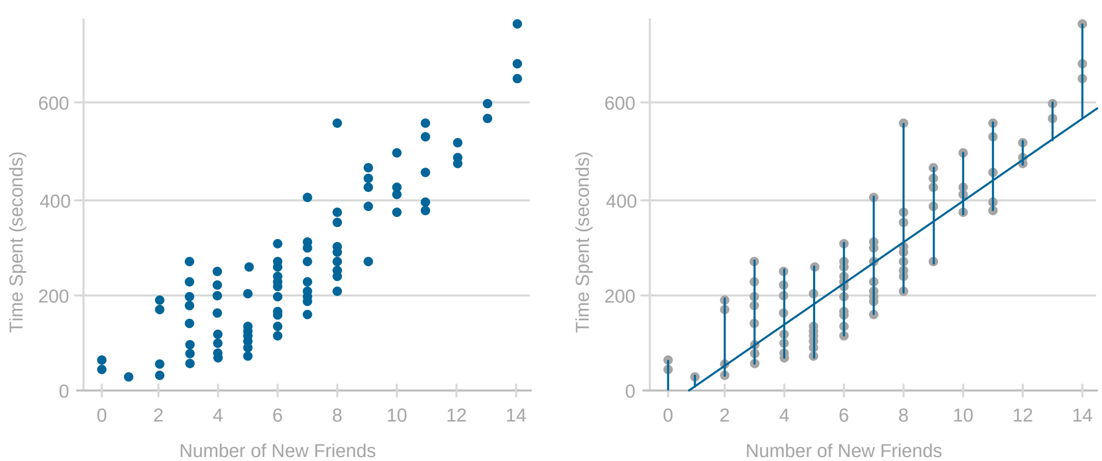

Figure 3 – Sample Linear Patter and Regression Line (source: Rachel Schutt & Cathy O'Neil)

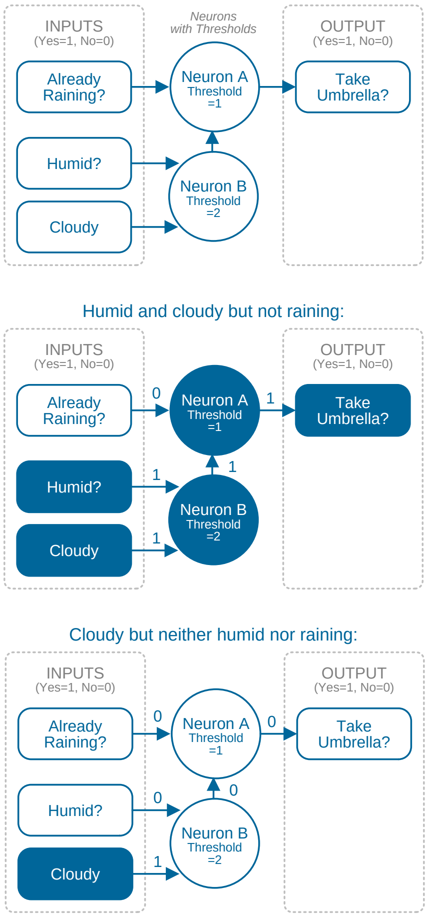

Figure 4 – Simple Neural Network Architecture and Examples

Figure 3 presents a simple example from a social networking site. The horizontal axis presents the number of new friends in a week for the website's user population. The vertical axis plots the time spent online by users.

It's clear that there is a linear relationship. Using linear regression, the task is to draw a "line of best fit" given the initial observed data. Figure 3 presents the scatter points utilizing a method known as least squares. The regression line is the predictive output, or model. That is to say, given another user with 6 new friends in the last week, the model predicts the user will spend about 200 seconds on the site.

Linear regression algorithms are just one type of machine learning algorithm. Technica's FUNL analytic tool includes numerous others including collaborative filtering, principal component analysis, support vector machines, and K-means. Importantly, Technica tuned each algorithm to take advantage of Graphics Processing Units (GPUs). This approach improves performance and reduces costs by orders of magnitude. It is not uncommon to see GPU accelerated systems cost 100-times less than their CPU-only competitors.

It is obvious that the bar to claiming the use of AI is set low. Many of the companies touting advanced AI are simply applying machine learning algorithms. But a certain category of machine learning has experienced significant publicity—deep learning.

DEEP LEARNING

Deep learning is one of the most exciting approaches to machine learning in the last decade. Other papers in this series provide more in-depth coverage of deep learning. This paper highlights the basic components of deep learning—neural networks and neurons. **Figure 4** provides a sample neural network and two examples.

Imagine trying to create an artificial neural network that predicts if you should take an umbrella to work. The following simple model consists of two neurons, Neuron A and Neuron B. At any given moment, each neuron is either "firing" or "idle." An idle neuron does not pass messages or signals. The neurons threshold determines whether it is "on" or "off." Neuron A's threshold is 1, while Neuron B's threshold value is 2. There are three inputs pictured to the far left (already raining?, humid?, and cloudy?). The output of "take umbrella?" triggers if it receives the value of '1' from Neuron A. The layout of this neural network is shown at the top of Figure 4.

The middle drawing in Figure 4 presents the first example. It is humid and cloudy, but it is not already raining. Therefore Neuron A's threshold of 1 is not met. However, Neuron B fires (its threshold of 2 is reached because it is both humid and cloudy) to trigger Neuron A. Therefore the neural network predicts the need of an umbrella.

The second example, the bottom drawing of Figure 4, shows that it is cloudy but neither humid nor raining. Neither neuron triggers and the model predicts that an umbrella is not needed.

Surprisingly the math behind deep learning is simple. But, current deep learning algorithms employ 20 to 100 layers (with millions of individual neurons), taxing even the most advanced Central Processing Units (CPUs).

The parallel processing power of GPUs is the real driving force behind advances in deep learning algorithms. Deep neural networks that used to take weeks to train with CPUs can be trained in hours utilizing GPUs. Technica has created a deep learning algorithm for FUNL called "DeepInsight." It uses deep learning to predict non-obvious relationships in graph data, like social networks and bodies of text. GPU-accelerated neural networks are one of the main factors for the renewed interest in AI.

Figure 5 – AI Predicts Panda with 60% Confidence

Figure 6 – AI Predicts Gibbon with 99% Confidence

AI CHALLENGES

AI centered around machine learning in general, and deep learning in particular, is not without issues. Currently, the greatest limitation of deep learning is that it requires massive amounts of labeled training data. While the growth of Big Data often solves the data quantity problem, to date deep learning has excelled only with labeled data. While it is impressive that deep learning algorithms for image recognition have identified cats in pictures faster and more accurately than humans viewing the same pictures, the algorithms relied on millions of images tagged as "cats." Deep learning's other great successes in speech and handwriting recognition were also due to high volumes of labeled data. However, much data remains unlabeled. Data scientists are currently seeking remedies for this dependency.

Another flaw is that by tweaking an input in a specific way one can fool a deep learning neural network to misclassify. This is a problem for not only deep learning but for almost every machine learning method that employs linear algebra to approximate its classifications. The human brain does not make this same kind of misclassification. Consider **Figure 5**. A deep learning algorithm estimated that the figure was a panda bear with a 60% confidence. **Figure 6** is imperceptibly different to the human eye, but includes subtle changes to a 32 point floating variable used to encode the picture. The algorithm guessed with 99% confidence that the animal was a gibbon.

One cannot speak of AI without noting dire warnings offered by several prominent intellectuals and business people. Elon Musk has stated that AI is "potentially more dangerous than nukes." Most of these fears deal with super-intelligence—the theoretical period of time in which computers get smarter than human beings.

Some believe it stands to reason that as scientists and developers continue to improve technology at an accelerated rate and society continues to increase our dependence on technology, those AI systems will eventually be designed to improve themselves (perhaps through advanced unsupervised deep learning techniques) to the point of becoming smarter than the human beings that programed them.

With intelligent machines having the ability to think flexibly across multiple domains they could design and build even better computers without the need for human intellect. Theoretically this raises some areas of concern and implications as we develop AI—how should ethics be programmed into automated or intelligent systems?

For example, consider the case where a driverless car's brakes fail and the only two possible paths for the car's current direction lead to a cliff or a crowd of people. Where does probability and predictive outcomes trump moral responsibility? Would we be able to shut down a super-intelligent machine if it violates our laws, values, or morals?

Without becoming alarmist, it does make sense to think about these future implications of AI. What is clear is that we have not even arrived at general intelligence for computers, the term used to describe the equality of human and computer intelligence. It is more likely that AI will be used to augment many human capabilities, as computers become more important tools for humans.

However, of great importance in the near-future will be the effect of AI on the workforce. Many jobs involving driving (delivery man, truckers, taxi cab drivers, etc.) will be displaced by self-driven AI vehicles. Newspapers are already using AI publishing platforms to author sports and weather stories.

SUMMARY

GPU computing, the growth of Big Data, and deep learning have supercharged the current interest in AI. All of these factors can only be expected to become faster, bigger, and better. Data scientist are currently looking to improve deep learning algorithms to deal with unlabeled data, in addition to devising other complimentary sets of algorithms.

As pictured in Figure 7, IBM foresees a world in which improving both Big Data predictions and AI merge to become one, in a discipline they call "cognitive computing."

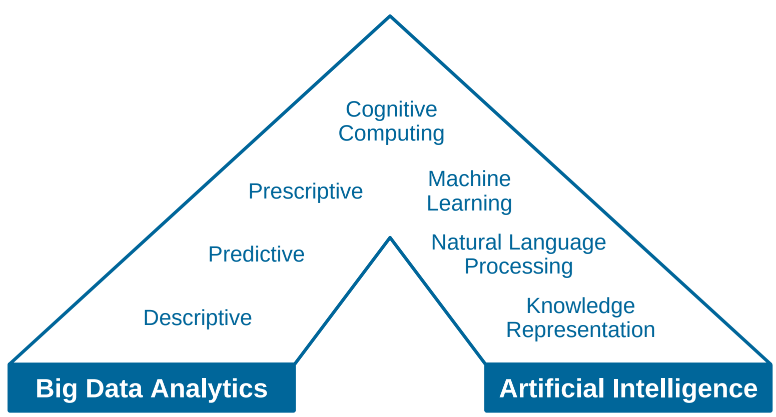

Figure 7 – Convergence to Cognitive Computing: source: "Cognitive Computing and Big Data Analytics," Judith S. Hurwitz, et. al.

Whether cognitive computing will actually arrive at general or super-intelligence remains to be seen. Academics and science fiction writers have been pondering this for close to 100 years. However it is clear that the future is bright for both Big Data analytics and AI, and both areas will see significant growth in opportunities and possibilities.

Technica®

22970 Indian Creek Drive, Suite 500
Dulles, VA 20166
703.662.2000