ARTIFICIAL INTELLIGENCE: A LOOK AT THE GOOD, THE BAD AND THE UGLY

By Jeffery S. Rowland

We have all heard many opinions around "artificial intelligence" and what it could mean to our professional careers. Because all industries stand to be impacted in one way or another by this quickly emerging technical capability, the truth is that we have yet to fully understand the total impact of what a mature deployment might really mean to us, either individually or collectively. The technology is gaining momentum in business and across society, and one thing is very clear: we will all need to learn to adapt.

As with many things, if something can be used for good, it can also very likely be used for bad. In this article, we will investigate some of the various use cases that have been contemplated. Reminiscent of the .com boom of the mid 1990s, it is possible that we will see similar "game changing" deployments in the 2020s that go well beyond what we can imagine today.

So, what is artificial intelligence? Most people, after thinking about it as more science fiction, go toward a common and simple definition along the lines of "machines that can think independently." A Google™ search on the question "definition of artificial intelligence" returns a series of websites. One of the better results was from The English Oxford Living Dictionary as follows:

"The theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making and translation between languages."

Most people, after thinking about it as more science fiction, go toward a common and simple definition along the lines of "machines that can think independently."

Although it's a complete definition, let's explore other sources. When asking Siri™ "Siri, what is artificial intelligence?" we receive a much less satisfying answer than before: *"Branch of computer science that develops machines and software with human-like intelligence"*

To the question of "Siri, are you artificially intelligent," she responds: *"I don't have the answer to that. Is there something else I can help you with?"*

Let's ask yet another source to determine the best definition. Google Home™ responded to the question "OK Google, what is artificial intelligence?" with:

"Artificial intelligence is the theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision making and translation between languages."

Google actually pulls the definition from Oxford. Just as with Siri, we followed the answer with, "OK Google, are you artificially intelligent?"

Answer: "*My intelligence is artificial.*" Well, at least that was a better answer.

Although this exercise might seem somewhat rudimentary, we can come to an understanding that artificial intelligence, at least in its current form (and perhaps in its future form as well), is limited by the information it is fed. However, the real power might not be apparent until the technology can demonstrate being able to compare multiple sources of information in the blink of an eye and then present the results back to the originator of the question those answers most common across the data sources. We need to consider the implications of feeding a sophisticated artificial intelligence program a healthy dose of "big data" and then we might begin to see the real potential power. We are not even to the point of talking about "thinking machines" and are just scratching the surface on what the technology might be able to do for us.

The Four General Types of Artificial Intelligence

The industry has classified artificial intelligence into the following four general types.

Type I – Reactive artificial intelligence: Reactive artificial intelligence will respond to the same situation in the same way, each and every time. A Type I deployment has no ability to store memories, an example being something along the lines of an automotive assembly robot that places a weld in a specific location each time it sees a chassis that it has been programmed to look for.

Type II – Limited memory artificial intelligence: Limited memory artificial intelligence can recall stored data, but new memories are not saved, meaning it cannot "learn" from its past experiences. An example might include a self-driving car or perhaps a home internet personal assistant (e.g., Google Home, Amazon Alexia or Siri).

As a side note, there have been cases in the news where these home personal assistants were temporarily storing what they gathered and authorities were trying to extract that data following the commission of a crime, as these devices are always "listening" until they hear their cue words. Even so, it would still be classified as Type II.

Type III – Theory of mind: Theory of mind is a reference to the notion that a machine can recognize others it interacts with, including the ability to have thoughts, feelings and even expectations. A machine equipped with this technology would be able to understand interactions with others and interpret conversations. This interpretation would include detecting the feelings and expectations of others, and adjusting its behavior and interactions accordingly.

Type IV – Self-awareness: This is the version of artificial intelligence that the general public seems to fear, mainly because a machine equipped with Type IV would, in fact, be self-aware. As the next progressive step beyond the "theory of mind," a self-aware machine would be completely aware of itself and have knowledge about its own internal state of being, while also being able to predict and interpret the feelings of others.

Taken to an extreme, Type IV thinking machines that are capable of decisions and of even programming other machines creates a dynamic, especially in the fields of accounting, law or even health care, where the ability to understand how certain conclusions have been reached may be difficult, if not impossible. It is for this reason that the blending of certain industries with the requirement that technical aptitude and knowledge be gained to understand this technology is clear.

For example, the CPAs, lawyers or doctors of the future might very well be required to not only understand accounting, law or biology, but might also be required to master data analytics and artificial intelligence as a part of their professional education.

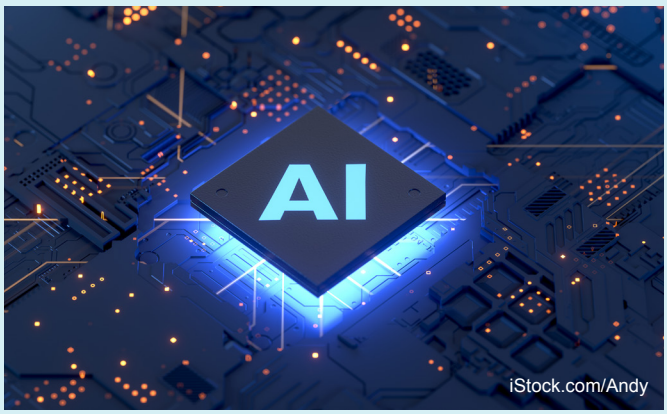

As another example, look at the car industry and how technology has driven the educational expectations of mechanics. A few decades ago, a car could be worked on with nothing more than a toolbox full of Craftsman™ wrenches. For modern automobiles though, trying to do anything major to your car without a diagnostic computer is difficult, if not impossible. Now, understanding computer diagnostics is a key part of a mechanics' formal training.

Now, let's examine some of the use cases.

The Good

Health care: The first and perhaps best "use case" is around the use of artificial intelligence for the good of human health. If you consider the ability of a machine that is fed data for millions of ailment symptoms, combined with the ability to analyze hundreds, thousands or even millions of reference pictures for any given health problem, you start seeing the power for improving our ability to reduce the time needed to come up with possible matches and then to accurately diagnose those ailments.

The first and perhaps best "use case" is around the use of artificial intelligence for the good of human health.

Cutting down the amount of research time in health care could drastically decrease diagnosis times so earlier treatment can begin and with the advent of artificial intelligence, it is more practically feasible than ever before for the medical community. In addition, imagine the potential application of remote diagnosis in towns and villages that might not be able to afford to have a physician nearby. Using technology to further a doctor's reach beyond physical limitations, expanding the ability to efficiently and effectively treat human ailments to any part of the world is an astounding prospect and certainly one of the most significant "greater good" aspects.

Convenience: Contemplating "the good" of artificial intelligence, we can look no further than the example of self-driving cars. The benefits of this technology, particularly to our senior population, would be crucial. Instead of being advised, either by family members or law enforcement, that you are no longer capable of driving safely and must surrender your right to do so, imagine the ability to maintain that independence by simply saying aloud "Take me home, car." Perhaps we will even be able to "name" our self-driving car!

Faster, Better, Cheaper: The prospect of artificial intelligence being designed to do jobs that are hazardous

to humans also comes to mind. In business, we constantly try to manufacture things faster, better and cheaper, and the automotive industry is already well on its way in using automated robots to weld parts together as they move through the assembly line.

In another use case, Amazon™ uses automation technology in its order fulfilment facilities and mining companies use partial automation to extract hazardous minerals, such as coal.

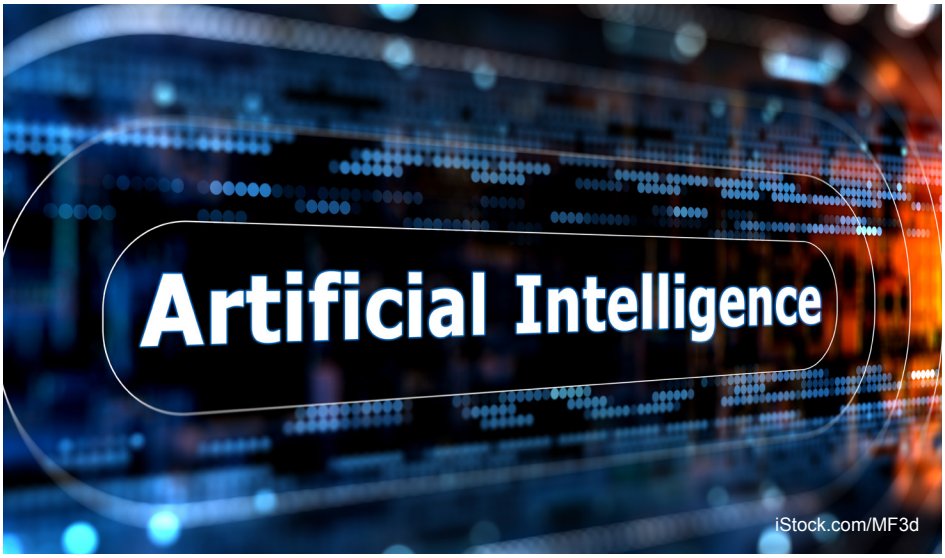

The many use cases are clear in using technology to help create greater efficiencies while also increasing safety to humans.

The Bad

In a 2014 interview with the late Stephen Hawking, arguably one of the most brilliant thinkers of our time, he said: "The primitive forms of artificial intelligence we already have, have proved very useful. But I think the development of full artificial intelligence could spell the end of the human race. Once humans develop artificial intelligence, it would take off on its own and redesign itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded."

In a 2014 interview with Elon Musk, he stated: "I don't think most people understand just how quickly machine intelligence is advancing. It is much faster than almost anyone realizes, even within Silicon Valley and certainly outside Silicon Valley, people really have no idea. If there's a super intelligence, particularly if it is engaged in recursive self-improvement, if there's some digital superintelligence and its optimization or utility function is something that is detrimental to humanity, then it will have a very bad effect. It could be just something like getting rid of spam email or something, and it concludes that the best way to get rid of spam email is to get rid of humans."

Although these examples from highly intelligent people may seem extreme and might be easily discounted as unlikely "science fiction," the boundaries that are necessary to avoid creating these extremes don't seem to exist or at least are not at the top of mind today. Consider when automobiles first began to gain prevalence in our society … there were no well-developed roads, signs or laws. There were quite a few accidents and fatalities that likely prompted the creation of the laws we have today.

Unfortunately, making a mistake with a technology such as creating Type IV artificial intelligence that could result in a cataclysmic demise of humanity would deserve our attention in terms of laws and regulations. Because laws have traditionally trailed innovation, such a lagging approach would not appear to be the best approach.

The Ugly

Consider the implication of artificial intelligence being used in the construction of a dedicated killing machine. Although this sounds "Terminator"-like, a very realistic but fictitious YouTube video begins with a professional sales-type conference setting. The speaker introduces a new "drone" type of killing machine that is imprinted with biometric specific instructions where it tracks its target, only to self-destruct itself by exploding a one-ounce charge of explosive sufficient to scramble the brains of its pre-programmed target. Minus the biometric imprinting, the attempted assassination attempt on the Venezuelan President using a drone laden with several pounds of explosive starts to make this sound much less like unlikely science fiction.

Finally, another ugly aspect is the inability of programmers to effectively program morality. The fundamental aspects of what is inherently right and wrong are subject to the attitudes and beliefs held by the programmer. If the "programmer" is one day an artificially intelligent machine, the concerns around the "ethics" of the machine are clear, especially if the machine

was making life and death decisions.

Because of the moral and ethical consistency challenges, the need to align the technical community on a common set of practices is compelling. If ever machines routinely program machines, it will take a machine to understand what is being audited.

Impact on Auditing

We must continue encouraging the auditors of the future to embrace an educational curriculum that requires understanding of this type of technology.

If ever machines routinely program machines, it will take a machine to understand what is being audited.

Imagine trying to audit a deployment of artificial intelligence in a traditional way. Some of our military deployments of this technology have over 10 million lines of code. It would take an army of programmers literally months to get a handle on what the code is doing.

Therefore, auditing artificial intelligence in the future will require an auditor who is adept at using the very technology that is being audited, which is not so different than the approach we try to take today.

Will the deployment of this technology replace us in lieu of a robot? Perhaps in some ways, but for those who embrace the change and start to sharpen their skills in understanding this emerging technology, the future begins to look more like an opportunity than one where jobs are eliminated.

ABOUT THE AUTHOR:

Jeffery S. Rowland, CPA, CFE, CISA, CIA, CRMA, practices internal audit in his capacity as Vice President, IT/Security Audit with USAA in San Antonio. A graduate of SMU and a proud alumni of EY, he has almost 35 years of experience in the audit field. He may be contacted at jefferyrowland@gmail.com.