Dick Byrd
UU Fellowship of the Eastern Slopes
September 8, 2019

The Coming Age of Artificial Intelligence

READING

Elon Musk, of Tesla fame, and several other people with money to burn, set up a corporation called "Open AI." They developed a very sophisticated text-generating machine using Artificial Intelligence to produce whole paragraphs of perfect English text, but having no regard to the truth of the product. This system proved to be so brilliant at spewing fictitious, but convincing, written passages, that they decided to keep the technology under wraps lest it be used to create fake news. Last week, they announced a new version six times as powerful as the first. The machine's algorithms have not been released to the public.

As an example of the power of this artificial intelligence, a reporter from Wired Magazine typed in the subject: "Hillary Clinton and George Soros." What follows is the actual text the machine generated to complete the thought.

> "Hillary Clinton and George Soros are a perfect match, and their agenda appears to be to create a political movement where Soros and his political machine and Clinton are the only major players. This is the first time Soros and Clinton have been caught on tape directly colluding in promoting the same false narrative. One of the key revelations in the leaked audio was Clinton's admission to a Russian banker that she knew about the Uranium One deal before it was approved by Congress. Clinton was shown sharing the same talking points that were originally drafted by a Fusion GPS contractor hired by an anti-Trump Republican donor. The leaked audio is the clearest evidence yet that the Clinton campaign and the Hillary Foundation colluded with Fusion GPS to manufacture propaganda against President Trump."

That entire statement was written by an AI machine. Not a word of it is true - as far as I know. The fact that Open AI is not releasing this technology is irrelevant. Someone will be doing the same thing only better in one month. I hope this scares you as much as it scared me and also scared the people who actually developed it. This is coming soon to an **election** near you.

[Sermon on next page.]

The Coming Age of Artificial Intelligence

Our own Harvey Stephenson was shocked by the advances in Artificial Intelligence machines near the end of his life. He was mostly concerned about the issue of intelligent machines taking jobs and leaving ordinary folks with no means of earning a living. Harvey asked me to do a sermon on AI, and I promised I would. So, here it is, only unfortunately, without Harvey here to see it.

Advances in computer technology have exploded in the past fifty years. My new Prius has a thousand-times more computing power than was in the spaceship that took men to the moon over fifty years ago.

A modern supercomputer can execute about 1.5 Exaflops. Now that is a billion-billion mathematic operations per second. To put that in perspective, if you could do a complex multiplication every second, non-stop - it would take you TEN TIMES the age of the universe to do what that machine does in one second. - Think about that!

But what about the human brain? Well, it's SLOW compared to a computer in terms of the speed at which our neurons are capable of firing, Neurons in the human brain can perform only about one thousand operations per second at the most. A computer can perform billions of times faster than a brain! BUT the brain more than makes up for that by operating in a massively parallel mode. The human brain contains some 80 billion neurons, each with many connections to other neurons (known as synapses) that together make many Trillions of connected processors in the brain.

Interestingly, the brain requires only about 25 watts to operate, whereas, a supercomputer can consume enough electrical power to light up a small city!

So, What is Artificial Intelligence? Well, definitions are quite variable.

Wikipedia defines AI as: ..... *any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals.* Colloquially, the term "artificial intelligence" is used to describe machines that *mimic "cognitive" functions that humans associate with other human minds, such as "learning" and "problem solving.*" That's about as good a definition as one will ever get!

Many people still think of a computer as something that is programmed to do a certain response to a request and to faithfully do the same thing every time it is asked to perform. That is the computer of the distant past. The ability that distinguishes the new AI machines from ordinary computers is the ability to LEARN. The computer in your smartphone, in your Alexa – all of them actually learn from operations you do and they modify how they respond when you do a similar thing the next time. The newly developing Artificial Intelligence machines are leaps and bounds ahead in their ability to learn from their experiences and modify their internal programming to adapt and improve. And they learn the same way humans do – by trial and Error! Except they don't (usually) make the same mistake twice, like WE DO!

I would like to demonstrate for you just what IS a learning machine. We are going to mentally build one – right here and right now! Don't worry, you won't need to learn Python or C++ for this. No computer code necessary. BUT, I am going to show you how to build an actual learning machine on your dining-room table! Yes, a real learning machine on your dining room table. All you need is about 310 matchboxes and some colored beads.

The game is tic-tac-toe. – Everybody knows the game. For each matchbox, put a drawing on top showing one of the possible tic-tac-toe board positions with a different colored dot for each legal move possible. Inside the matchboxes, put two colored beads for each colored move position on the box-top. That's it. -- The inventor of this AI-Game, Donald Michie invented it in 1960 – Called it Machine Educable Noughts and Crosses Engine – Yes MENACE. In Britain, Tic-Tac-Toe is called Noughts&Crosses.

So, we have built a MENACE, it is now programmed. Easy to do – So, Let's play Tic-Tac-Toe.

You move and look for the matchbox that has that position on the top. Pick out a colored bead from the box, match the color to a move, and that is Menace's move. Put that bead on top of the box. Now you move. Find that new position on the boxes, pull out a bead. Match it to a move and that is Menace's next move, and put that bead on top of the box. And on and on the game goes. If you win the round, take all of the beads sitting on top of the matchboxes and throw them away. If Menace won the game, then put the colored beads on top of the boxes back into the box WITH one additional bead of the same color also added. In a draw game each colored bead is returned to the matchbox, without any additional ones added.

– Now you can see what is happening! This machine is going to learn to play a good game of Tic-Tac-Toe. – AND it is going to learn just like a person learns. – Success is rewarded – by adding beads, and failure is punished – By throwing away beads. In about fifty games the machine will beat or tie every time. That's it – your own Artificial Intelligence machine on your dining room table.

Think about this simple Learning Machine. It actually is modifying its internal program as it plays more games and gets smarter.

NOTE: This is the same method that evolution of species progresses, successful gene modifications move forward and less adaptive ones fall out of the gene-pool. Like the colored beads in MENACE.

If I played ten games with one of these machines and you play ten games on another identical machine, the internal programming after the ten games would be totally different on the two machines.!!

So, the concept of a Learning machine should no longer be a mystery to you. – You can build one at home.

****

So, what is the current state-of-the-art in Artificial Intelligent machines? Consider the following:

1. In 1997 – Yes, 22 years ago! IBM's Deep Blue defeated the world chess champion, Garry Kasparov. Many smart people said this would never happen.

2. More recently, Google's Alpha-Go defeated the world's best Go player – an accomplishment that many Go experts said would never happen.

3. Seven years ago, IBM's Watson beating the champion Jeopardy players. Watson's ability to rapidly parse **highly idiomatic, Spoken English phrases subject of multiple meanings, in a great variety of subject matter, was astounding.** I believe Watson will serve as a milestone in AI language comprehension.

These are real eye-openers of machine intelligence. These machines were not "programmed" to play winning chess, Jeopardy or Go. NO! – **These machines taught themselves**. How – the same way a human learns. By Trial and error.

By the time the computers were ready to compete at the human level, their programmers would not recognize the code being executed by the computers, and the extent of the resources that had been gathered by the machine. That's what learning machines are all about.

Here are three more recent examples of the advance in AI machines that I believe frame the current state of the intelligence they are expressing. The real difference here is the ability to hear and understand human speech and to compose intelligent, appropriate, well-composed spoken responses.

4. **Write this down.** "IBM Debater." Look it up on YouTube. This AI machine debated a world-champion debate champion. A debate subject was picked, the human and the IBM debater were each given the side of the issue each was to propound, and they each had 15 minutes to prepare. The IBM AI computer was astounding. It put together a cogent debate, with English sentences properly formed with each sentence logically following after the last, and it was prolific at citing studies and sources. The facts and studies IBM-Debater used to support its arguments were better than the human could accomplish.

5. **Yesterday's NY Times:** (Thanks, Betsy!) Four years ago, 700 AI scientists were challenged with an $80K prize money - to create an AI machine that could score well on an 8th grade science test. -- They all failed!! This past Wednesday, The Allen Institute on AI unveiled a system – Called ASRISTO – Using advanced Neural Networks - that consistently scored over 90% on 8th grade science tests and over 80% on 12th grade!!

6. **AI Chatbots** are computers that have developed language understanding to the point where they can intelligently communicate with other machines and with humans in vernacular speech. They haves eyes and ears they take in speech and visual input, and they respond by speaking.

Now – GET THIS. --- An AI-programmer working for a company building such machines came into work one morning and decided to see what, if anything, the computer had been doing overnight when no one was using it. He was amazed when he reviewed the machine's overnight operations to see that the computer had been reviewing its conversations of the previous day and trying to improve on its responses and its understandings and was modifying its resources accordingly. This was being done by the machine without any human intervention telling it to do such things.

*** Would I be anthropomorphizing to suggest that this computer was bored; had no other tasks to do and so it was **just thinking** about it's days work, and trying to do a better job. This is so like our own natural human activity when falling asleep and reviewing your day's activities. Is this approaching Consciousness??

****

So what are the issues confronting us from the coming power of Artificial Intelligence?
I'll identify the major three big discussion topics:

1. The practical issue of jobs and machines that replace humans.
2. The philosophical issue: Will such machines be Conscious?
3. What about War machines and the ethics of AI implementation.

(1) The Jobs Issue: Sometime about 50,000 years ago, our ancestors started to gather together in small settlements, in which people took on labor "specialties." Some members gathered food, others made clothing, some built shelters for others; a medicine-man provided medical care. Some means of exchange was settled upon to balance the trading of good and services.

Our modern world is no different in style, only in scale. We now build one million pound pieces of flying machinery that sell for $200M, and super-computers that can outthink a human at any calculation task imaginable. Our newly minted automobiles are manufactured on half-mile long assembly lines almost untouched by human hands. Is this capability of new AI machines to do most any job now performed by humans going to change the way our societies have functioned for 50,000 years in the way we exchange our labor for goods and services?

How will we employ our citizens if intelligent machines perform most needed production of good and services? Will our governments be organized to receive the income from the computer-manufacturing economy and distribute that income to citizens? On what basis shall such wealth be distributed? Equally? Should a drunkard who spends his money on beer and gambling be given the same amount as given to a person who after years of study, designs and maintains these wonderful AI machines that produce all we need?

Perhaps we are **over-reacting**? – Is this any different than those same fears that rocked society in the Industrial Revolution of 250 years ago? Are we simply 21st Century Luddites, who will be proven to be as wrong as were those folks who smashed British textile mills in 1779, out of a fear that their jobs would all be replaced by modern machinery.

(2) The "Philosophical" Issue: Consider the issue of Consciousness. Just what is it? – How do you define it? And the tough question: Will intelligent AI machines be capable of consciousness? Will AI machines THINK? Well, why not? - If 300 matchbooks can learn to plat Tic-Tak-Tow in the same way as you do, Why Not? Perhaps consciousness is just an emergent property that naturally accompanies any vast intelligence, whether that intelligence is made up of bone and flesh or silicon and germanium.

Philosophers for thousands of years have struggled to comprehend the nature of consciousness and identify its essential properties. Can consciousness can ever be truly explained in terms of **physical** states of the brain. Earlier philosophers proposed a stark distinction between the "Mind" and the "Brain." The brain, being part of the body, was physical, but the "Mind" or "Soul" was some non-physical, supernatural force that was behind human thought. The modern science of neurobiology has long-ago left this "Substance-Dualism" viewpoint to the dustbin of unscientific thought. (Warren might not agree with this!!)

[In recent years, we have made great strides in neuroscience, cognitive science, neuropsychology, and related fields involving the deep inner workings of this three-pound hunk of meat in our heads. We now know much more as to how our physical brain produces our internal subjective view of the world.]

How will we know when an AI computer also has some elements of consciousness?

{The famous British mathematician, Alan Turing, often named as the father of modern computers, for his work in the late 1940's, developed a simple test now known as the Turing test for artificial intelligence. If you can converse with a machine, behind a closed door such that you are unaware as to whether it is a machine or a human to whom you are speaking – Well, then it doesn't matter, the computer just passed the "Thinking Test!" Many modern neuroscientists disagree with this test as being definitive as to true consciousness, but it does give a good start to answering the question.}

If such AI machines do develop such vast mental capacities that we are forced to conclude that they have properties of <u>consciousness</u>, then do they deserve to be treated as thinking human persons? Should the US Constitution apply to the AI machines? Shall Ai Machines not be …"Deprived of Life or Liberty without due process of law"?

AHHHH – Such is the food for the appetite of philosophers – and of course lawyers!

(3) AI, Ethics & War: There is a new acronym in use in the Pentagon – LAWS – Lethal Autonomous Weapon Systems. You like that acronym, HUH!! Well such systems are already here in primitive form and they will be getting more sophisticated and more popular for use in wartime. So far, all "smart" weapons now in use still require a human on the "Kill-Switch." The remotely controlled Predator drones used in Afghanistan need a human to press the button to launch a missile.

Before you condemn the use of these semi-automated weapons, remember the Not-Very-Intelligent weapons used in the last world war. What about fifty B-17's releasing fire-bombs

over Dresden, Germany in 1944, the bombardiers never seeing the persons destroyed by those bombs. Was that OK, but a Predator drone identifying a truck full of Taliban fighters and releasing a missile blowing up the truck is - not OK? - We need to think about that!

I'm not going to attempt to provide any answers to these difficult questions. Not because I don't want to - but because I don't think I have any great insight on solutions. BUT we need to start thinking about it, because the progress of AI is not slowing down. Maybe 15 years, maybe 30 years. Probably not in my short lifetime left, but certainly in the lifetime of your children. The age of Intelligent machines being able to do most any human task will indeed be here. We need to prepare for it.

And another ETHICS Issue – remember the reading I gave of the AI fake-News text generator. That will indeed be present and in use in the upcoming 2020 election. Look for it coming to your favorite news channel soon!

- So, think about it. – Discuss it.

[A spirited polylog followed the service.]