Artificial Intelligence (AI): A Mind of its Own?
First in a two-part series.

September 15, 2019

Earlier in my career, I spent several years as a software programmer. Much of my work was writing code for simulation models. This software attempts to mimic some aspect of the real world and allows its users to post "what if?" questions. The software, based on rules coded into it, then comes up with hypothetical (ideally, realistic) answers.

My work, first with the Department of Defense, was coding models to simulate battles between U.S. Navy destroyers and Chinese submarines. We programmers joked about this assignment. At that time, the Chinese had no submarines.

Jokes aside, it was not an easy assignment. How were we to write code about something that for all practical purposes did not exist? Such was (is) the funding given to the organization in the Department of Defense now called DARPA, for the Defense Advanced Research Projects Agency. (For which, in spite of the Chinese submarine project, I have great respect.)

A few years later, working at the Federal Reserve, I was assigned as the person to write programs to simulate the United States money supply in relation to the Fed's interest rate (the Fed Funds rate). Without going into the arcane aspects of these subjects, I was hired by the Fed because of my "extensive" experience in coding simulation software for U.S. and Chinese naval battles.

This statement comes across as irrelevant now. But in those times (the late 1960s and early 1970s), software simulation models were practically non-existent. I placed on my application to the Fed something to the effect of: "wrote simulation software while in the Navy." I was snatched-up by the Fed after only one cursory interview.

Enough history, which is actually relevant to this article. In those early days, which have persisted until quite recently (a few months ago), programmers could examine the output of their software and compare the output with the software that produced it. Did the output seem correct? If it did not, the programmer could examine the code that produced the output, and make adjustments---fix the "bugs" as they were called.

Transparent Operations, Under Control

The programmer controlled the software. It was the programmer's responsibility to make certain the software conformed to the output it produced. The output was analyzed and evaluated by the users of the software. For my examples: at the Navy, warfare experts; at the Fed, gifted economists.

A Mind of its Own?

Consider this event, which occurred a few months ago, one that is both exhilarating and disturbing (see *The Atlantic*, August 2019, 24-26).

In 2018, the company AlphaZero developed a chess-playing program that was only given the rules of the game, with no input from well-known and accepted chess-playing strategies. It trained itself by self-play to become the best chess player in the world, even beating chess grand masters. The program accomplished this feat, unaided by

humans, in less than 24 hours. It did not use classic chess strategies or practices that had been developed over hundreds of years.

Chess experts considered its moves to be "counterintuitive, if not simply wrong." One of the owners of the company said AI "…is no longer constrained by the limits of human knowledge."

Even more, the programmers of this software could not trace their own code that led to the program's non-traceable output! They could not associate the code of the program with what it had done: defeating all chess players it took on. It was an unfathomable accomplishment. To this former software programmer, it was the opening of another world for us humans.

The second article in this series discusses artificial intelligence experts' views on the subject.

+++

Artificial Intelligence (AI): A Mind of its Own?
Second in a two-part series.

In the first section of this article, I discussed a recently developed software program that plays chess. In less than 24 hours, given only the rules of the game, it taught itself to play with such expertise that it defeated all comers, even chess grand masters. We learned it defied conventional rules, made counterintuitive moves, yet essentially kicked-ass of any challengers. (Forgive the slang. The chess term is *checkmated*.)

Returning to my earlier and now antiquated software background, I trust you can appreciate that this event staggered me. I hope it will stagger you as well.

We humans have created a world of AI software, a world that, with each day, is entering into the mainstream of our lives. AI has fantastic capabilities, especially when used with the massive databases open for access to the public.

AI is also called deep learning. As one expert puts it (*Foreign Affairs*, July/August 2019, 192-198): "Deep learning has its own dynamics, it does its own repair and its own organization, and it gives you the right results most of the time. But when it doesn't, you don't have a clue about what went wrong and what should be fixed."

Do you find this assertion bothersome? I do.

Let's consider the modern-day issue of humans attacking other humans with nuclear weapons. Left to its own, who can guess what AI might be capable of a few years from now? Will we have the wisdom to build AI platforms dealing with weapons of mass destruction (WMD) to control their self-taught logic which, like the chess playing program, will likely exhibit opaque operations?

Will we have the capability to control an AI machine's manipulation of genes as DNA alteration technology progresses by leaps and bounds?

What happens if, as stated above, "…it gives you the right results most of the time." But on those occasions when it does not, then what? We cannot recall the launched missile. We may not be able to undo the genetic alteration. We cannot even determine the code and logic that went into the missile launch or a gene cut-and-insert operation.

If these questions seem far-fetched, keep in mind that AI is in its embryonic stage. Like genetic engineering, it is at the tip of an iceberg.

However it evolves, artificial intelligence will have a transformative effect on the human race. Whether we guide AI, or have AI guide us is, at this time, still up to us.

Do you think I am over-reaching? I ask you to re-read the material about the self-taught chess program cited in the first part of this article.

Also, consider this observation from the experts writing in *Foreign Affairs*: "…no one, not even the people who design these tools, really understands how they work." And: "…AI technologies, at their most advanced levels, do not merely assist human knowledge; they surpass it."

Sleep tight, Chicken Little. Hello, Hal.