

Evolution Versus Revolution in AI: Why the Distinction?

ALAN M. REZNIK, MD, MBA, FAAOS

EDITOR'S NOTE: This is the fifth article in an ongoing series on artificial intelligence (AI) in orthopaedics. Previous articles tackled topics such as the origins of AI, how AI reads X-rays, natural language processing in medical decision-making, and medical ethics related to AI. This article explores the difference between shallow and deep AI, as well as the phenomenon of AI creep. Visit www.aaosnow.org to read more.

The beginning

In the 1800s, Charles Babbage designed a mechanical calculating machine (Fig. 1) that could, by design, solve complex math equations with unheard of accuracy. Ada Lovelace proposed a way of programming algorithms for his design. Of note, due to a lack of funding and machining technology, Mr. Babbage's machine was not built until 2002. When completed, the machine consisted of 8,000 parts and weighed 10 tons. Its hand-cranked action can be viewed on YouTube. Mr. Babbage's machine made him the father of computers, and Ms. Lovelace's prescribed algorithms made her the mother of programming. One could say the original computer revolution was Mr. Babbage's machine, and Ms. Lovelace's programs were the next leap. The original gears have since evolved into microchips in newer machines. The mystery remains where the next revolution or evolution in artificial intelligence (AI) will take us. It may already be here. Quantum computing has the power to perform calculations in 200 seconds that would take the world's most powerful current supercomputer more than 10,000 years, and this processing power could revolutionize the ability of future AI to handle unimaginable quantities of data, while solving the most difficult problems.

What separates Babbage/Lovelace AI from modern AI?

One answer is the availability of vast amounts of memory at a very low cost. In the early 1960s, one megabyte cost approximately \$2 million. Today, one megabyte costs less than one-hundredth of a penny (Fig. 2). In addition to the increasingly affordable memory required for AI, the development of electronic models of neurons has led to neural computing networks. These "electronic nerve" networks can learn from large amounts of data rather than just solve programable algorithms. The

first electronic model of the human neuron, the "perceptron," was designed in the 1950s (Fig. 3), and it has progressed to the modeling of many types of different computer neurons with a variety of unique functions. These electronic neurons have allowed the formation of mathematical versions of neural networks that operate like a small child's brain (Fig. 4)—very pliable and eager to learn. However, this is not the full story, because AI has a preexisting and coexisting sibling that does not require a neural network—shallow AI.

Shallow AI

The main reason we are not concerned about all of the AI devices around us is that they are mostly shallow AI devices. Like Ms. Lovelace dreaming about programming Mr. Babbage's machine, they make a lot of sense to us. We write an algorithm or program to measure time, the clock works, the algorithm does not change, and we take it for granted and trust it.

For example, math nerds know that a body mass index (BMI) calculator gives us the answer to the equation $BMI = 703 \times (\text{weight in pounds}) / (\text{height in inches})^2$. We enter information into an app, and it gives us a number. Most patients have no understanding of how we get to the number. Yet we use the number for medical decision-making without any doubts about its value.

Similarly, we trust end-tidal CO₂ monitoring and pulse oximetry in the operating room. We may not know or understand the many assumptions and approximations used in the calculations, nor do we care. They are often based on the math, chemistry, and physics of the devices we are modeling. Again, the models are tested and proved, and the error margins are known; therefore, they are trusted. Only in special outlier cases do they fail us, and with better technology, such failures are becoming exceedingly rare. Shallow AI is everywhere, and we use it every day without a thought.

AI creep

AI, in the strictest sense (mechanical, electrical, and even quantum-based), processes information in a way that a human might without a human brain doing the work. Everything we touch that has a program in it or that does a task that we would have done with a pencil and paper is AI. This includes




Fig. 1 Charles Babbage invented a calculating machine, known as the difference engine, in the early 1880s.

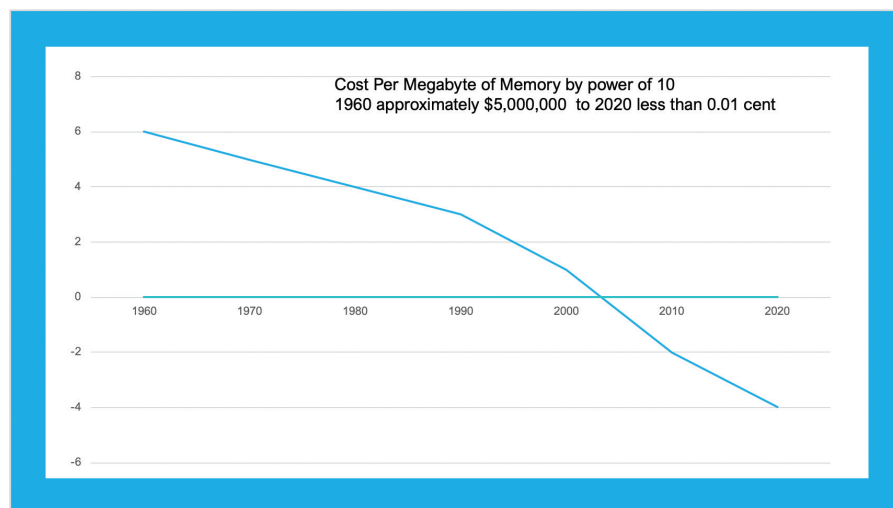


Fig. 2 Cost per megabyte per decade

The cost of megabytes of memory over time in dollars by powers of 10 (e.g., 6 = \$1,000,000; 0 = \$1; -2 = one penny). In the 1960s, the original core memory was made by hand using tiny iron "rings" with two wires, one to flip the magnetic direction for zero or one and the other wire to read the set value. Each ring was hand wired, wrapped, and then placed by hand on a circuit board. It took thousands of workers to make a few megabytes.

COURTESY OF ALAN M. REZNIK, MD, MBA, FAAOS

performing long division on a calculator, making change at McDonald's, plotting a course with a GPS (global positioning system), creating an ultrasound image from sound waves, graphing an electrocardiogram, using a picture archiving and communication system to measure a Cobb angle, and even winding up a simple calendar watch—all shallow AI. It turns out, the more familiar AI becomes, the less we

think of it as AI. Hence, the phenomenon of AI creep: The simpler the device that performs the AI task, the more we rely on it, and the less it seems like the device is doing a significant brain task. In short, AI is everywhere and is evolving every day without much concern.

Deep AI

Deep AI is different. Deep AI is truly like a child who learns from the world

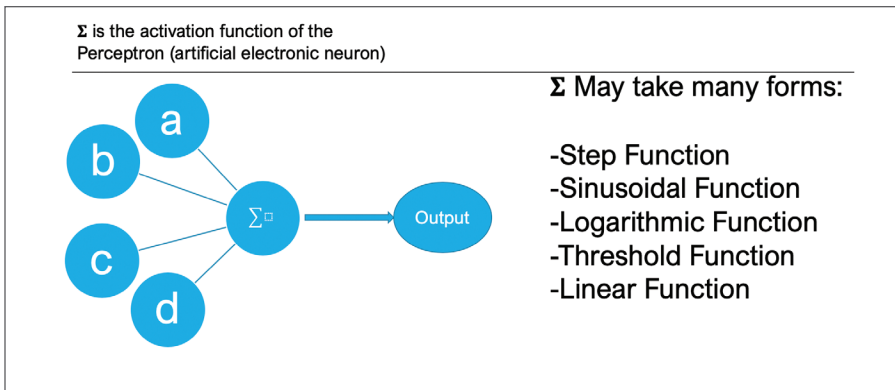


Fig. 3 Perceptron functions

The perceptron as a mathematical model of common neurons in the brain. Each neuron may have a specific function and as such may require a different mathematical model. Examples of the possible models are here. In more complex systems, more sophisticated models are possible, including models with internal mini-neural networks that have local memory in order to resolve the use of more complex inputs.

COURTESY OF ALAN M. REZNIK, MD, MBA, FAAOS




Fig. 4 Simple, three-layer neural network

This is a standard diagram of a simple neural network with an input layer, one hidden layer, and an output layer.

COURTESY OF ALAN M. REZNIK, MD, MBA, FAAOS

around him or her. Sometimes we use “supervised learning”: Humans review the data and point out the information needed. For example, teams of thousands of people identify abnormalities to an AI system every day so it may learn how to “see” a tumor from a gastrointestinal study, read an X-ray, or interpret a command from Siri. We experiment with different deep AI techniques and specialized electronic neurons, configured in layers, like the retina in real life, and see how well these electric brains work (pun intended).

Deep AI assumptions

The data we use to train a network represent all population data fairly. It has a cross-section of normal and abnormal findings that match the “pathology” of interest. The presentation of the data is the same in the unknowns to be examined later. A past dataset represents the future population data, which will not change with time. The actions of the AI itself on the population will not change the population data in the future.

To “trust” an AI system when making

medical decisions, we need to know whether our trained network can ignore unimportant differences in the data and still function properly. In addition, as discussed in a previous article in this series about AI ethics, we need to ensure that these technologies do not propagate bias or disparities in health care. Some of those challenges are still being solved. Lastly, the method of collection of the data must have internal consistency with the collected training set. For example, when training systems to read X-rays, different machines may have different background electronic “noise” based on power supply, age of the X-ray beam, and the quality of the image-collecting device. Such differences produce many types of noise and can throw an AI engine off when test cases are not from the same equipment as the original training set.

The FDA’s role in medical AI

The Food and Drug Administration (FDA) assesses devices with fixed algorithms in one way, improvements in algorithms in another way, and learning devices in a third way. Ultimately, a device’s ability to learn as it goes

Software as a medical device

The 510(k) software modifications guidance focuses on the risk to users/patients resulting from software changes. Categories of software modifications that may require premarket submission include a change:

- that introduces a new risk or modifies an existing risk that could result in significant harm
- to risk controls to prevent significant harm
- that significantly affects clinical functionality or performance specifications of a device

When applied to artificial intelligence (AI)/machine learning (ML)-based software as a medical device (SaMD), the above approach would require a premarket submission to the Food and Drug Administration when an AI/ML software modification significantly affects device performance, safety, or effectiveness; the modification is to the device’s intended use; or the modification introduces a major change to the SaMD algorithm. For a Parts Manufacturer Approval-approved SaMD, a supplemental application would be required for changes that affect safety or effectiveness, such as new indications for use, new clinical effects, or significant technology modifications that affect performance characteristics.

To address the critical question of when a continuously learning AI/ML SaMD may require premarket submission for an algorithm change, we are prompted to reimagine an approach to premarket review for AI/ML-driven software modifications. Such an approach would need to maintain reasonable assurance of safety and effectiveness of AI/ML-based SaMD, while allowing the software to continue to learn and evolve over time to improve patient care.

References for the information cited in the sidebar can be found in the online version of this article, available at www.aaosnow.org.

has a higher concern for future patient safety than fixed algorithms. See the sidebar for differences between shallow and deep AI devices. Even so, the FDA may require a trained network to use the validated initial training without changes. The network cannot “improve” itself unless patient safety is protected. For now, a true, FDA-approved AI neural network is frozen in that state of its trained function until a new FDA application is made.

The future

As we continue the revolution of AI and watch the evolution of shallow and deep AI, we need to be cognizant of AI creep. We should be mindful of assumptions in using current training datasets as predictors of future population conditions. An algorithm that makes calculations like Cobb angles and a BMI app are relatively safe because doctors usually make final diagnoses and treatment plans. The mathematical formulas do not change over time. The shallow AI systems are not making absolute diagnoses or direct treatment decisions.

However, deep AI has the promise of doing both the “math” and the decision-making. As a positive, deep AI may be more consistent than humans in applying diagnostic concepts to the vast amounts of data that now exist, which may be beyond the human scope to process in clinical decision-making. AI also may uncover relationships in data far beyond what

humans can imagine. Still, because of self-learning, how a self-trained neural network of many layers of differing perceptrons would make medical treatment decisions may remain a mystery to doctors and patients—an AI black box, if you will.

Understanding the black box may become much more important in the near future. The FDA and European law may require all medical AI to explain itself—explainable AI, or XAI. Designing AI neural networks that can explain how they make medical decisions is a possible solution. Deep AI has no inherent language or syntax, and making it is no small task. Understanding the difference between shallow and deep AI, as well as XAI, is a significant challenge for the future of AI technology and will be the subject of the next article in this series.

References for the studies cited can be found in the online version of this article, available at www.aaosnow.org.

Alan M. Reznik, MD, MBA, FAAOS, specializes in sports medicine and arthroscopic surgery and serves on the AAOS Now Editorial Board. He is also a former member of the AAOS Communications Cabinet, as well as the Committee on Research and Quality. Dr. Reznik is chief medical officer of Connecticut Orthopaedics, is an assistant professor of orthopaedics at Yale University School of Medicine, holds multiple patents, and consults on new AI applications in medicine.