Accepted Manuscript

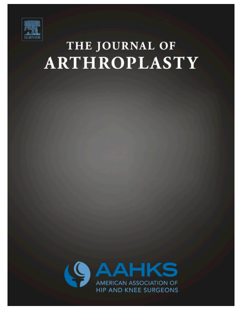

Artificial Intelligence, Machine Learning, Deep Learning, and Cognitive Computing: what do these terms mean and how will they impact health care?

Stefano Alec Bini, M.D.

PII: S0883-5403(18)30215-8

DOI: 10.1016/j.arth.2018.02.067

Reference: YARTH 56463

To appear in: *The Journal of Arthroplasty*

Received Date: 12 February 2018

Accepted Date: 14 February 2018

Title Page

Artificial Intelligence, Machine Learning, Deep Learning, and Cognitive Computing: what do these terms mean and how will they impact health care?

Stefano Alec Bini, M.D.
University of California, San Francisco
500 Parnassus Avenue, MU 323-W
San Francisco, CA 94143-0728

Tel: 415.476.3900
Fax: 415.476.1304

E: Stefano.Bini@ucsf.edu

Artificial Intelligence, Machine Learning, Deep Learning, and Cognitive Computing: what do these terms mean and how will they impact health care?

Abstract

This paper was presented at the 2017 annual meeting of the American Association of Hip and Knee Surgeons to introduce the members gathered in the audience to the concepts behind Artificial Intelligence (AI) and the applications that AI can have in the world of health care today. We discuss the origins of AI, progress to Machine Learning (ML) and then discuss how the limits of ML lead data scientists to develop Artificial Neural Networks and Deep Learning algorithms through biomimicry. We will place all these technologies in the context of practical clinical examples and show how AI can act as a tool to support and amplify human cognitive functions for physicians delivering care to increasingly complex patients. The aim of this paper is to provide the reader with a basic understanding of the fundamentals of Artificial Intelligence. Its purpose is to demystify this technology for practicing surgeons so they can better understand how and where to apply it.

Gartner is a research and advisory company that publishes the yearly “technology hype cycle”[1,2]. Spend any time in the world of digital health and you’ll see the term “hype cycle” used frequently. Essentially the concept describes what happens after new technologies enter the public sphere and generate excitement and hyperbolic claims of “disruption”. The expectations and fervor that are generated are often unmatched by the initial capabilities of the technology. As described by the Gartner researchers, the stage of “innovation” yields to that of “inflated expectations”. Soon thereafter, when the technology does not meet hyped-up expectations, it falls into the “trough of disillusionment” (otherwise known as the “valley of death”). With time, the technology matures, achieves its promise, yields to the “phase of enlightenment” and eventually to the “plateau of productivity”. For most startups, the challenge is to raise enough money during the peak of “inflated expectations” to survive through the “valley of death” and make it to the “plateau of productivity”.

Currently, the technologies which are the top of the Gartner Hype Cycle are all basically associated with Artificial Intelligence (AI)[1]. The hype is so high, that some are suggesting that AI will be the harbinger of doom for humanity and call forth a dystopian world view where machines run the planet [3,4]. Others offer a more optimistic view with millions of new jobs and greater economic growth spurred by smarter decision making [5]. Laying all that aside, there are many areas in healthcare where we are already seeing a positive impact from this technology.

So: what is AI? There are multiple definitions for artificial intelligence. One definition from Wikipedia.com is that it is human intelligence exhibited by machines. In computer science, the field of AI research defines AI as the study of “intelligent agents” which are devices that “perceive their environment and take actions to maximize their chance of success at some goal”[6]. The “and” part of the definition is an important distinction.

There are many examples of artificial intelligence in our lives. Apple's Siri is one such example. Another is Amazon's Alexa. Natural Language Processing (NLP) technology, a form of AI, is used to translate languages in Google Translate. Indeed, up to $30B has been invested in AI in the past 5 years and 90% of it on research and development by companies such as Google and Microsoft [7]. Because there is a great deal of interest in this technology and its applications in healthcare, many hospital systems and individual practitioners of being approached by vendors who claim to be using AI in their products. In some of these instances, this practice is called "AI washing", the idea of adding the label of AI to all and every software platform. Buyers should therefore beware as, often, what is being sold as AI is nothing more than a basic algorithm.

True artificial intelligence has great applications in health care because it can handle and optimize very complex data sets residing in very complex systems. Caring for patients requires controlling many steps, each of which is highly variable, and each of which is dependent on or connected to multiple other steps. Further, these steps involve both machines and people in processes that are more stochastic and less deterministic than a traditional assembly line. To manage such variation one needs a predictive and centralized command and control system that can handle such complex data and learn continually from its experience by optimizing (read: re-write) the very algorithms it uses to deliver recommendations. While tracking so many variables is challenging for people (see for example the many reports of wrong medication administration or even wrong side surgery), it is something that computers are particularly adept at.

Some examples of where we are today with AI in health care include work done at UC Health in Colorado where an AI based scheduling tool was used to optimize surgical schedules[8]. The software evaluated historical data and future demand to create recommendations around OR block allocations. The OR increased revenue by $15M (4%) by increasing the number of blocks released by 47% and doing so 10% sooner. Further, they absorbed six new surgeons by optimizing but not increasing their total staffed OR block allocation.

Another example involves New Yok Presbyterian Hospital's 30 Bed Infusion Center[9]. By using one of several commercially available AI solutions (www.leentaas.com, Santa Clara CA) to optimize schedules and templates, their patient wait-times dropped by 50%. Sharp Health Care in San Diego wanted to decrease the time it took to transfer an emergency room patient to the ward. They used the same platform to forecast the need for admissions, provided visibility to workflows, and prioritize which patients to discharge first from which unit and which room to clean first. In so doing, they decreased the time from the admission order to patient transfer by three hours.

Although these examples may be a far cry from the humanoid run cyber-worlds of many Science Fiction movies, these examples are real, pragmatic and effective health care applications of Artificial Intelligence in place today.

AI has been around for a while. The term was coined by John McCarthy at a lecture at Dartmouth College in 1956. Since then, the idea of sentient computers has captured and terrorized our collective imagination. When the HAL 9000 computer on board Discovery One, the spaceship in Stanley Kubrick's movie *2001: A Space Odyssey* decides to sacrifice human life to ensure the success of its mission, fear of AI and computers was seared deep into our collective psyche.

And yet, it has taken quite some time for the technology to catch up with the promise of AI and the hopes of computer science specialists. Moore's Law, defined in 1965 by Intel co-founder Gordon E. Moore, predicted that the number of transistors on integrated circuits would double approximately every two years [10]. The law so far has held true and created an exponential increase in computing power that is hard to comprehend. In 1971, the Intel 4004 was proud of its roughly 2300 transistors, by 1989 Intel's 80486 chip had 1,000,000 transistors, by 2010 an Intel® Core™ processor with a 32 nm processing die and second-generation high-k metal gate silicon technology held 560 million transistors. In 2016 the fastest chip in the world had 10 billion transistors [10].

The price of computing has similarly plummeted. In today's dollars, a computer that could perform comparably to a $100 iPad2 in 2010 would have cost $10B in 1970. Similarly, if we normalize calculations per second to 10 for a computer built in 1970, today the computing power is 10,000,000,000 calculations per second [11].

The problem we have as humans is that while we are relatively comfortable understanding linear growth, we have a hard time getting our head around exponential growth. We look back at the pace of change in our lives and believe that we can understand and adjust for it. Until one day we realize that the amount of change we experienced over one decade just occurred in one year and that, looking to the future, a decade of change will occur over just a few months. That is difficult to accept, never mind to understand[11].

Exponential growth is what is expected from AI in the coming years[5]. Yes, it is at the top of the Gartner Hype curve, but many people believe it is with good reason. And while it is important to note that we are still a long way away from the Sci-Fi version of cyborgs that are human in all but physiology, many AI applications are coming to healthcare which will have profound impact in the care that is delivered and how it is delivered.

To understand where AI is going, let's first go back to one of the earliest forms of AI, a piece of software developed at IBM by Arthur Samuel that could play checkers autonomously after gathering data and looking at all available options prior to making a decision. For it to work, every possible move in checkers had to be programmed into the algorithm. Only then could the computer decide which amongst the options available it should choose. Fortunately, checkers was a relatively simple game with a limited number of options. The programmers were able to write the code for each possible option and the hardware could handle the computational load. It would take a while before computers could handle the complexity of games like chess. But to handle games like Go, which can have more possible permutations than there are atoms in the known universe (that's a lot of options), an entirely different form of AI had to be invented, namely "deep learning". But we will come to that later.

Let's consider Machine Learning (ML) first. We had to climb a long way along the technology curve before ML was even possible. ML is best considered as a subset of AI. ML learns from experience and improves its performance as it learns. As we saw in the earlier examples, it is a field which is showing promise in helping to optimize processes and resource allocation. Machine learning works as follows: let's say we want to teach a computer how to recognize the species of an Iris[12]. First, the programmers figure out which features are relevant to the various types of Iris. Because over the past several centuries botanists have figured out exactly what differentiates one Iris from another (petal length, width etc.), the programmer can create a long table of known flowers (each flower would be considered an "instance") and list their characteristics (or "features" such as petal length and width, image histogram, color distribution, unique color counts, etc). This data set is called the "training data set" or "ground truth". After "looking" or "learning from" this training set of say 150 high quality images, the software will have learned what combination of attributes are associated with any type of Iris. The magic occurs when software is shown a picture of an Iris it has never seen before and accurately recognizes which species it belongs to. The accuracy goes up with the number and size of the training set (which is why everyone is so excited about "big data": the bigger the data set of known variables, the more accurate the software can become when presented with unknown variables). Further, if trained through feedback loops (right decision/wrong decision) it can adjust its own algorithm and "learn". The software literally re-codes itself.

Probably the best-known example of Machine Learning software in healthcare is IBM's Watson Health. Watson Health has been fed everything ever written in any language at any time related to cancer diagnostics and treatment. Not only that, it continues to ingest all new data as it is published. When presented with a specific patient with cancer, Dr. Watson will recommend the treatment trial most likely to cure that individual patient's cancer by considering their genome, history, imaging, and pathology; coupled to all the information known about the treatment of that cancer. The more information Watson has about the patient, the more accurate it will be. Basically, IBM's vision for AI is that it will support physician decision making rather than supplant it by doing those things that it does best: manage and consider massive amounts of data and present only the relevant information to the physician[13,14]. By knowing all the data and sorting through it at incredible speed, AI can make the physician smarter no matter where they are geographically and what resources are available to them. This application of AI, where AI is in a support role, is referred to as Cognitive Computing [15]. It should be noted here, that IBM Watson Health has had some challenges such that its partnership with the MD Anderson Clinic came to an end in 2017 due primarily to cost overruns and challenges with migrating data across electronic health records[16]. It's not all perfect yet.

Note that, in our ML example about flowers, the programmer must know which features are relevant or important to define what a flower "is". Further, the "ground truth" in the data must be accurate for the software to work. If a *versicolor* Iris was consistently and inaccurately labeled as a *setosa*, the algorithm will initially be inaccurate. Furthermore, with machine learning the programmer must get the data to the computer in a form that it can understand. That usually means converting data into a large spreadsheet consisting of many rows and columns each holding unique numbers or values (this is called "structured data", something that is relatively scarce in Health care where so much of the information is in "unstructured" data sets consisting mostly of clinical chart notes). Since designing and selecting the features is very time consuming, software engineers building machine learning algorithms need to be smart and only extract the features (from the data) that can improve the model. But since in most real-world scenarios engineers don't know which features are useful until they train and test their model, developers can get into long development cycles where the team has to identify and develop new features, rebuild the model, measure results, and repeat the cycle until the results are satisfactory. This is an extremely time-consuming task and while, over time and with enough new data and training, the algorithm can become increasingly accurate, the process can take a long time. Machine learning is thus challenged when dealing with data sets with multiple dimensions where extracting the most predictive features is not obvious or where there are large numbers of both inputs and outputs that need to be addressed.

To handle large data sets and unstructured data without clear, known, features computer scientists turned to biomimicry for inspiration; i.e.: they copied nature. Thanks to the advent of increasingly powerful computer chips and microprocessors, scientists have created statistical models called Artificial Neural Networks (ANNs) that can process data inputs much the same way as the human brain does: by assigning a logical construct to the information. This form of AI, "the state of the art of the state of the art" as some call it, has been around conceptually for some time (basically since the 70's) but in a very limited fashion (2 layers of ANNs at most) due to the limits of computational power. When plotted out, ANNs look very much like neurons with dendrites going in many directions and connecting to more neurons, each of which has more dendrites connecting to more neurons etc. The idea is that each node can accept an input and then store some information about it before passing the information up to the next level. Thus, each layer has an increasingly complex "understanding" of the information it receives than the previous layer. Although the concept was there[17], it was not until the early 2000's with the introduction of the powerful NVIDIA chips that neural networks could to be stacked one on top of the other to create more than one or two connections and achieve their full potential. Today, some engineers are creating ANN networks that are as many as 100 layers deep, hence the origin of the name "deep" learning for these algorithms. With so many layers, the ANNs can tackle and master increasingly complex data[18].

So: why is that interesting? Because these algorithms can be shown raw data and extract their own features without human input. Shown a large enough number of flowers, they can identify the features that define each species without being told what those features are. They can differentiate one face from the next. They do not need structured data sets to learn from. It's pretty amazing: they learn very much like children.

But what does that actually *mean*? It means that if we want a deep learning algorithm to learn how to recognize and differentiate faces in a photograph we don't have to create a massive spreadsheet with features like "nose" and "eyes" and how these (and hundreds of other features) differ from one person to another as we would for a traditional machine learning algorithms. Indeed, Google's algorithms can use deep learning based facial recognition to identify the same person (Uncle Bob) in two photographs taken a decade apart without any human input. And although it was never told that Uncle Bob is the one with the blond hair and crooked nose, if you identify that person as "Uncle Bob", the software can go and find him in all your other pictures (or any on the web).

But AI can go many steps further. Google DeepMind's Deep Q-learning software can learn how to play Atari's Breakout Video game without being told anything other than to maximize the score [19]. It is not told what a ball or a bat is. It has no concept of tennis or ball games. And yet, after a series of trial and error games from which it learns that to maximize the score it must hit the "ball" with a "bat" to knock out some targets, in a matter of minutes it can play at the level of a very skilled person and achieve very high scores. And if that were not enough, it can take those learned skills and apply them to other video games. In May of 2016, Google's program called AlphaGo beat the Korean GO master at the most complex game known to man [20]. However, as with humans, the software is only as good as the data it is trained. A person trained as a lawyer is not likely to be a great physician or a good plumber. Similarly, the Atari playing software Deep Mind would not be all that great at face recognition, and I doubt AlphaGo is all that great at selecting the right medications for patients.

There are applications for deep learning algorithms in healthcare that are quite interesting. For example, the FDA now allows pharmaceutical companies to model drug interactions using AI algorithms that accurately identify drug toxicity without the usual decade of animal testing [21] In fact, they encourage it. Cancer detection in pathology slides is getting very close to outperforming trained pathologists and similar pattern recognition software is getting quite accurate at reading radiographs [22–24]. The reader is encouraged to search on-line for the latest advances in this space because undoubtedly what may seem earth shattering at the writing of this paper may be old hat by the time it is being read.

With that caveat in mind, what does the immediate future promise with respect to AI? As mentioned above, on the clinical front pattern recognition will improve the accuracy of image based diagnostics and AI platforms will provide decision support to clinicians (and patients and insurers). Analytics based on historical data will be deployed by hospitals and clinics to optimize workflows and resource allocations. And while the impact on healthcare of massive new amounts of data acquired by patients from their interaction with the Internet of Things (otherwise known as the IoT, this term refers to the web of digital data collected by the world around us from our cars to our refrigerators, from our browsers to surveillance cameras at the mall, from our DNA sequence to our Fitbit data) has yet to be fully appreciated, there is no question that it will be an AI engine trained on that data that will provide the insights. Artificial Intelligence may be at the peak of the Gartner Hype Cycle, but it is unlikely to fall deep into the Trough of Disillusionment in the health care realm before coming out to the Phase of Enlightenment and into the Plateau of Productivity. And, by that time, it will be the new "normal".

Bibliography

[1] Panetta K. Top Trends in the Gartner Hype Cycle for Emerging Technologies, 2017. Gartner 2017:1–5.

[2] Gartner Inc. Gartner Hype Cycle. Reproduction 2011;9:240–52. doi:10.1016/j.accinf.2008.09.001.

[3] Gibbs S. Elon Musk : artificial intelligence is our biggest existential threat. The Guardian 2014:2014.

[4] Rory Cellan-Jones. Stephen Hawking warns artificial intelligence could end mankind. BBC News 2014:1–5.

[5] Purdy M, Daugherty P. Why AI is the Future of Growth. 2016. doi:10.1016/j.techfore.2016.08.019.

[6] Poole DL, Mackworth A, Goebel RG. Computational Intelligence and Knowledge. Computational Intelligence: A Logical Approach 1998:1–22.

[7] Bughin, Jacques, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlström NHMT. McKinsey Global Institute: Artificial Intelligence The Next Digital Frontier? 2017.

[8] LeenTaas. U Colorado OR Optimization 2017. https://leantaas.com/uploads/booklets/iQueue_for_Operating_Rooms_UCHealth_Results_Handbook.pdf.

[9] LeenTaas. Infusion Center Optimization 2017. https://leantaas.com/infusion_centers/case_studies_booklet.

[10] Rupp K. Moore’s Law: Transistors per microprocessor. Ourworldindata.org 2017. https://ourworldindata.org/grapher/transistors-per-microprocessor (accessed November 2, 2018).

[11] kurzweill ray. The law of accelerating returns. Nature Physics 2008;4:507. doi:10.1038/nphys1010.

[12] True J. Watson Analytics Use Case: The Iris Data Set 2015. https://www.ibm.com/communities/analytics/watson-analytics-blog/watson-analytics-use-case-the-iris-data-set/ (accessed February 11, 2018).

[13] Watson Health. 5 Ways Cognitive Technology Can Help Revolutionize Healthcare - Watson Health Perspectives. Watson Health 2016. https://www.ibm.com/blogs/watson-health/5-ways-cognitive-technology-can-help-revolutionize-healthcare/.

[14] Somashekhar SP, Sepúlveda M-J, Norden AD, Rauthan A, Arun K, Patil P, et al. Early experience with IBM Watson for Oncology (WFO) cognitive computing system for lung and colorectal cancer treatment. Journal of Clinical Oncology 2017;35:8527. doi:10.1200/JCO.2017.35.15_suppl.8527.

[15] Ahmed MN, Toor AS, O’Neil K, Friedland D. Cognitive Computing and the Future of Health Care Cognitive Computing and the Future of Healthcare: The Cognitive Power of IBM Watson Has the Potential to Transform Global Personalized Medicine. IEEE Pulse 2017;8:4–9. doi:10.1109/MPUL.2017.2678098.

[16] Mole B. {IBM}’s {Watson} proves useful at fighting cancer—except in {Texas}. Ars Technica 2017.

[17] Hopfield JJ. Artificial neural networks. IEEE Circuits and Devices Magazine 1988;4:3–10. doi:10.1109/101.8118.

[18] Schmidhuber J. Deep Learning in neural networks: An overview. Neural Networks 2015;61:85–117. doi:10.1016/j.neunet.2014.09.003.

[19] Mnih V, Kavukcuoglu K, Silver D, Graves A, Antonoglou I, Wierstra D, et al. Playing Atari with Deep Reinforcement Learning Volodymyr. arXiv Preprint arXiv: … 2013:1–9. doi:10.1038/nature14236.

[20] Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, et al. Mastering the game of Go with deep neural networks and tree search. Nature 2016;529:484–9. doi:10.1038/nature16961.

[21] Aliper A, Plis S, Artemov A, Ulloa A, Mamoshina P, Zhavoronkov A. Deep learning applications for predicting pharmacological properties of drugs and drug repurposing using transcriptomic data. Molecular Pharmaceutics 2016;13:2524–30. doi:10.1021/acs.molpharmaceut.6b00248.

[22] Fakoor R, Ladhak F, Nazi A, Huber M. Using deep learning to enhance cancer diagnosis and classification. Proceeding of the 30th International Conference on Machine Learning Atlanta, Georgia,USA 2013;28.

[23] Wang D, Khosla A, Gargeya R, Irshad H, Beck AH. Deep Learning for Identifying Metastatic Breast Cancer. arXiv Preprint 2016:1–6.

[24] Cruz-Roa A, Gilmore H, Basavanhally A, Feldman M, Ganesan S, Shih NNC, et al. Accurate and reproducible invasive breast cancer detection in whole-slide images: A Deep Learning approach for quantifying tumor extent. Scientific Reports 2017;7. doi:10.1038/srep46450.