

## Artificial Intelligence: What’s in it for medical professionals?

<sup>1</sup>Narumanchi, T., <sup>1</sup>Narumanchi, J., <sup>2</sup>Boots, T. and <sup>3,4,5</sup>Naveena C. Yanamala\*

<sup>1</sup>West Virginia University School of Medicine, Department of Pediatrics, 1 Medical Center Drive, Morgantown, WV – 26506; <sup>2</sup>BioAnalytics Branch, Health Effects Laboratory Division, National Institute for Occupational Safety and Health, 1095 Willowdale Road, Morgantown, WV 26505; <sup>3</sup>Exposure Assessment Branch, Health Effects Laboratory Division, National Institute for Occupational Safety and Health, 1095 Willowdale Road, Morgantown, WV 26505; <sup>4</sup>Institute for Software Research, School of Computer Science, Carnegie Mellon University, 5000 Forbes Avenue, Wean Hall 5216, Pittsburgh, PA – 15213; <sup>5</sup>West Virginia University School of Medicine, Heart and Vascular Institute, 1 Medical Center Drive, Morgantown, WV – 26506. (\*Correspondence: Dr. Naveena Yanamala, Exposure Assessment Branch (MS-3030), 1095 Willowdale Road, Morgantown, WV 26505 USA. Phone: (304) 285-6313. Fax: (404) 638-5515. E-mail: naveena@andrew.cmu.edu)

Medical professionals struggle to cope with a daily flood of new information and technology. Computer technology is maturing to help manage and improve our ability to care for patients. Artificial Intelligent systems with their problem-solving approach based on pattern recognition are one such example in the medical field. Their intelligent architecture, which incorporates learning and the ability to act autonomously is alluring; ideally allowing us to offload the day to day drudgery and spend more time with our patients. Unfortunately, there is an ingrained fear that AIs will not just help us but might one day replace us. This review serves as a primer for clinicians to better understand the basis for the new technology coming our way and distinguish what it can and cannot do.

**Keywords:** Deep Learning, Machine Learning, Expert systems in medicine, Medical decision-making, automation in healthcare.

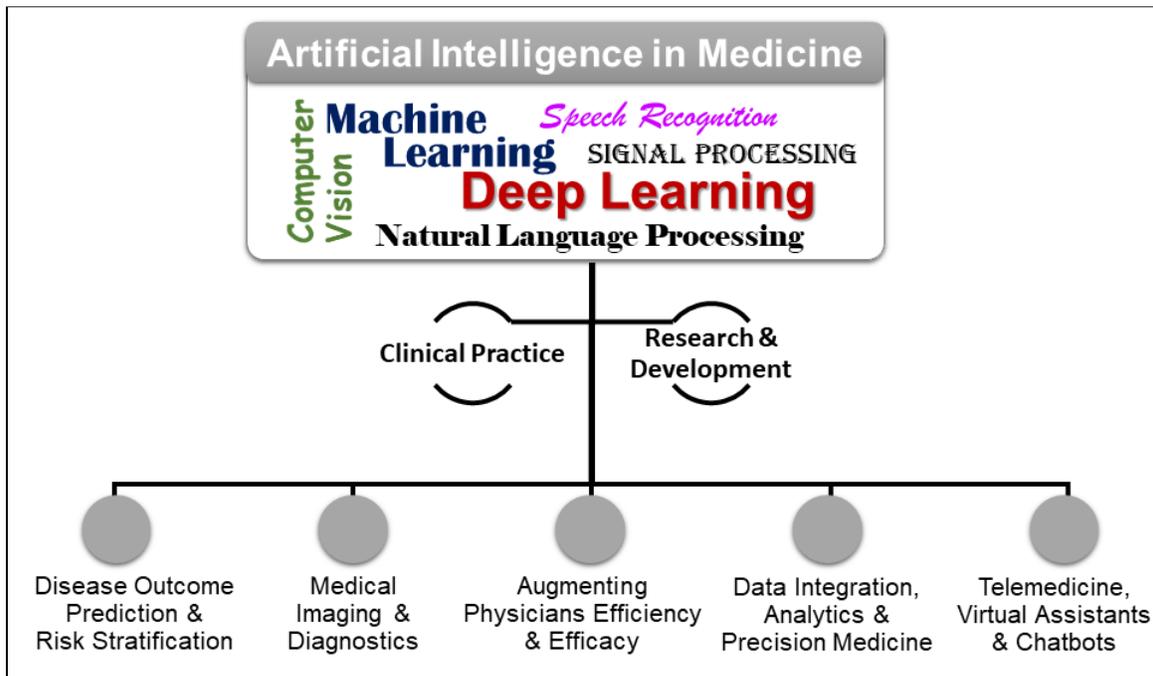


Figure 1. Potential application areas and role of Artificial Intelligence in Medicine.

As we were preparing to write this paper, one of our authors remembered the advice he had been given by his surgery professor: Plan to start your day with a healthy breakfast and a stack of journals, expect to skim through 300 pages a day, and read at least 30 pages in detail. We can guess what most of you are thinking at this point: “When was the last time I actually had time for a leisurely breakfast, much less a chance to read my journals? I spend most of my time dictating, reviewing results, or trying to make sense of the latest dictate from the Dark Lords in the Administrative suite. It just gets worse when I am actually at work, as I am expected to do more with less support.”

Nevertheless, there is hope. Technology is starting to advance into what would once have been reserved for science fiction. Personal computers and cloud computing can now do tasks that were once strictly limited to supercomputers hidden in the dim recesses of the department of defense (DOD). Software is beginning to take over or assist with tasks, once delegated to our assistants or that require us to depend on other experts (e.g., radiologists, pathologists, etc.).

This review is meant to give the reader a primer to the concept of Artificial Intelligence (AI), the theories behind AI programming such as machine learning, neural networks, pattern recognition and natural language processing, and briefly introduces various real-world applications of AI in modern medicine. It also summarizes what we need to watch out for when relying on automated systems i.e., the limitations of AI-based models (or as one of us put it - “Don’t dream about being put out to pasture, they’ll always need us to actually see patients when they’ve installed the next upgrade”). For the purposes of this paper, we will refer to everything as AI, regardless

of the underlying methods, software and architecture that powers it.

**Artificial Intelligence.** Humans are considered to have “natural intelligence”, making them self-aware. The general purpose of AI is to create software and computer programs that allow computers to perform tasks or functions that normally would require human intervention. Fundamentally, the goal is to create a machine that can interact with its environment, learn, and complete the tasks it has been assigned. In 1950 Alan Turing wrote his seminal paper on Machine Intelligence (1), in which he proposed that if an impartial judge reviewing a conversation between a computer and a real person cannot tell one from the other, the machine is “intelligent”. However, “intelligence” does not constitute awareness. A simple example of this is Optical Character Recognition (OCR). Most scanners come with a software package that can recognize the letters on a typed page and convert them into an editable format. With training, some packages can convert handwritten notes - but that training requires a real person to verify and teach the software what the various squiggles mean.

When someone refers to AI, people usually associate the term with a humanoid robot. However, the term is originally credited to John McCarthy and the Dartmouth project - a 1956 project at the Dartmouth College based on “the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it” (2).

AI is already pervasive in many parts of our day-to-day lives. It is the magic behind autocomplete and identifying our preferences so that we are presented with that perfect online purchase or a favorite video. However, AI’s application to the field

of medicine is relatively recent and still a work in progress. The idea of utilizing AI in medical applications was first presented in the 1960s, but did not begin to be successful in classification or diagnosis (aka solving “practical problems”) until the 1990s (3-5). In the last five years, the explosion of data and advances in computational methods have led to an exponential increase in research on these practical applications in medicine. But even with recent advances in the field it still has several limitations that require attention (6). Thus, while AI is not going to replace you at your workplace, it can enhance a doctor’s performance by supporting the never-ending administrative and repetitive tasks, aid in making faster and more accurate diagnoses, early detection of disease or even help in predicting the risk of a disease in time to prevent it.

*Explain some of those buzzwords: AI hot terms:* These days phrases like: AI, deep learning, machine learning, neural networks and much more get tossed around in various settings such as grand rounds, conferences and other professional gatherings. Let us have a quick glance at what they are really talking about.

*AI vs. Machine Learning (ML):* AI is a popular buzzword that is seen throughout the news but is often used as a blanket term that encompasses a variety of different methods and applications. AI actually refers to machines performing cognitive-like tasks analogous to human intelligence (e.g., playing chess, autonomous driving, speech/image recognition), while machine learning - another hot term among academics and researchers - refers to the application of statistical models and mathematical algorithms to perform these tasks, which typically includes pattern recognition to make inferences.

*Machine Learning: Supervised, Semi-supervised and Unsupervised:* Supervised ML tasks are typically considere

-d to be the “best-case” scenarios for prediction and classification of problems as the ML algorithm is “trained” on a set of data that has been fully labeled by a human expert - learning patterns in the data that pertain to each label and enabling it to apply these patterns to new data (7). Semi-supervised ML is closely related to supervised ML tasks; however, they do not have the luxury of being fully labeled, with usually only a small fraction of the data being labeled. Unsupervised ML tasks are the counterpart to supervised ML in that the entire dataset is unlabeled (7). These may not seem useful at first, but these algorithms are pivotal towards searching for overall patterns and natural clustering in data that may not be obvious to the human mind.

*Neural Networks & Deep Learning (DL):* Neural Network is another obscure term tossed around quite a bit. In general, this term refers to a computer system modeled to function and process information similarly to neuronal connections in a human brain. However, there are several different classes of neural networks - artificial neural networks, recurrent neural networks, convolutional neural networks, etc. - each with a unique structure and function that attempts to mimic different types of thinking and memory. Each of these architectures can have one or several layers of processing for transforming input data to predict output results. When there are multiple layers that occur within a neural network, it is often referred to as “deep learning”.

### **AI in Medical Data analysis & Diagnosis:**

How can we convert the large influx of raw data into actual improvements in patient outcomes? The use of AI for processing and effectively utilizing large datasets has expanded throughout the medical field over the last decade (Figure 1). AI has shown to be very beneficial in the medical specialties

that involve identifying a ‘pattern’, whether it is an image, scan, sound wave, genetic information, speech, or text (e.g., pathology, radiology, dermatology, cardiology, ophthalmology and oncology). In most of these cases, DL and ML can enhance a physician’s performance in two ways: (a) seeing the data and patterns in ways that are close to impossible to identify or visualize; and (b) quickly sifting through large datasets to identify new knowledge from those patterns and trends. Below are some examples of how different AI techniques are being used in this new era of medical “big data” (e.g., omics, electronic health records (EHR) and image analysis) (Figure 1).

*Feature selection: Finding Importance in the sea of big data.* Many datasets in the medical field are plagued with the curse of high dimensionality, where the number of variables (or features) greatly outnumbers the number of samples or subjects. When a dataset has high dimensionality, the need to find the features of most importance or that carry the most information can be key. This can not only reduce the time and cost of running the algorithm, but also significantly increases the accuracy (8). Filtering through the massive set of variables that often accompany medical datasets is a step that is utilized in almost every application of AI in medicine. This is true, especially for ‘omics data’ (e.g., genomics, epigenomics, proteomics, etc.) where a single dataset could contain hundreds of thousands of gene/protein feature sets or genetic sequences from only a few dozen different samples. For example, in Personalized Cancer therapy, researchers have developed a method for taking the hundreds to thousands of mutations found in a tumor sample and combine that with a patient’s ability to metabolize drugs to determine the best combination of medicines that could lead to remission (9).

*(Semi-) Supervised Learning: Disease vs. Non-disease classification.* Supervised ML refers to algorithms that are trained using a dataset in which all samples have a corresponding group label, such as “healthy” versus “disease” or tumor type (7). These types of classification algorithms are currently the most prevalent ML tactics being used in the field of medical research, as the primary goal of many studies is the ability to differentiate and predict diseased states (Figure 1). ML classification methods have been successful in a variety of oncology applications, from susceptibility of an individual to the prediction of recurrence or survival of a patient (10). These methods have also been shown to be useful in the field of psychiatry. Imaging scans, such as fMRI and MRI scans, have been used to help classify schizophrenic patients (11, 12). Techniques such as these could help physicians diagnose and track disease progression in disorders that are not yet well characterized and involves a very heterogeneous population and set of symptoms. Another common issue among medical datasets is the problem of missing class (or data) labels (13, 14). In such cases, semi-supervised versions of a Generative adversarial network (GAN) have been proposed to deal with issues related to labeled data scarcity and were shown to be effective in searching for chest abnormalities from X-ray images (15).

*Unsupervised Learning: Discovering patterns amongst the sea of big data.* The goal of pattern discovery, as opposed to classification or prediction, is often the case for many ‘omics’ style studies (16, 17). In pharmacogenomics, several applications of unsupervised ML algorithms have been successfully developed to help uncover new patterns in gene expression/regulation for drug development or repurposing old drugs, patient stratification within clinical trials, and identification of new target regions in

the genome (18). These advances can positively impact both the drug discovery process and personalized treatment. Unsupervised ML has also been used in combination with supervised methods in order to find patterns in medical records and reports that could help a physician expedite the diagnosis process (19).

Neural Networks: Going Deep into Medical Images. DL algorithms are becoming increasingly popular for pattern recognition tasks using images (radiology, pathology and dermatology). For example, a deep convolutional neural network (CNN) was used to identify skin cancer from images of skin lesions and was found to be equal to a panel of board-certified dermatologists (20). In another study, a DL classification network was successfully used for diagnoses and physician referrals based on 3-dimensional retinal scans (21). In both cases, DL and CNNs have been used to reduce false positives from PET/CT scans (22), real-time cardiac arrhythmia detection from EKGs (23), high-precision analysis of EEGs (24); and creating more accurate images from low-dose X-ray CT scans (25). CNN networks have also been combined with Natural Language Processing to classify clinical notes into subdomains within medical specialties and redirection of patient care based on patterns in a patient's EHR (26).

These are just a few examples that demonstrate the role of AI techniques (e.g., DL, CNN, supervised/unsupervised ML) in assisting clinicians. While we have seen that AI has a great potential to improve medical data analysis either by enhancing the accuracy of diagnosis or by making the processes more efficient from early detection of disease to prognosis, it is hard to imagine that it could completely replace the expert involvement in the medical decision-making process. Patterns identified using AI techniques are mere correlations

among the data, not causal relationships, and cannot be directly translatable to replace clinical decision-making from experts in the field. Thus, the narrow AI applications in the field of medicine - the only form of AI that we have achieved so far - still rely on expertise of medical professionals.

**AI-in-the-making for making your day easier.** AI has the capacity to do parallel computations, process and store large amounts of data, and build better algorithms to simplify repetitive tasks. So why not use AI's help in helping ourselves in the workplace? AI can power dictation to produce notes, and update social media and billboard on local ER wait times. AI can also help decrease wait times and create a more efficient schedule for both patients and physicians or help us build systems for addressing patient messages. Chat bots like HealthTap, Your.MD or Safedrugbot already exist to help patients and physicians streamline the process by offering information and providing solutions.

Imagine applications like Alexa for medicine. Voice assistants are another example of the emerging use of DL in the field of medicine. Present in several popular voice assistant programs (e.g. Siri and Alexa), the use of similar algorithms is now expanding into the world of EHR. Companies are working to use the same types of algorithms to fill prescriptions or schedule follow-up visits to reduce the need for human staff (26).

Participating and keeping up with latest research is also a large part of every academic physician's life - another task AI is adapted to help with. With PubMed having over 23 million papers, imagine the difference between manual searching versus using Watson (IBM) when conducting a literature search. Especially, Watson (IBM) has the capability of going through millions of pages in seconds. Its speed of processing

the latest information is the reason it is being used in oncology centers for making evidence-based treatment decisions. Potential research collaboration is another option that it can offer as the database can store the information of other researchers working on a similar idea. Imagine being presented with just the articles that matter to you with your morning cup of coffee. With the help of AI in handling repetitive/administrative tasks, physicians can enjoy a more important task: healing (27).

**Mind the data that is mined.** Finally, yet importantly, where is your data coming from and what is it being used for? Data provenance is key to the inferences drawn by any AI that uses it. One should not ignore the critical details that pertain to the background, source and origin of datasets employed in developing AIs or other intelligent systems (28). A recent Google AI - Automated Retinal Disease Assessment (ARDA) tool developed using high quality retinal scans taken from developed countries has shown the ability to identify early signs of ocular conditions that doctors could potentially have missed. But it was unsuccessful in detecting diabetic retinopathy from lower quality images taken in field clinics in rural India (29). In another example, markers of Down Syndrome developed using low-resolution images were not able to reflect similar inferences when remodeled using modern high-resolution images (30).

The predictive abilities of AI systems are only as good as the input data and are often anchored to the framework, subjective assumptions, and biases of the people conceptualizing the algorithm (28). In image recognition analysis using deep/machine learning methods, there is the risk of flooding the AI system with images heavily partitioned towards a particular

condition and creating a bias towards identifying that condition. Similarly, EHR data are also subject to biases or oversampling, as it tends to contain more sample data from sicker populations (31). Models trained on oversampled or biased datasets will most likely fail to accurately generalize to populations outside the study design. Deliberate efforts to carefully probe models for various vulnerabilities such as unreliable behavior or bias due to shifts in population, practice, patterns, technology or other characteristics of data are critical to mitigate or reduce the risk of incorrect decision-making in the field of medicine (31). This is not entirely surprising given that the size, quality, dimensionality and resolution of sample images, which AI algorithms are trained on, can be different from those used in practice.

**Trust and Dependence on Automation in Healthcare.** Application of AI can only advance to its highest potential when there is acceptance and cooperation within the medical community. Despite the development of very successful prototype systems, many clinicians are not yet ready to use AI-based systems in clinical practice. Embracing a symbiotic relationship is key in enhancing performance and preserving provider-patient interactions (e.g., communication, presence, empathy and comfort) in healthcare (27).

Human interaction with AI-based systems is often intended to result in performance gain. This is not always true because of a disproportionality between the true capability/reliability of AI systems and the clinician's trust of automation. However, over- or under-trusting task automation could result in sub-optimal decision-making (32, 33). While many human factors such as belief, behavior, intention, and attitude can be used to define trust, in the field of intelligent systems trust is an attitude where

humans rely on AI-systems in situations characterized by ambiguity/uncertainty (32). In contrast, automation usage or dependence is a behavior attribute. While they could be highly correlated with each other, they are not the same.

It is also important to recognize that trust is influenced by time. In the real world, trust evolves and stabilizes over time. Similarly, prolonged and repeated encounters with AI-enabled tools and technology can improve a physician's trust, attitude and behavior towards such emerging technologies. In addition, involving clinicians and other medical professionals in the process could also significantly improve their trust towards AI technology. The more you use something, the better you understand it, and the more accepting of it you will likely become; whether it be a new technology (GPS, social media) or machines (smart phones, tablets, grocery self-checkouts). We need to stop being afraid of AIs and start working on making them improve our practice and life.

**Summary & Outlook.** Overall, it is evident that AI can define clinical patterns and insights by utilizing big datasets beyond current human capabilities. However, this requires careful and appropriate integration of AI into healthcare and medical practice. It becomes important to recognize that there are certain tasks and interactions that can neither be digitized nor quantified for an AI system to adapt. However, the current development of any AI-based automated system still requires a human expert. These systems are effective at learning patterns, correlations, and rules in complex datasets, but they do not offer meaning, purpose, or a sense of justice/fairness for the data and its context. Algorithms trained on oversampled data tend to amplify the bias further. By design, AI systems cannot be aware of the source of input data and thus require human

intervention before relaying the clinical implications to a patient. Humans and machines can each excel in distinct ways that the other cannot, meaning that the two combined can accomplish what neither could do alone. Thus, a harmonious collaboration between emerging AI technologies and human intelligence, or augmented intelligence, is likely the most powerful approach to enhance a clinician's real task in medicine. So, rest assured that while AI might drive your next car for you, the art of practicing medicine will still require the sometimes illogical but critical leaps that remain in the domain of human intuition and empathy.

**Search methods.** Articles were obtained by using search strings such as “artificial intelligence in medicine” and “supervised/unsupervised machine learning in medicine” from PubMed and Google Scholar.

## References

1. Turing AM. *Mind* 1950; 49:433-460.
2. McCarthy J *et al.* *AI Magazine* 2006; 27(4):12-14.
3. Schwartz WB *et al.* *N Engl J Med.* 1987; 316(11):685-688.
4. Baxt WG. *Neural Comput* 1990; 2(4):480-489.
5. Wu Y *et al.* *Radiology* 1993; 187(1):81-7.
6. Jiang F *et al.* *Stroke and Vascular Neurology* 2017; 0:e000101.
7. Tarca AL *et al.* *PLoS Comput Bio* 2007; 3(6):e116.
8. Haury AC *et al.* *PLoS ONE* 2007; 6(12):e28210.
9. Wood DE *et al.* *Sci Trans Med* 2018; 10:eaar7939.
10. Kourou K *et al.* *Comput and Struct Biotech* 2015; 13:8-17.
11. Dluhos P *et al.* *NeuroImage* 2017; 155:10-24.
12. Yang H *et al.* *Front Human Neuro* 2010 4:192.
13. Liang Y *et al.* *BMC Med Genom* 2016 9:11.
14. Liao R *et al.* 2019; arXiv:1902.10785v3.
15. Madani A *et al.* 2018. San Jose: IBM Research.
16. Perakakis N *et al.* *Metabolism* 2018; 87:A1-A9.
17. Kiselev VY *et al.* *Nat Rev Genet* 2019; 20:273-282.
18. Kalinin AA *et al.* *Pharmacogenomics* 2018; 19(7):629-650.
19. Zech J *et al.* *Radiology* 2018; 287:2.
20. Esteva A *et al.* *Nature* 2017; 542:115-118.
21. De Fauw J *et al.* *Nat Med* 2018; 24:1342-1350.

22. Teramoto A *et al.* *Med Phys* 2016; 43(6): 2821-2827.
23. Yildirim O *et al.* *Comput Biol Med* 2018; 102:411-420.
24. Acharya UR *et al.* *Comput Biol Med* 2018; 100:270-278.
25. Kang E *et al.* *Med Phys* 2017; 44(10):e360-e375.
26. Molteni M. 2010 [<https://www.wired.com/story/does-your-doctor-need-a-voice-assistant/?verso=true>].
27. Verghese A *et al.* *JAMA* 2018; 319(1):19-20.
28. Miller DD. *NPJ Digital Medicine* 2019; 2:62.
29. Abrams C. August 13th, 2019 [<https://www.wsj.com/articles/googles-effort-to-prevent-blindness-hits-roadblock-11548504004?mod=searchresults&page=1&pos=1>].
30. Shiney J *et al.* *Biomed Pharm J* 2017; 10:1559-1568.
31. Gianfrancesco MA *et al.* *JAMA Intern Med* 2019; 178(11):1544-1547.
32. Hoff K., Bashir M. *Hum Factors* 2015; 57(3):407-434.
33. Parasuraman R and Riley V. *Hum Factors* 1997; 39(2):230-253.

**How to cite this article?**

Narumanchi *et al.* (2021). *TCABSE-J*, Vol. 1, Issue 2:1-8. Epub: Oct15<sup>th</sup>, 2021.

**Disclaimer:** The authors alone are responsible for the content of this manuscript. The findings and conclusions in this manuscript are those of the authors and do not necessarily represent the official position of the National Institute for Occupational Safety and Health, Centers for Disease Control and Prevention.

**Funding:** This work was supported by the National Institute for Occupational Safety and Health [NIOSH 939051E].

**Competing interests:** The authors declare that they have no competing interests.