AI Education Roadmap

An introduction to our series

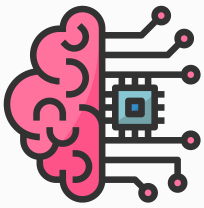

Artificial intelligence (AI) is a fascinating, yet highly complex, field. Throughout the past several decades, it has witnessed significant ups and downs in its adoption. However, we now seem to be in an era in which AI is being extensively deployed within every sector, even healthcare, an industry often known to be the 'laggard' in adopting new technologies. Along with the increased prevalence of AI comes a growing awareness among clinicians of the need to become familiar with this emerging technology.

We are Whitley and Christy, pharmacists with a passion for digital health and innovation, especially as it relates to artificial intelligence (AI). Realizing the challenges of tackling the subject of AI, we have made it our mission to help demystify the topic and make AI accessible to all pharmacists. When it comes to learning about AI, we have found that it can be challenging knowing where to begin with such a vast topic. With all the articles and educational resources available online to choose from, how do you know which ones will meet your needs and which ones to invest your time in? To help answer these questions, we created an **AI Education Roadmap**. It's our vision, through this education series, to guide others through the intricacies of AI and its applications. Much of the information provided has been gleaned through our own first-hand exploration of numerous AI resources. However, please keep in mind that this is not meant to be an exhaustive or comprehensive resource guide; rather, it aims to provide a suggested approach on how to focus your time and energy so that you get the biggest bang for your buck. We hope you will find some value in the lessons we learned from our own educational journeys.

AI Education Roadmap

An introduction to our series

Whitley Yi graduated from the University of Colorado and completed her ambulatory care residency at UNC Hospitals with an emphasis on informatics and analytics, where she developed a passion for artificial intelligence and the unique complexities of applying AI to patient care. She has spoken at several events on the clinical considerations and challenges of AI implementation in healthcare and continues to advocate for greater collaboration between clinicians and data scientists. Whitley currently works at a healthcare start-up in Chapel Hill, North Carolina. Her mission is to be apart of disrupting the healthcare system by leveraging the use of technology and AI to empower patients.

Christy Cheung graduated from the University of British Columbia and is currently completing an Industrial Fellowship at Sanofi Genzyme Canada. As a pharmacy student, Christy attended healthcare hackathons, which sparked her passion in digital health. Like Whitley, She believes that the healthcare industry is ripe for disruption and what better way than to look to and collaborate with other disciplines. Christy sees the promise in technology, not as a replacement to current practice, but rather as a complementary tool, which can be used to enhance patient care. Her goal is to integrate her enthusiasm for digital with her clinical background and to contribute to the digital transformation of the healthcare industry.

Part 1 - Artificial Intelligence Introduction

Where did AI come from?

Society first toyed with the concept of artificial intelligence in the early 20th century when science fiction introduced human-like robots into the creative industry. It was around the 1950s when academics began to seriously explore the notion of AI[1]. The famous mathematician and computer scientist, Alan Turing, widely known for his contribution to what we now know to be a computer, was one such academic[2]. His paper, “Computing Machinery and Intelligence,” highlighted the logic behind building intelligent machines and how to test their intelligence[2]. The term “artificial intelligence” was officially coined by John McCarthy, a computer scientist and one of the founding fathers of AI, alongside Alan Turing[2]. He organized the famous Dartmouth Summer Research Project on Artificial Intelligence in 1956, which is considered the conference that started artificial intelligence as a field[2].

Hype vs. Reality? - Keeping it in Perspective

How do you differentiate between hype and reality? When looking at headlines in the news, it is not apparent which claims make illogical leaps in their assumptions or which falsely represent the current state of the technology. To put it in historical context, this is not the first time there has been a lot of hype surrounding AI; the field has actually gone through many hype cycles, each followed by an "AI winter." These were periods of time when AI became unpopular with the general public due to its failure to deliver on its promises and to live up to expectations. AI research still occurred, but it sometimes went by other names. We know, therefore, that there is a precedent for the public to overestimate AI’s ability. Keeping this in mind, we should approach the headlines with a heavy dose of curiosity and skepticism. After all, there is always the possibility that the current hype cycle could lead to another AI winter. Only time will tell.

Getting Familiar with the Concepts and Jargon

To gain a broad overview of what AI is, the best place to start is basic concepts and terminology. This is often the most challenging part, and it is okay if it takes a little while for the concepts to click; part of the reason for this is that definitions often change and are not used consistently by everyone to refer to the same things.

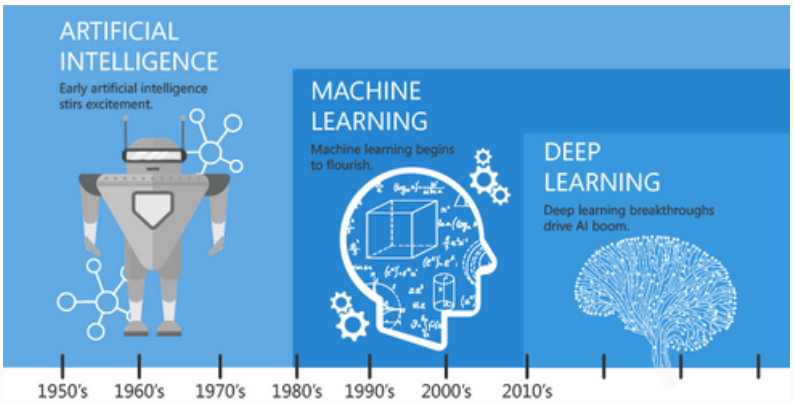

https://hackernoon.com/difference-between-artificial-intelligence-machine-learning-and-deep-learning-1pcv3zeg

It is important to keep in mind that "artificial intelligence" and "machine learning (ML)" are separate terms with overlapping, but distinct, definitions. The term "artificial intelligence" is somewhat nebulous. Although you will find that there are generally agreed-upon definitions of what the term means as a concept, when it comes to determining whether something should be classified as artificial intelligence or not is another matter entirely.

Broadly speaking, AI refers to the idea of mimicking human intelligence in regards to the way we act and think; in other words, it refers to completing tasks or functions previously requiring human intelligence to complete. To put this into perspective, when the calculator was first invented, it fell under this definition of AI, as it was able to complete a task previously only completed by humans. However, the definition of what it means to mimic human intelligence has evolved significantly since the 1950s, when the term "artificial intelligence" was first coined. As the definition of human intelligence evolves, so will the definition of artificial intelligence.

Machine learning, on the other hand, is easier to define and should be viewed as a technique of AI. It refers to the use of algorithms to analyze and identify patterns within datasets without having to be explicitly programmed. In other words, it gives machines the ability to learn directly from data. However, something important to note is that machine learning is on a spectrum; it is less about whether machine learning is present or not, and more about the degree to which it is present in an algorithm. On one end of the spectrum are algorithms such as logistic regression, while on the opposite end of the spectrum you would find deep neural networks. The less a model or algorithm needs to be explicitly programmed and the more it learns from examples, the more like machine learning it is. Deep learning, then, is a subset of ML, which we will describe in later parts of this series.

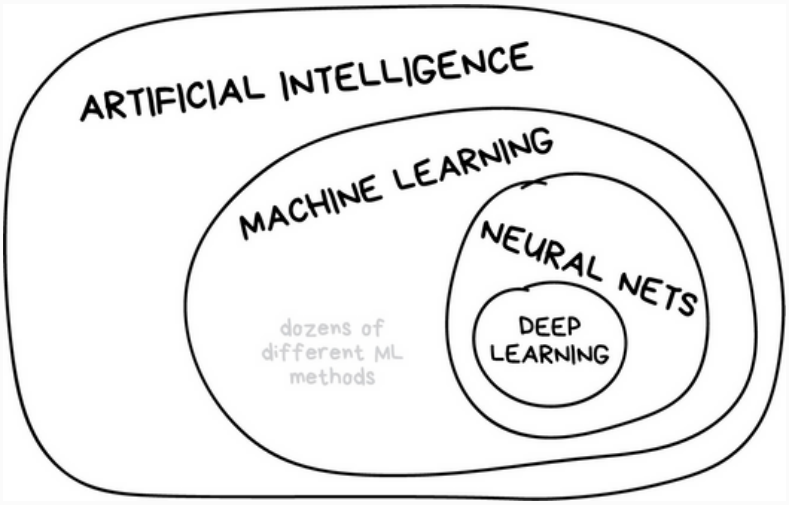

https://vas3k.com/blog/machine_learning/

How does this relate to what you already know?

Fortunately, there are some machine learning tools you may already know. Predictive models and other statistical methods you may be familiar with also utilize machine learning. A multivariate regression analysis, for example, can be used to look at a large number of variables, and determine which of them are significantly associated with the outcome of interest and the degree of that association. The outcome of the algorithm is an equation representing the line that best fits the data. This is an example of machine learning because the final equation (i.e. output) is not determined by the individual writing the algorithm; rather, it is determined by the actual data the algorithm sees. In fact, the equation may change if it is run with additional data or run a second or third time on a different dataset. You can imagine how such a capability can be very helpful to us. If an algorithm can learn on its own, then, the more data it has, the smarter it becomes. Additionally, there is less manual effort involved in development.

https://medium.com/swlh/how-can-ai-impact-road-network-operations-6cfb1de26c7f

When it comes to traditional statistics, well-established best practices are available to guide the selection of an appropriate statistical method based on the outcome of interest, the type of data (categorical vs. continuous), data distribution (parametric vs. non-parametric), and the number of variables. For example, a chi-square test would be used to compare categorical data between two unpaired groups. This concept is important because there are, in fact, many different types of AI/ML algorithms, some of which are being explored as an alternative to traditional statistical methods for certain use cases in healthcare, such as predictive modeling. However, since the field of AI in healthcare is still relatively young, we do not yet have well-established best practices for which types of AI models are most appropriate for a given problem. To illustrate how many types of algorithms exist to choose among, the diagram below shows a few examples of algorithms in each major category. As the field of AI in healthcare continues to mature, we should begin seeing more effort in developing best practice guidelines.

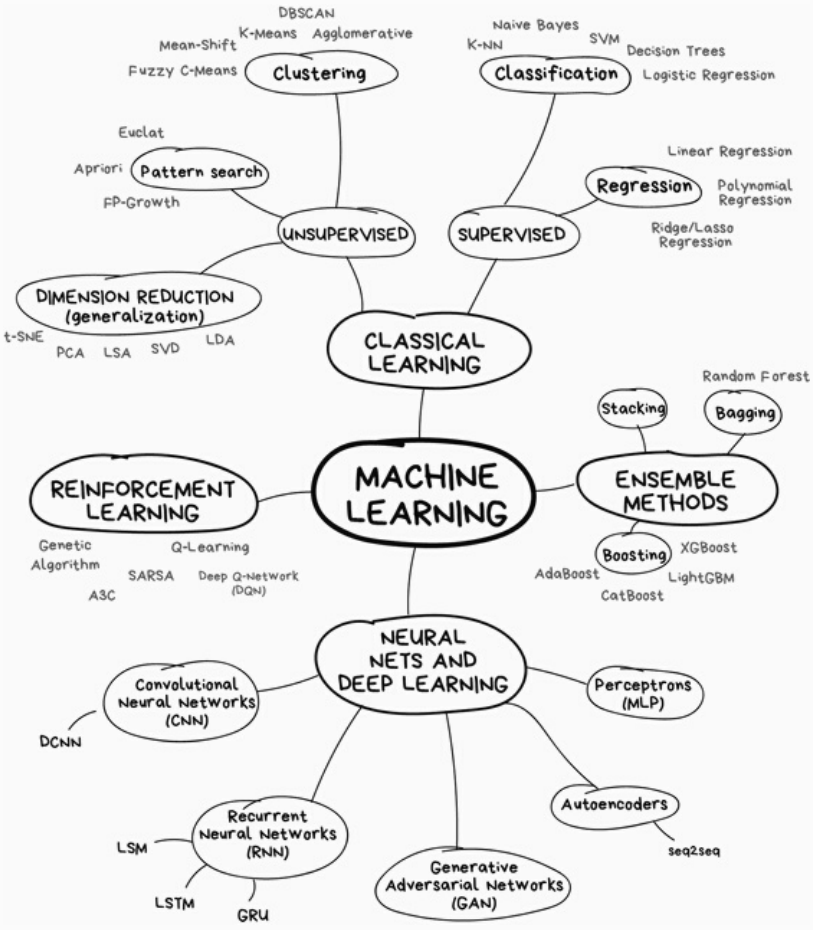

https://vas3k.com/blog/machine_learning/

Limitations

Though AI has been touted as revolutionary, it is not without its limitations. The implementation of AI is dependent on the data given. Much like in humans, if we are taught the wrong information, we will behave and act in accordance to that misinformation. Similarly, if machines are trained with “bad data,” the predictions they make will be largely inaccurate. This is a phenomenon often referred to as “garbage in, garbage out."

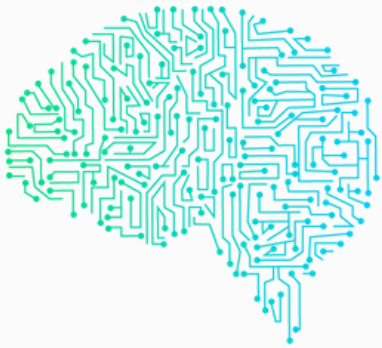

Bias – The quality of data can be impacted by the presence of bias. Let’s give an example. If your intention is to train an algorithm to identify pneumonia from radiological data (ie. from X-rays) and it is trained on data that only includes adult patients, it may not be able to determine the presence of pneumonia for adolescent patients to the same degree of accuracy as in adult patients. This is what we call ‘sampling bias,’ whereby the training data are not representative of the data one would need in that particular situation.[3] Other types of biases are gender and racial biases, which are prevalent in our history, and even in society today; they are a significant concern in the AI community, as they have cultural implications[4]. As you can now see, there are numerous types of biases, and they all become important considerations when building an algorithm.

Black Box – Certain types of deep learning algorithms, such as neural networks, are considered “black box” algorithms, meaning we do not know what the machine is learning. You can think of it this way – as humans, we tend to take the easiest path to learning something new, which includes taking mental "shortcuts". When learning a list of new drug names with similar endings, we may initially memorize only the name endings since they specify the drug class and can, thus, be learned more quickly. Similarly, a machine will pick up on whatever details it finds most helpful, regardless of how relevant they are.

Example Case

A company trains a machine to classify a chest x-ray as normal or abnormal by using pre-annotated images from online datasets. The model was very accurate until it was tested on real patient data, when it began making incorrect predictions. The company discovers that their model did not actually learn how to read a chest x-ray; instead, the model made its predictions entirely on the amount of annotated text on the image. In the training data, the more abnormal the x-ray, the more text it had overlaid on the image. Unfortunately, in real life, x-ray images are not pre-annotated, and since the machine had only learned the association between volume of text and chest x-ray classification, it could not be relied upon to accurately classify non-annotated images from real patients. In the end, the company had to start over and retrain the model with different data. Keep in mind that we want AI algorithms to not only get the right answer, but to get it for the right reasons.

AI/ML models classified as "deep learning" use complex mathematical mechanisms to learn in a non-linear fashion. By definition, deep learning models are not interpretable. Even if a deep learning model is highly accurate, there is currently no way to prove why it is accurate. Just like the example above with the x-ray images, a model might be learning from the noise in the data instead of what we want it to use to learn. This inability to know how an AI/ML model makes its decisions is called a "black box." To illustrate how important the implications of this concept are, Europe passed a law prohibiting black box algorithms from being used to make automated decisions about or on behalf of an individual (i.e. legal profiling) without that individual's explicit consent[5].

References used in this article

1. http://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/
2. https://medium.com/@exastax/the-difference-between-ai-and-machine-learning-32cfef316372
3. https://towardsdatascience.com/5-types-of-bias-how-to-eliminate-them-in-your-machine-learning-project-75959af9d3a0
4. https://medium.com/thoughts-and-reflections/racial-bias-and-gender-bias-examples-in-ai-systems-7211e4c166a1
5. Article 22. General Data Protection Regulation (GDPR). Automated individual decision-making, including profiling. 2018

Resources List

Self-study Resource Guide

When just starting out, online courses are a great resource. Most free online courses begin with a basic overview of AI/ML and its applications. A helpful approach is to watch/read the introductory lesson of a few different online AI courses, and see where that takes you. The rest of the course can be completed if you then decide to dive deeper into the subject, but the introduction lecture will provide most of the conceptual information you are looking for.

Review articles/editorials

- Top AI algorithms for Healthcare 🔗
- Advice to Medical Experts Interested in AI 🔗
- Artificial Intelligence (AI) vs Machine Learning (ML) vs Big Data 🔗

Online tutorials

- Machine Learning for Everyone 🔗
- Machine Learning from First Principles 🔗
- A 6-minute intro to AI 🔗

Online courses

Fast.ai

- **Introduction to Machine Learning for Coders** — This is a great course to watch even for non-coders. It helps demystify AI by providing a behind-the-scenes look at what it means to build an ML algorithm 🔗

Coursera

- **Machine Learning** by Andrew Ng (Stanford University) — Consistently ranked as one of the top online machine learning courses offered for free 🔗
- **AI for Everyone** by deeplearning.ai 🔗
- **Introduction to Artificial Intelligence (AI)** by IBM Watson 🔗

Podcasts

- **Building Robodoc** (Episodes 1 through 12) — Designed for an AI beginner, this podcast walks through the challenges of AI implementation in healthcare and how it is different compared to other industries. It explains fundamental concepts, such as covariate shift, utilizing examples that healthcare professionals can relate to
- **The TWIML AI Podcast** (This Week in Machine Learning and AI) — Discusses cutting edge developments in machine learning and AI. It includes interviews with AI experts in fields ranging from healthcare to marketing 🔗