Part 2 - Understanding the Impact of AI in Practice

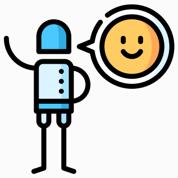

AI/ML is already being used in Pharmacy and Medicine, but its utility varies greatly depending on the purpose. In Part I, we reviewed the basics of AI methodology, terminology, limitations, and barriers. In Part 2, we will lay the foundation needed for understanding the existing and potential impact of AI in practice. The impact of any technology is dependent not just on its capabilities, but also the degree to which it is used. What factors influence the perceived utility and adoption of an AI application? To answer these questions, we will take a closer look at AI models and dive into model design and deployment. Lastly, we will review existing examples of AI in healthcare in order to provide an overview of ways AI is already impacting practice.

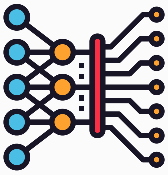

Artificial Intelligence Models

What is an AI model? From a clinical perspective, it is helpful to first think about models as being composed of 3 major components:

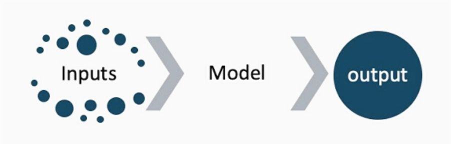

1. The type of inputs (or data) that go into the model. This includes training data, as well as live data used when the model is deployed
2. The model's algorithm(s)
3. The type of output (or inference the model makes) based on the data (i.e. a line of best fit or classifying something as 'yes' or 'no')

In other words, an AI model is the vehicle used to generate actionable information from the data collected. There are numerous models that exist and which one you select is largely dependent on the type of data you are working with and what you hope the model will glean from that data.

The potential utility or value of a model is often discussed in terms of its accuracy, which is a function of how precisely and reliably the model provides the correct output. In other words, how well does the model map the inputs to the right output? However, there is one more element to consider.

Just as the validity or applicability of a clinical trial's results are always evaluated within the context of the study population, trial design, and specific outcomes that were measured, an AI model cannot be looked at in isolation. We must think of a model in the context of the environment in which it is designed to operate. In other words, we must consider all the intended and potentially unintended consequences that might occur when a model is operationalized in real life. The ultimate value, therefore, of an AI application is a product of both its design and its deployment.

AI application = Design + Deployment

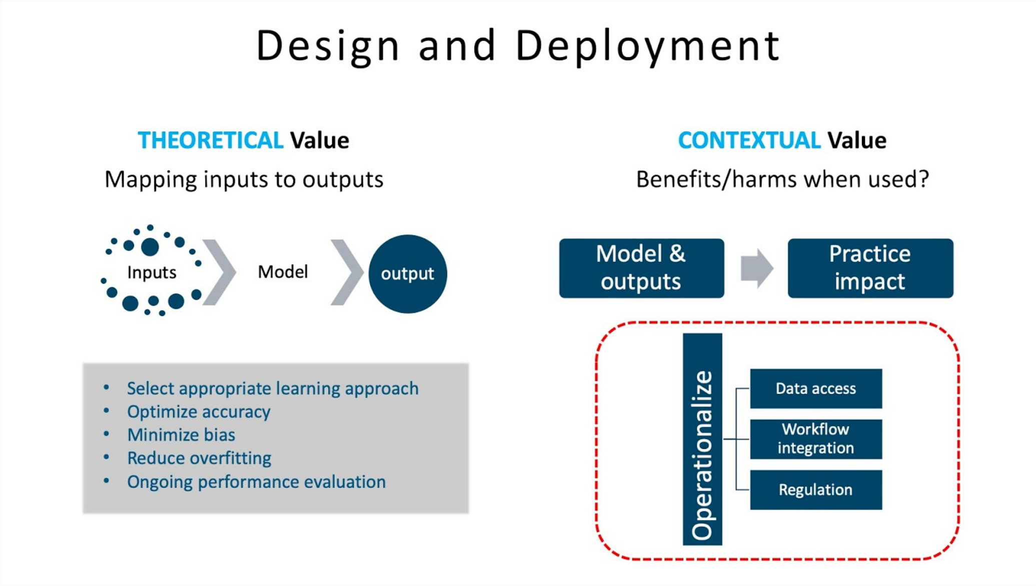

In this image, the value of an AI application is broken down into its theoretical value, which is defined here as the model's ability to accurately generate the output it was designed to provide, and its contextual value, which is defined as the clinical impact of the model when integrated into a complex clinical workflow and environment.

Deeper Dive – Defining models by their parts

Input Data: It all starts with data, but what kind?

Many types of data can be used in a model, including unstructured data such as images or natural language, as well as structured, or tabular data. The ability to look at unstructured data is a major distinction between machine learning and basic statistical analytics. However, machine learning usually still involves considerable data pre-processing, and often the data are transformed from one form to another. For example, a model used for identifying fraudulent online transactions does not use the data you might initially think it would[1]. Rather than looking at transaction logs or keyboard strokes, the model actually looks at the movements of the mouse's cursor on the screen. It converts those movements into an image (see below) and uses image-recognition to detect patterns in the image. The unique visual patterns provide a way to represent complex behaviors by the user that would be hard to represent through any other data type.

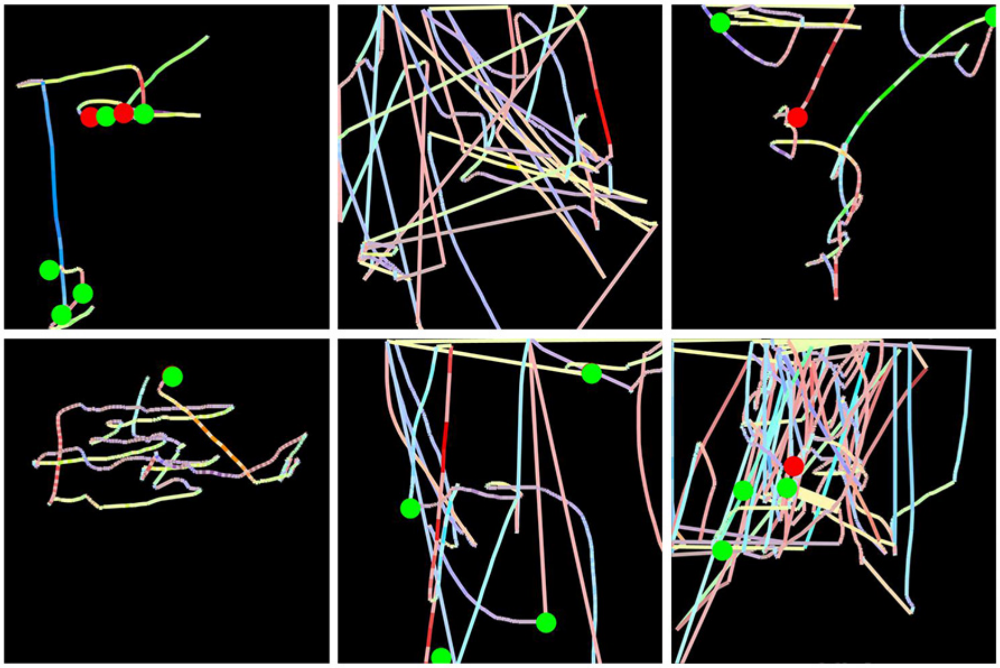

https://paperpile.com/c/9akFfA/whqS

Each image above represents the summation of mouse movements on the computer screen, utilizing color to represent acceleration and direction. These unique patterns are able to help detect behavior patterns associated with fraudulent activity[1].

Types of Data

Images

Image-recognition is where deep learning really began to make a name for itself, taking on pattern recognition tasks that previously had not been able to be automated. Neural network models in deep learning, for example, are able to pick up patterns in images that may not be perceptible to the human eye. An important point to note when it comes to image recognition is that the model architecture is often dependent on the size (i.e. number of pixels) of the image. Each pixel of the image is an individual data point that makes up the input going into the model; therefore, a different image size means a different data input size, which could impact results. When training a model for healthcare applications, you want to train it with images of the same resolution and size of those that the model would see in real practice.

Natural language

Natural language refers to the unstructured, nuanced nature of human language, while natural language processing (NLP) is used to describe the way machines make sense of human language. NLP is the ability for humans and machines to interact through the use of everyday language as opposed to code. It is what enables applications such as Siri, Google Home, Amazon Alexa, and Google Translate. Human language is complex. It is not enough to know the meaning of individual words; the machine must translate the intent and meaning of words based on the context in which those words are being used, which can vary widely.

How does a machine detect tone? Jargon? Sarcasm? Informal slang? NLP is a challenging field within machine learning that still has much to learn. Similar to how pixels are extracted from an image and transformed into data points that a model can process, NLP models also have to distill language data into a machine readable form. It involves a number of techniques for evaluating syntax (i.e. grammar) and semantics (i.e. meaning of words). Ultimately, regardless of the specific technique, the end result is that each word, or group of words, will be represented by a mathematical vector. It is this vector that is passed through the NLP model. When it comes to clinical application, it is important to understand that the data the model is trained on is critically important to the model being able to interpret information accurately. In the medical and scientific world, if an NLP training dataset (called a "corpus") does not contain sufficient representation of a specific topic or subspecialty, it is likely to be less accurate understanding text from that topic. Unfortunately, this is often the case with medication information and pharmacy-specific jargon, which is lacking in the large corpuses of many popular NLP models.

Tabular data

Tabular data refers to data that can be represented in a table, such as a spreadsheet. One that many are likely the most familiar with is an Excel spreadsheet. Tabular data are structured and organized within rows and columns. Two types of data should be considered in this case: categorical data, which refers to discrete variables (e.g. colors, cities), and continuous data, which has an infinite number of values (i.e. numbers). Tabular data are abundant in all industries and machine learning can be leveraged for a number of reasons, from optimizing business operations to the forecasting of sales.

Models — Algorithm learning styles

When we talk about algorithms, there are two main things to consider: the learning style and the algorithm "type." We will talk about learning styles first, which is a way to describe how the algorithm uses data to gain information. There are four major learning styles. Ensemble learning is used to describe when multiple learning styles or algorithms are used in combination.

Supervised

In supervised learning, the model is being trained with data that has been labelled/categorized correctly. For example, we tell the model that spruce, cypress, and evergreen fall under the category of "tree", while tulip, rose, and lily fall under the category of "flower." In this case, what the algorithm is doing is learning to make connections, autonomously, by mapping the input data to its correct label or category. When the model churns out an incorrect answer, the model will make incremental adjustments so that it is less likely to be wrong the next time. Once it is trained to achieve an acceptable level of accuracy, it can then be applied to new data. The model will extrapolate what it learned from experience (i.e. previous data) to make determinations about new data it has not encountered before. This may include classifying the new data or making a prediction.

A brief note on model over-fitting: when a model over-fits, it means that what the model learned is so specific to the data it was trained on, that it does not generalize very well to new data.

Over fitting may be seen as a model having extremely high accuracy when validated with the data on which it was trained (e.g. EHR records from "Hospital A"), but low accuracy when used with another dataset (e.g. EHR records from "Hospital B"). Ultimately, it is about balancing accuracy and generalizability.

Examples: facial recognition; training a model to recognize handwritten letters; training a model with labelled x-ray images to automatically detect abnormal chest x-rays

Unsupervised

As you may have guessed, this is a type of learning in which we do not have the available answers prior to training the model. This sort of learning might be used in settings in which we simply do not know the answer or it is not immediately apparent or available. The model independently identifies patterns and heuristics to determine an output that is most suitable based on the data it was given. Unsupervised learning is often used for clustering or pattern identification. It can find patterns that may be difficult to identify by hand, especially when working with large datasets. For example, it can be used to define new groupings of data or to segment patients into novel “phenotypes” based on hundreds or thousands of variables per patient. It can be used to detect anomalies, such as identifying an inpatient medication order potentially made in error because of how dissimilar it is to other orders made under similar clinical circumstances. In these examples, it is considered unsupervised learning because there is no ground truth. The model makes associations and creates implicit “rules” about the data without knowing what the answer is, as there is no predefined answer. Many of these types of activities fall under the umbrella of data mining. Unsupervised learning can also be used for dimensionality reduction, which is the process of reducing the number of features in a large dataset.

Examples: market or customer segmentation; word embeddings and language model creation; learning rule associations; anomaly detection; gene clustering; data mining activities

Semi-supervised

Now we get into the grey zone of somewhere in between. This type of learning employs aspects of both learning styles described above. We may have some knowledge of the correct answers, but perhaps the data set is incomplete or unreliable. Semi-supervised learning will use both labelled and unlabelled data. The labelled data can be used to create and apply labels to the unlabelled data, thereby increasing the total amount of labelled data available to be used in training the model. This strategy is helpful when you need a large quantity of training data, but may only have a small amount of labelled data.

Hand-labelling data, also called annotating, is a time-consuming task. It is often a rate-limiting step in building and researching new machine learning models. The option of using a semi-supervised approach helps to scale the development of AI; however, it also requires collaboration between data scientists and subject matter experts to ensure that the rules and algorithms used to create the new labels are appropriate. Otherwise, the model may inadvertently learn the wrong information and its accuracy may suffer.

Examples: web content classification; speech analysis, augmentation of NLP models; protein sequence classification

Reinforcement

Reinforcement learning is a type of supervised learning used for building decision models to be used within a time series or a sequence of events. Instead of classifying what something is, the model determines what action (i.e. decision) should be taken to achieve a predefined outcome or to maximize some type of reward.

It can be thought of a little like supervised learning, except instead of the model knowing right away whether it got the right answer, it has to wait until it has made several decisions to find out if the sum of those decisions led to the outcome it was trying to achieve. A great example of this is playing a board game like chess. The ultimate goal is to win. In the process of playing the game and trying to win, multiple decisions are made without any immediate feedback of whether each move was the right decision or if it made winning more likely. The player does not know if the summation of all their decisions were successful until the end of the game. This is known as the concept of "delayed reward." It is analogous to many healthcare processes in which the desired health outcome is not immediately known. For example, reducing cardiovascular events or preventing a heart attack is a long-term outcome that is impacted by numerous decisions and interventions over time. Similarly, many actions take place during a patient's hospitalization, but the ultimate outcome of survival to discharge is not known until the patient is discharged.

Lastly, reinforcement learning draws from aspects of childhood (and human) learning with the concepts of reward and punishment. The model trains within its environment and is continually modified so as to maximize reward and minimize penalty[2]. You can argue, however, that this sort of learning largely depends on the environment the model trains in. If put in another environment, this trained model may not behave in the desired way.

Models — Algorithm type

There are numerous types of machine learning algorithms. While we will not go into depth on the topic at this time, we do want to mention a few of the most common ones you may see.

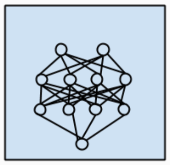

Deep Learning Algorithms

Neural networks, modeled after the human brain, use multiple layers (or processing steps) to take a variable from an input to an output in a non-linear fashion. There are many different types of neural networks. Some that you may hear of are recurrent neural networks (RNN), convolutional neural networks (CNN), long short-term memory networks (LSTM), or generative adversarial network (GAN).

As mentioned in Part 1 of this series, deep learning is a subset of machine learning that includes neural networks. It involves algorithms that use hierarchical feature learning, whereby learning is based on first finding lower-level features in an input variable and building upon those to detect more complex, higher-level features of the variable.

Another way to think about it is that the model's representation of the input is sequentially becoming more and more complex as it travels through the layers. Each layer of the model is able to add on to the information in the previous layer. Suppose we have a neural network with 3 layers that is trained to recognize images. The first layer might only detect the presence of simple shapes, such as lines, circles, or squares. The next layer might look for patterns made by the combination of two or three shapes. Lastly, the third layer would then look at all the combinations of all the patterns to understand the image as a whole. Because of the level of complexity that it can process, deep learning has been attributed to greater success for tasks such as image recognition and natural language processing.

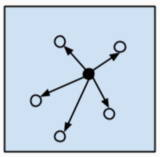

Instance-based Algorithms

A **Support Vector Machine (SVM)** is a type of classifier that tries to find a hyper-plane in a multi-dimensional space that accurately separates two groups from each other. Once this hyper-plane is defined, it can be used to classify a new data point as belonging to one group or another based on which side of the hyperplane it sits.[3]

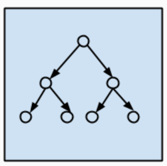

Decision Tree Algorithms

Decision Trees can be used for classification or regression analysis. A decision tree takes labelled data and repeatedly splits the data into smaller and smaller groups, based on the attributes of the data, until it reaches a "leaf" or a terminal node. The goal is for each split to progressively divide the data into groups in which the members within each group are more similar to each other and less similar to other groups.[4]

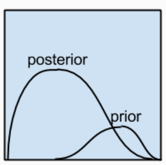

Bayesian Algorithms

https://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/

Naive Bayes is a simple, but powerful, classifier based on Bayes' Theorem that distinguishes between variables based on their unique features. Drawing on Bayes' Theorem, it calculates the probability that something belongs to a specific class based on what is known about the probabilities of its features. It can be used as a binary or multi-class classifier.

Models – Output

When we think about the output of the model, it all depends on how the model was designed to provide information. It can do this in many ways. If a model is a classifier, it might be designed to classify multiple elements or only one element. For example, if a binary model were designed to determine if an image contained a dog, the model will only provide one of two outputs – "yes" or "no." If a model is a multi-class classifier, designed to recognize multiple types of animals, its output will be the name of one of the animals from the list of animals it was trained on. Some models also provide a probability score for their outputs. Other models, such as the reinforcement learning models discussed earlier, might provide a recommended action step. Below are three common categories of model outputs:

- Classification: Often used in diagnostic applications
  Examples: Classifying a chest X-ray as abnormal, or a CT scan as indicative of the presence of a tumor, or a retinal fundus photograph as negative for diabetic retinopathy

- Risk Prediction: Often used for risk stratification, to identify individuals who might benefit from intervention

  Examples: Predicting the risk of experiencing an adverse drug event, of being readmitted to the hospital within 30 days, or the risk of mortality from a cardiovascular event

- Recommendation: Based on a desired outcome, this type of model provides a recommended intervention, action, or series of actions.

  Examples: For a given patient, providing the recommended dose of IV heparin most likely to achieve target serum levels, or the initial drug combination that will be most effective at reducing symptoms of depression

Models – Deployment

Model deployment is all about implementation in practice. How would the model be operationalized in a complex clinical workflow? There are many factors that must be considered, including regulation, ease-of-use, trustworthiness, and understandability by the end-user. The first question that must be answered, however, is what role the AI model will play. There are two main approaches when implementing an AI model:

- AI Model Replaces a Human Process: An autonomous system is a stand-alone system that does not require human input. It often replaces a repetitive, well-defined, and/or standardized task.

  Example: Closed-loop insulin pump that automatically adjusts insulin based on continuous glucose monitoring.

- AI Model Augments a Human Process: Often referred to as "human-in-the-loop" models, these models generate information to be used by a clinician in their decision-making process. The clinician acts as an important checkpoint to validate the model output and its applicability to the individual situation.

  Example: Clinical decision support tools

It is worth noting that for the purposes of Part II of the series, we are focusing on AI applications created to be used within healthcare, whether from an operational, administrative, clinical practice or clinical research perspective. Model deployment is a critical aspect of these types of AI applications. However, another area of AI/ML that applies to health is knowledge discovery, which is an important field of research. This is an area of analytics research in which machine learning is applied to big data to generate new insights that can then inform the subject of future research and hypotheses. For example, by using machine learning to analyze >60,000 prescriptions of more than 200 drugs, several potential drug-drug interactions were discovered that had previously been unknown.[5] When used solely for knowledge discovery, deployment is not a major consideration for these AI models.

Key considerations

While we have touched on a lot of information about AI models, ultimately, how can you start applying this knowledge to help you sift through all the news and data about emerging AI tools and applications in healthcare?

The following 3 key questions can be used as a quick method to gain a high level understanding of the potential clinical utility of a model.

- Do the training data match the type of data the model uses to make inferences when deployed?
- Do the data the model uses when deployed, align with the data that makes the most sense based on the clinical practice workflow in which the model is being used?
- Is the model intuitive and does it provide value to the end-user based on its role in the workflow?

Current state of AI in health

With the exponentially increasing amount of literature published every year, it is challenging to stay up-to-date on both state-of-the-art AI techniques and emerging AI use cases (i.e. applications). Keep in mind, possibility does not equal reality. The story of AI in healthcare is riddled with proprietary algorithms, slow adoption, and challenges to implementation. In 2018, there were ~6,000 publications on AI in healthcare, while a recent article outlined only 32 major examples of AI currently in use in healthcare.

The popularity of different AI techniques has evolved and cycled over the years. To provide an in-depth look at every AI-powered technology currently being implemented or explored in healthcare would be beyond the capabilities and the scope of this series. Our intent is to provide you with a high level overview of the current state of AI in healthcare and to do so, we have chosen several key examples to illustrate existing AI capabilities and trends in healthcare.

Applications and use cases – Biopharmaceuticals Industry

Drug discovery

Companies leading this area are looking at drug discovery from multiple angles. For example, one start-up is amassing data from research papers, patents, clinical trials, and patient records to delineate potential relationships between genes, proteins, diseases, candidate drugs, and other biological elements.[6] The goal is to use multiple machine learning techniques, including NLP, in order to find new patterns by combining different types of information spread across numerous sources and domains.[6] These insights can then be used to inform methods to better target disease processes.

Drug re-purposing

The traditional process for drug discovery involves high throughput screening (HTS) of large compound libraries. By today's standards, HTS is an inefficient, costly, and arduous process. Artificial intelligence has emerged as an approach that wins on accuracy and efficiency; because AI models can find patterns in data, they are able to sift through the compounds in a manner that makes scientific and clinical sense. Putting this technology to work in a compound library of already approved drugs can help researchers identify molecules that are fit for repurposing, while also doing so at a much faster rate than humans.

A recent study evaluated a deep learning approach for in silico identification of molecular targets from existing drug compounds. deepDTnet, as it is called, draws from chemical, genomic, phenotypic, and cellular information to determine biological and pharmacological relationships.[7]

From their algorithm, they are able to profile the vast drug-target interactions of each drug molecule which often contribute to 'off-target side effects', but in this case, may have potential for treating other disease states.

Yet another example would be Pharnext, a startup born ahead of its time, headed by Daniel Cohen, who helped produce the first-ever human genetic linkage map. They recently demonstrated positive results in a Phase III trial with a drug combination made up of already available drugs, baclofen and naltrexone.[8] This compound, PXT3003, aims to treat Charcot-Marie-Tooth disease, a rare neurodegenerative condition, and has been granted fast-track status by the FDA.[8]

Clinical trials

Clinical trials can be a painstakingly slow undertaking, which is exacerbated by multiple bottlenecks that exist early on in the process. One of the largest bottlenecks is the patient recruitment process, which is both time-consuming and costly. Oftentimes, fewer than the goal number of patients are found, which delays the start of the clinical trial tremendously. Fortunately, technology is able to assist. With the amount of data that resides in the healthcare industry, digital tools such as artificial intelligence and machine learning, can help optimize the process by picking out the most relevant information to act upon. If we look at the recruitment process, this would mean having an AI model predetermine the most suitable candidates within a particular disease state who are most likely to respond to a candidate drug.[9] From the patient's perspective, this could also simplify their search for an appropriate clinical trial. The below image depicts additional areas of the clinical trial process in which AI can be leveraged.

Exhibit 27: Predictive Analytics and AI Driving Value for Clinical Development

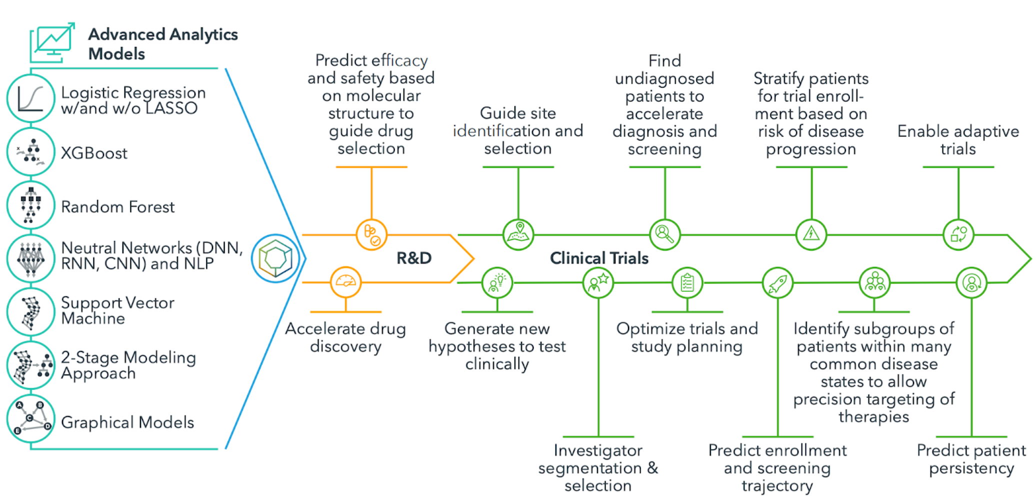

Source: IQVIA Advanced Analytics, Feb 2019; IQVIA Institute, Mar 2019

https://www.iqvia.com/-/media/iqvia/pdfs/institute-reports/the-changing-landscape-of-research-and-development.pdf?_=1586016604253

From the IQVIA Report "The Changing Landscape of Research and Development," April 2019 (Advanced Analytics models simply refer to Machine Learning models)

Applications and use cases – Direct Patient Care

Predictive models

Predictive modeling involves scanning historical data to predict future outcomes, and it can be applied across a range of healthcare operations, from enabling earlier diagnosis to optimizing hospital triage for patients. It allows clinicians to shift from traditional reactive care to a proactive approach. In an industry in which data are abundant, AI is transforming the way we build and use predictive models, moving us towards greater individualized care. It is anticipated that this technology will improve the way diagnostics, prescriptions, and treatments are handled. Below are a few examples:

- Palliative Care: Penn Medicine developed a system, called Palliative Connect, which is a mortality prediction model. Leveraging clinical data from electronic health records, it aims to identify hospitalized patients who are at risk for poor health outcomes and would, therefore, benefit from early palliative consultation.[10] Additionally, NLP can be used to examine the content and sentiment of ICU clinical notes to predict mortality risk in patients; use of NLP-based methods performed better for mortality prediction than using structured clinical data alone.[11]

- Chronic Kidney Disease: Cricket Health is a startup focused on kidney care. Armed with a multi-year dataset that includes information such as demographics and patient co-morbidities, the team developed a machine learning model to identify patients at-risk for or living with chronic kidney disease and to predict an estimated glomerular filtration rate for those patients.[12]

- Oncology: A significant challenge in the field of pathology is the heterogeneity that exists among pathologists' interpretation of biopsies. One can imagine how this may lead to negative downstream effects relating to the patient's treatment. Current use cases of predictive analytics in this field use tumor characteristics to predict prognosis and response to therapy.[13]

- Risk of Opioid unintentional overdose death: NarxCheck Score uses machine learning to predict risk, at the individual-level, of death from an unintentional opioid overdose. The score was designed to help guide clinician decision-making in real time.[14]

Diagnostics

IDx-DR is the first AI model approved by the FDA to diagnose autonomously (i.e. without human intervention).[15] IDx-DR uses its software to analyze retinal images of the eye for the presence of more-than-mild diabetic retinopathy (mmDR).[15] Currently, the AI-powered software is only approved for use with the Topcon NW400 fundus camera.[15] Similar to medications requiring approval for additional indications, IDx-DR must submit to the FDA for approval of any additional camera models they would like to use with their software. As mentioned previously, because cameras can produce images of different size, resolution, and noise, the model has to be trained on and/or validated with images specific to each camera for which it will be used.

Radiology

There are many examples of companies, such as Arterys, Zebra Medical Vision, or qure.ai who provide AI-enabled diagnostic suggestions, abnormality detection, or predictions based on radiographic imaging. As a task well suited for deep learning, image recognition in healthcare is a growing field. A key difference in these applications compared to IDx-DR is that they do not operate autonomously.

Voice analytics

Machine learning can be used to detect the likelihood someone has a disease or condition based on the sound of their voice.[16] It looks at the audio data and can discern patterns that would be undetectable to the human ear. Several startups, including Vocalis Health (previously, Beyond Verbal) and Sonde Health, are exploring voice analysis to inform disease diagnosis; specifically, these companies are exploring association of voice with depression, respiratory health, and emotional health.[17]

Drug recommendations

AI may be able to help optimize drug therapy selection. Treatment with antidepressants frequently yields low response rates and there is often a need to trial several medications before seeing efficacy. As a result, we are constantly trying to identify better ways to understand who might respond to which medication, and AI has opened up new avenues to explore. Researchers have found that AI can help predict antidepressant response based on an individual's electroencephalography profile, which is a measurement of the brain's electrical activity.[18]

Pharmacogenomics (PGx) is an emerging field that explores the affect of an individual's genes on their response to drugs. As it stands today, we have accumulated vast amounts of data on drug-gene interactions, but only for a small number of those interactions do we have actionable recommendations to follow (e.g. increase dose by 25% for ultra-rapid metabolizers; or contraindicated in individuals with loss-of-function metabolizing genes). There are significant challenges with antidepressant therapy owing to the heterogeneity of psychiatric disorders among patients. Athreya et al. explored the integration of machine learning with validated PGx single nucleotide polymorphism (SNP) biomarkers, which accurately predicted the 8-week treatment response to two SSRI antidepressants, citalopram and escitalopram.[19] Their research highlights that leveraging machine learning methods in PGx analysis can help to improve treatment choice, and ultimately, efficacy and safety outcomes for patients. When we look at direct patient care, PGx has implications in medication optimization, but from the perspective of the biopharmaceutical industry, it remains an untapped area within drug discovery and development.

Clinical decision support

Unsupervised learning models built around anomaly detection are being used to help develop clinical decision support (CDS) tools, including drug alert notifications in the EHR. Due to the limitations of current rules-based CDS tools, such as alert fatigue, there is a growing need for more "intelligent", dynamic, and adaptable CDS technology.

MedAware, for example, has built a CDS tool which identifies potentially inappropriate medication orders based on what is typically ordered for similar patients; it does so by looking at the patient's clinical status, baseline characteristics, and medical history.[20] If the medication order – whether the drug, route, or dose – deviate significantly from what is commonly ordered for similar patients, an alert will pop up to advise the user to double-check the order. It also provides the percentages of medication(s) that are most commonly seen ordered in a given situation for similar patients, in case the user accidentally entered the wrong medication or perhaps missed a key patient lab, such as a hepatic or renal function indicator. This model is predicated on the idea that the most common therapeutic interventions are appropriate and something that deviates from that could be an error. However, as we are aware, this is not always guaranteed to be the case. There are always exceptions to the rule.

Wearables

Wearables are becoming more and more prevalent in society today and have rapidly taken off amongst consumers. In healthcare, these devices employ a number of algorithms to assist the user in detecting, tracking, and monitoring biometric data. Wearables can also be used to augment physiological functions. OrCam MyEye, for example, intended for visually impaired individuals, is one that, when attached to eyeglasses, will help with reading text and facial recognition.[21] Livio AI, developed by Starkey Hearing Technologies, is a hearing aid product that leverages natural language processing to translate speech for people with hearing impairments.[21] Other wearables currently in the market monitor biometric data such as heart rate and ECG; their utility is in providing real-time, continuous data and, depending on the purpose and regulation-status of the device, helping inform diagnostics, treatment, and patient outcomes.[22] Many may be aware of innovative continuous glucose monitoring systems such as the Freestyle Libre, the data of which, if combined with laboratory values, blood work, patient demographics, and other pertinent health data, can be used for algorithm development, with the goal of improving treatment paradigms.

Adherence

Companies like AiCure use facial recognition and deep learning to monitor patient adherence to investigational drugs. When a patient in a study is about to take a dose of their medication, for example, the patient would open the AiCure app on their phone and hold the camera towards their face. The app uses image recognition to confirm patient identity and analyzes patient movements to confirm that the medication was ingested. The company also uses machine learning techniques to monitor facial expressivity and capture level of patient engagement.

Chat bots

Chat Bots allow companies to create new and scalable interfaces for interacting with patients and providing important, current, and relevant medical information or education. Although they have been used in industries outside of healthcare for a while, healthcare is beginning to see more adoption as the technology improves. Common use cases for health-related chat bots include diagnostic symptom checkers, triage applications, and applications designed to provide clinical drug information.

How to stay current and stay smart

As mentioned in Part 1, we need to differentiate between hype and reality. Revisiting the concept of AI winters (introduced in Part I), it's important to understand the role that public opinion plays on AI marketing and terminology. During an AI winter, when there is a decrease in general interest in AI, it is often accompanied by changes in marketing and branding to offset public perception. For example, the term "AI" might be replaced by other terms, such as "cognitive computing" or "cognitive machines." Some companies have already made the decision to avoid using the term "artificial intelligence" to describe their operations in anticipation of the possibility of a shift in public opinion and the onset of another AI winter. Other companies, meanwhile, are capitalizing on the current hype cycle and branding everything as "AI-powered." However, as we have seen, not all AI models are the same. Deep learning models are very different from logistic regression, but both could be labelled by a company as state-of-the-art artificial intelligence tools. Ultimately, it is important to keep the following in mind:

- Just because there is a lot of hype around AI, currently, does not mean that it will be able to live up to potentially inflated expectations (as evidenced by its rocky history). Be skeptical of any too-good-to-be-true claims about AI's ability.
- Never take a company's (or researcher's) claim of using AI/ML in healthcare at face value. "Artificial intelligence" by itself is not specific enough and should be accompanied with a description about the type of AI model or methodology used. Be cautious about companies or articles that do not provide those details.

Resources List

What are things you encounter on a day-to-day basis in your field that might be impacted by AI? How do you get up to speed on what's currently out there? The field of AI is growing rapidly, with new publications and products coming to market on a daily basis. It is a vast discipline with obscure and overlapping subfields, constantly evolving information and endless new applications. It can be challenging to navigate your learning of AI, but do not be discouraged! A helpful suggestion in staying up-to-date is to focus your literature surveillance on AI literature within a specific disease state of interest to you, especially at the beginning, when you are still familiarizing yourself with the topic. Here are some areas to get you started:

Curated news feeds to help stay up to date

- Healthcare IT News
- Health IT Analytics
- HIT Infrastructure
- AI in Healthcare

Academic journals focused on digital health

- Digital health (Sage Journals)
- npj Digital Medicine (Nature Partner Journals)
- The Lancet Digital HealthDigital Medicine (International Society of Digital Medicine)
- Journal of Medical Internet Research (JMIR)
- International Journal of Digital Healthcare (IJDH)
- Frontiers in Digital Health

Helpful articles to expand upon sections above

- What you need to know before building predictive models
- Biopharmaceutical companies using AI in their drug discovery processes
- Comprehensive list of startups exploring AI in drug discovery
- Use of AI in various Pharma settings along the drug development cycle
- Helpful resource for AI use cases, from marketing and sales to fintech
- An inspiring read on the use of AI in drug repurposing
- A detailed and informative read on Deep Learning in Pharmacogenomics
- A look at how NLP is being used to address problems in healthcare

References

1. Esman G. Splunk and Tensorflow for Security: Catching the Fraudster with Behavior Biometrics. Splunk-Blogs. https://www.splunk.com/en_us/blog/security/deep-learning-with-splunk-and-tensorflow-for-security-catching-the-fraudster-in-neural-networks-with-behavioral-biometrics.html.
2. Types of Artificial Intelligence: A Detailed Guide - Certes. Certes. https://certes.co.uk/types-of-artificial-intelligence-a-detailed-guide/. Published December 20, 2018.
3. Gandhi R. Support Vector Machine — Introduction to Machine Learning Algorithms. Medium. https://towardsdatascience.com/support-vector-machine-introduction-to-machine-learning-algorithms-934a444fca47. Published June 7, 2018. Accessed March 29, 2020.
4. Yse DL. The Complete Guide to Decision Trees. Medium. https://towardsdatascience.com/the-complete-guide-to-decision-trees-28a4e3c7be14. Published April 17, 2019.
5. Hansen Peter Wæde, Clemmensen Line, Sehested Thomas S.G., et al. Identifying Drug–Drug Interactions by Data Mining. Circ Cardiovasc Qual Outcomes. 2016;9(6):621-628.
6. Fleming N. How artificial intelligence is changing drug discovery. Nature. 2018;557(7707):S55-S57.
7. Keshavarzi Arshadi A, Salem M, Collins J, Yuan JS, Chakrabarti D. DeepMalaria: Artificial Intelligence Driven Discovery of Potent Antiplasmodials. Front Pharmacol. 2019;10:1526.
8. Zeng X, Zhu S, Lu W, et al. Target identification among known drugs by deep learning from heterogeneous networks. Chem Sci. 2020;11(7):1775-1797.
9. Woo M. An AI boost for clinical trials. Nature. 2019;573(7775):S100-S102.
10. Courtright KR, Chivers C, Becker M, et al. Electronic Health Record Mortality Prediction Model for Targeted Palliative Care Among Hospitalized Medical Patients: a Pilot Quasi-experimental Study. J Gen Intern Med. 2019;34(9):1841-1847.
11. Jin M, Bahadori MT, Colak A, et al. Improving Hospital Mortality Prediction with Medical Named Entities and Multimodal Learning. arXiv [csCL]. November 2018. http://arxiv.org/abs/1811.12276.
12. Cricket Health News. Cricket Health Announces New Machine Learning Model for Detecting Chronic Kidney Disease – Cricket Health. Cricket Health. https://www.crickethealth.com/cricket-health-announces-new-machine-learning-model-for-detecting-chronic-kidney-disease/. Published November 6, 2019.
13. Parikh RB, Gdowski A, Patt DA, Hertler A, Mermel C, Bekelman JE. Using Big Data and Predictive Analytics to Determine Patient Risk in Oncology. Am Soc Clin Oncol Educ Book. 2019;39:e53-e58.
14. Huizenga JE, Breneman BC, Patel VR, A. R, Speights DB. NARxCHECK Score as a Predictor of Unintentional Overdose Death. https://apprisshealth.com/wp-content/uploads/sites/2/2017/02/NARxCHECK-Score-as-a-Predictor.pdf.
15. DeNovo Summary 180001. https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180001.pdf.
16. Marmar CR, Brown AD, Qian M, et al. Speech-based markers for posttraumatic stress disorder in US veterans. Depress Anxiety. 2019;36(7):607-616.
17. Brown B. Voice Analysis to Diagnose and Monitor Health. Health Tech Insider. https://healthtechinsider.com/2017/12/13/voice-analysis-diagnose-monitor-health/. Published December 13, 2017.
18. Jaworska N, de la Salle S, Ibrahim M-H, Blier P, Knott V. Leveraging Machine Learning Approaches for Predicting Antidepressant Treatment Response Using Electroencephalography (EEG) and Clinical Data. Front Psychiatry. 2018;9:768.
19. Athreya AP, Iyer R, Wang L, Weinshilboum RM, Bobo WV. Integration of machine learning and pharmacogenomic biomarkers for predicting response to antidepressant treatment: can computational intelligence be used to augment clinical assessments? Pharmacogenomics. 2019;20(14):983-988.
20. Rozenblum R, Rodriguez-Monguio R, Volk LA, et al. Using a Machine Learning System to Identify and Prevent Medication Prescribing Errors: A Clinical and Cost Analysis Evaluation. Jt Comm J Qual Patient Saf. 2020;46(1):3-10.
21. de Jesus A. Using Wearable Data for Artificial Intelligence Applications – Current Use Cases | Emerj. Emerj. https://emerj.com/ai-sector-overviews/using-wearable-data-for-artificial-intelligence-applications-current-use-cases/.
22. Tran V-T, Riveros C, Ravaud P. Patients' views of wearable devices and AI in healthcare: findings from the ComPaRe e-cohort. NPJ Digit Med. 2019;2:53.