Analysis on the Effects of Machine Learning within the Healthcare System

Thesis
Presented to the Faculty of the
School of Engineering and Applied Science
University of Virginia

By

Worthley Burke

May 1st, 2020

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Signed: ____

Approved: ______________________________________ Date ________________________
Rider Foley, Department of Engineering and Society

Introduction

Machine learning (ML) is a new field of computer science that has rapidly grown over the past decade. ML is a subset of artificial intelligence that is commonly used to find patterns within data. It uses old data to train its systems to assess the probability that new data fits historical patterns. If the algorithm is given enough training data, it can perform better than a human at certain tasks. ML algorithms can even be trained to dynamically react to new situations like a human. As such, it is currently being used for many functions in everyday life. Face recognition and Google's self-driving car are two high profile examples (SAS Insights, 2019). The code that the algorithm generates to interpret data is unknown to the coder, removing some forms of human error from the coding process. While machine learning has the potential to be more accurate than a human generated algorithm, if the ML algorithm is given an inaccurate data set by the human coder, the algorithm would be ruined.

Machine learning algorithms can be split into four groups based off of the algorithms' purpose: supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning. Supervised learning takes labeled data that has an input and a correct output. The algorithm will try to input data into a function that will produce the correct output, such as in a linear regression algorithm. It could also try to identify certain features that would denote a certain output, such as in neural network algorithms. Unsupervised learning is when the algorithm is fed data without a desired output. The algorithm then detects patterns from this data. Semi-supervised learning is fed a mix of labeled and unlabeled data. Reinforcement algorithms run on a feedback loop where if the algorithm chooses the "right" action, it is given an award (Fumo, 2017). While each of these algorithms are used for different purposes, they can all be used in the medical field.

Currently in the healthcare system both patient volume and patient cost are increasing. Bhardwaj and his colleagues (2017) argue that machine learning can reduce the costs of the healthcare system while providing patients with better care. As machine learning begins to assist doctors with various pre-diagnostic analyses, the cost per patient visit is expected to decrease. This will also allow doctors to spend less time on each patient, enabling more patients to be treated. Additionally, machine learning can reduce the need for patients to go to hospitals, decreasing patient cost and volume. Almost 90% of Emergency Room visits are preventable. Machine learning can diagnose and direct patients to proper care without the patients needing to go to expensive hospitals (Bhardwaj et al., 2017). Since machine learning will continue to be adopted by the health field in order to access its potential benefits, it is necessary to analyze how it will alter the healthcare system.

Case Context

A wide variety of machine learning algorithms have been recently introduced to multiple areas throughout the medical field, from personalizing individualized patient care to cutting costs for massive organizations. As the medical records grow, machine learning algorithms are able to take advantage of the vast amount of information. KenSci, a company from Seattle, uses predictive machine learning to determine which patients are most likely to become sick. This will allow doctors to proactively treat their patients before they become deathly ill (KenSci, 2017). KenSci trains their algorithm with data from electronic medical records. While the vast storage of data is very helpful, correctly implementing it can sometimes cause problems. Quotient Health uses machine learning to lower the costs of healthcare providers. They analyze a company's collection of electronic medical records to reduce the cost of maintaining the system.

Specifically, they use a deep learning algorithm to find issues within the system before they detrimentally affect the company (Welcome AI, 2019).

While many machine learning algorithms are used as a predictive tool, some algorithms help patients in real time. Patients are often monitored by wireless sensor networks to ensure their well-being while they are away from doctors and nurses. While these wireless sensors provide the doctors with valuable patient data, they can fail. When they fail, corrupted patient data is sent to the doctor instead. A linear regression machine learning algorithm has been developed that determines whether a patient is in poor health, or if the sensor is recording faulty readings (Salem et al., 2014). This technology will determine whether the patient will need a life-saving ambulance ride to the hospital, or if they need to simply replace their monitoring device.

One of the most exciting places of machine learning use in the medical field is within the radiology department. Using neural networks, machine learning algorithms can diagnose patients by finding patterns that doctors cannot identify. Enlitic, Inc. is currently at the forefront of combining machine learning with image diagnostics. "By pairing world-class radiologists with data scientists and engineers, we collect and analyze the world's most comprehensive clinical data, pioneering medical software that enables doctors to diagnose sooner with renowned accuracy" (Enlitic, Inc., 2019). Enlitic's algorithm can interpret X-Ray, MRI, and CT images in less than 15 milliseconds (Enlitic, Inc., 2019). This allows doctors to accurately prioritize patients in real-time, while making their diagnosis process more efficient. They combine these images with patient histories and physician notes within the algorithm to most accurately analyze the pattern. The precision of the algorithm is groundbreaking: "Enlitic technology detected lung cancer nodules in chest CAT Scan (CT) images 50 percent more accurately as compared to an expert panel of radiologists" (Bhardwaj et al., 2017). This algorithm always ensures that the doctor is aware of the reasoning behind insights (Bhardwaj et al., 2017). By incorporating radiologist feedback afterwards, the algorithm will iteratively improve itself depending on whether it is over or under-calling features (Enlitic, Inc., 2019). Currently at UVA there are projects going on in the radiology department that use images to predict future conditions. Students are designing an algorithm that can determine whether or not a person will develop brain tumors from MRI generated images. These examples of machine learning offer a good summary of how machine learning is used, but they just scratch the surface of the depth to which it is used. It is crucial to see how integrating machine learning affects the stakeholders of the healthcare system.

Analysis of Machine Learning in Health Care

This project may contribute to the ongoing introduction of machine learning into the healthcare system, which is a complex sociotechnical system. The main stakeholders within the healthcare system are hospitals, doctors, patients, insurance companies, and research institutions. Theoretically, adopting machine learning should be beneficial for all four stakeholders:

- hospitals and doctors would be able to provide better patient care,
- patients would receive a better service for a cheaper price,
- insurance companies could charge more accurate rates, and
- research institutions could analyze experiments more efficiently.

As shown in Michael Harrison's paper the theoretical effects of a new artifact do not always align with the actual effects. Unfortunately, Harrison's Interactive Sociotechnical Analysis (ITSA) approach relies heavily on the recursive interactions of the system (Harrison et al., 2007). As machine learning has only been recently introduced in the healthcare system, ITSA cannot best analyze the effect of machine learning. Instead, I will use Latour's (1992) Actor Network Theory to analyze how machine learning interacts with stakeholders in the healthcare system. Actor Network Theory (ANT) proposes that people, institutions, and organizations affect the development of a sociotechnical system and that artifacts have an equal role in this process, which is called symmetry. Latour shows this by analyzing how installing a door closure device alters the behavior of people walking through the door (Latour, 1992). As machine learning is such a powerful tool, it will change the behavior of the key stakeholders within the system. The process of a technology affecting multiple actors within a system is known as *delegation to technology* (Waelbers, 2009). For example, it could lower costs to patients resulting in more patients being able to afford healthcare and increasing doctors' patient throughput. It could also possibly be used to alter health insurance policies to maximize their profitability. It is crucial to use ANT as it will allow me to analyze machine learning's interaction with other technologies. Most machine learning in healthcare will build-off of and assist currently existing technologies. For example, radiologists are using machine learning to detect tissues from imaging scans, machine learning algorithms are analyzing medical databases to form disease prediction software, and machine learning algorithms are integrated with ICD10 records to find hospital trends (Bhardwaj et al., 2017; Cabitza et al., 2017).

I have provided an ANT model below of a machine learning algorithm created by MIST LLC, a biotech startup. The algorithm is used to detect the curvature of the spine. The blue boxes denotate technologies, the green ovals denotate actors, and the orange diamonds denotate the relationship. This model is simplified as it assumes all doctors will be employed directly by hospitals. As shown in the diagram the machine learning algorithm both assists doctors, and analyzes X-rays to diagnose patients. These two actors can then affect the rest of the system besides the government. Patients will pay the insurance company to cover their bills in case of bad health. Insurance companies will pay the hospitals for the cost of the patients' visits. Doctors will work at the hospitals, and the hospitals will pay the doctors. Hospitals will also have certain rules for their doctors and enforce a standard of care. The hospital is then regulated by the government to abide by certain practices. By just analyzing this specific algorithm, ANT highlights how machine learning can affect the interactions between the artifacts. This model emphasizes how the introduction of the machine learning algorithm could affect all other actors, besides potentially the government. For example, if the algorithm assists doctors, this could reduce the time they spend treating the patient, which would reduce the cost the insurance company pays the hospitals. This would cause the hospital to pay the doctors less, and the insurance company would charge the patient less. Without ANT, it would be more difficult to see the intricacies of this complex network of relationships.

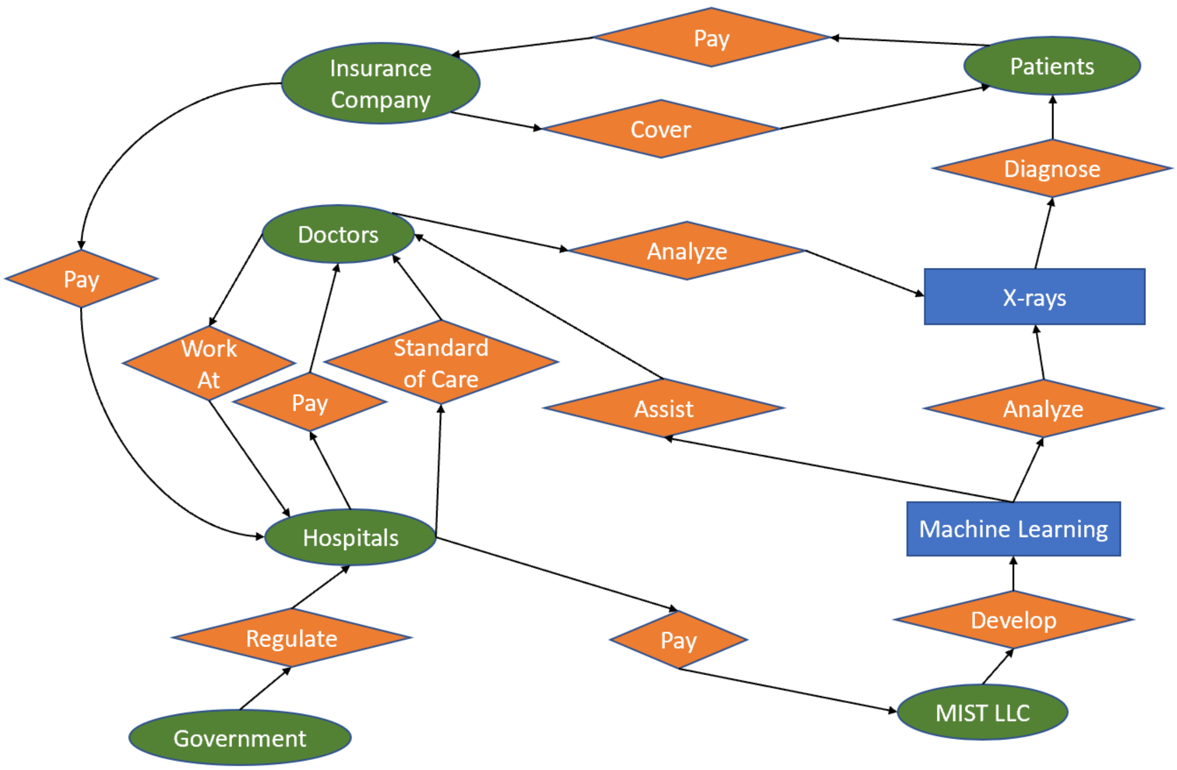

Figure 1. Actor Network Theory Diagram of MIST Diagnostic Algorithm. This diagram shows the ANT web of the machine learning algorithm that my team is producing. (Burke, 2019).

Research Question and Methods

Due to the potential of machine learning, the healthcare industry, along with most others, has begun to incorporate ML into their practices. To date companies have built ML algorithms for myriad healthcare functions such as smart records, medical imaging and diagnostics, drug delivery and development, medical data, and treatment and prediction of disease (Thomas, 2019). As with any new technology introduced into an existing sociotechnical system, its adoption is bound to result in changes within the society. The incorporation of ML into the healthcare system has the potential to produce massive change, as such it is important to analyze how it will alter relationships between stakeholders and technology. Some researchers speculate that machine learning could even be harmful for the healthcare system. One detriment of machine learning could the deskilling of the physicians. In one study residents had a 14% decrease in diagnostic accuracy when images were marked with computer-aided detection (Cabitza et al., 2017). This raises two important questions. While machine learning has made the healthcare industry more efficient, has it introduced unintended consequences? As machine learning has been incorporated into the healthcare field, how has it altered interactions amongst people and technology?

To answer these questions, I conducted interviews and researched case studies. I conducted interviews with doctors, government researchers, and IT consultants working in the health field. Their firsthand accounts gave me insight concerning machine learning adoption across multiple areas of healthcare. I supplemented this research with case studies on a variety of machine learning uses. I specifically looked at organizations that utilize ML throughout the healthcare industry and analyzed how machine learning is changing the healthcare system.

Name	Employer	Job Title	Experience with ML	Responsibilities
Dr. Sana Syed	University of Virginia	Assistant Professor	4 Years	75% Research 25% Clinical Work
Allision Horenberg	National Institute of Standards and Technology	Undergrad Research Fellow	6 Months Over 2 Internships	Wet Lab Research
Ashley Walton	Deloitte	Technology Consultant	5 Years	Worked with: NIAD National Cancer Institute National Institute of Health
Brendan Abraham	MITRE	Data Scientist	7 Months	Healthcare Analytics

Table 1. Interview Details. This table contains background information about each person I interviewed. (Burke, 2020).

Results

Surprisingly, I found that while machine learning has shown an increased trend within the health field, it has not yet been strongly incorporated in clinical practice. Rather, it has been most often used by research organizations, hospitals, and insurance companies to improve operations. The adoption of machine learning has most noticeably changed the healthcare system by introducing machine learning algorithm providers into the system. I have also uncovered the main complications that are halting the adoption of machine learning: low data accessibility and low artificial intelligence literacy.

While there is talk of using machine learning on the clinical side of the medical field, the day when that is commonplace is not here. Dr. Sana Syed, as an Assistant Professor at the University of Virginia (UVA), splits her time between research and the pediatric clinic. According to Dr. Syed there are two main reasons machine learning does not yet have a larger clinical presence: “the black box quality of machine learning and the high stakes of a human life” (Dr. Sana Syed, personal communication, March 20, 2020). Many ML algorithms function within a black box. Even though a programmer will set up an algorithm, he or she will have no idea how the algorithm makes a final decision once it has been trained. This makes doctors warry to trust these algorithms as there is no way to know what the algorithms’ deciding criteria are. Doctors are less likely to adopt a new technique, particularly such an opaque algorithm, as one error could result in a patient’s death. That is why machine learning algorithms are being so strongly tested before they are put into practice.

Even a machine learning algorithm used at UVA only in an advisory capacity, underwent a long iterative adoption process before it was integrated into the system. The radiology department has adopted a beta version of a machine learning algorithm called *Radiology Assistant*. The algorithm works in an advisory fashion: it flags CT scans if they contain any abnormal features. This algorithm “is helping the radiologists identify findings that they couldn’t see before” (Jones, 2017). It can detect low bone density within the patient, catching the condition before the fractures that would notify a radiologist occur. This early detection improves the quality of life and helps lower costs to the health system (Jones, 2017). Although this technology has the potential to diagnose patient complications earlier, greatly improving the quality of care, the radiology department has only very cautiously begun to incorporate it.

Research organizations, hospital operations, and insurers have been much quicker to adopt this technology. In an interview with Allison Horenberg, she spoke of the clear efficiency boost her government lab at National Institute of Standards and Technology (NIST) received from adopting a machine learning approach (A. Horenberg, personal communication, March 4, 2020). Ashley Walton also spoke of how the National Institute of Allergies and Infectious Diseases (NIAID) has begun to use machine learning to classify disease images and within their databases (A. Walton, personal communication, March 19, 2020). This quicker adoption stems from the desire to harness machine learning’s proven ability to increase efficiency, without worrying about the high stakes of death.

The use of machine learning within the medical field has opened a niche for new companies to interact in the health field. These companies have a strong technical background which they use to build ML algorithms for healthcare organizations. Consulting firms, like Deloitte, have been hired to improve database operations across several research institutions (Walton 2020). Information Technology companies, such as MITRE, have also shifted to using machine learning for cases such as hospital quality analytics (B. Abraham, personal communication, March 5, 2020). While Consulting and IT companies have adapted their business practices to incorporate ML, other companies specialize in providing ML solutions. Inovalon is one company that specifically utilizes ML to analyze large datasets within the healthcare field. Inovalon uses ML approaches to optimize operations within insurance companies. By analyzing clinical records, they were able to save one insurance company 3.54 million dollars in risk accuracy improvement (Inovalon, 2020a). They also cut time required for the clinical record review process by half for another undisclosed organization (Inovalon, 2020b).

While the current relationship between actors seems to work well on the surface, from my research I have found two major barriers to adopting ML in the healthcare space. The regulations and systemic storage of patient data has resulted in extremely low data accessibility. Additionally, the process of machine learning adoption skipped a crucial step: promoting AI literacy within the system. These two qualities make it difficult to adopt ML.

Data accessibility proves to be a unique problem within the healthcare field for two reasons: the use of data is strongly regulated by the Health Insurance Portability and Accountability Act (HIPAA) and healthcare's inadequate adoption of digital records. HIPAA regulations were created to protect the patient's privacy and set "limits and conditions on the uses and disclosures" of the patient data (U.S. Department of Health & Human Services, 2008). This means that patient data cannot be used for most purposes, including the creation of ML algorithms, without patient approval. Additionally, even if patients have permitted one organization to use their data, due to the HIPAA laws, data cannot be freely shared amongst other hospitals and organizations. This situation is known as data siloing. Thus, developers of ML algorithms face a difficult dilemma: use a smaller, likely biased, database resulting in a less accurate algorithm, or spend extra time and effort obtaining the rights to more data, slowing down the project. Brendan Abraham emphasized the complications this added to his work as he analyzed quality models across multiple hospitals (Abraham, 2020).

Driven by a desire to keep patient data safe, the healthcare system has been slow to adopt digital records. Only in 2009 did the US government make a concentrated effort to digitalize the healthcare space. By offering stimulus money to hospitals, the US government was able to promote quick adoption of digital records, but at the cost of usability. A decade later while “96% of hospitals and 86% of physicians’ offices in the United States have access to electronic health records,” the systems do not work effectively. In a 2018 survey, 59% of physicians said these electronic records needed a complete overhaul (Hecht, 2019). These inefficient systems make it difficult for doctors to enter data, resulting in patient data loss. During the development of the systems “various vendors had separately developed systems that formatted data in different ways, which made it hard to share records” (Hecht, 2019). This directly contributes to data siloing. Even once the data was collected, the erratic formatting made it extremely difficult to normalize the data for use by a ML algorithm. Combined with the strict HIPAA regulations, the low accessibility of data in the healthcare space make the use of machine learning a daunting task.

Another obstacle within the health space is a lack of AI literacy. Allison Horenberg recounted how she “was the only one within [her] lab that knew how to use a machine learning algorithm” (Horenberg, 2020). This disrupts machine learning adoption in two ways. Ashley Walton explained how a lack of AI literacy creates a major discrepancy in projects between the organization team and machine learning build team and results in wasted time and inferior algorithms. On the healthcare side of things, there is an overreliance on the technical staff, and a hesitation to learn about ML due to the perceived complexity. The technical staff knows the exact limitations of their product; however, "they do not know enough about the health domain to have their algorithms answer the right questions. The medical field does not know the restrictions of the algorithms" (Walton, 2020). This often restricts the ability of a ML algorithm to add the most value: either due to the type of question asked, or the ML approach used to answer it.

The other problem that arises from low AI literacy is the incorrect application of a built algorithm. In the extreme sense, a ML team could build out the algorithm, only to have its use denied by upper level management (Walton, 2020). In other cases, it will take convincing to get key stakeholders to accept the results of an ML algorithm (Abraham, 2020). Even once the algorithm is incorporated someone not well versed in statistics could draw an incorrect conclusion from the algorithm's results. Ashley Walton referenced a machine learning algorithm she had worked on concerning a tuberculosis database. Some of the cross references between other diseases, such as breast cancer, were only logged for women. This could result in an uniformed user unknowingly receiving biased statistics. In another study, a machine learning program drew a correlation between asthma and a higher chance of surviving pneumonia. This was due to people with pneumonia and asthma being admitted directly into the ICU. So, the ML program misidentified asthma as a protective condition against pneumonia (Cabitza et al., 2017).

Discussion

The adoption of machine learning has thoroughly altered the healthcare system as shown above. To analyze how it has altered relationships amongst actors I will incorporate a generalized ANT model. The adoption of ML created a new type of company in the healthcare system. These “machine learning companies”, like Inovalon, exist to develop algorithms for healthcare entities. To compete for profit with these ML companies, existing IT and consulting companies expanded their capabilities in order to build algorithms for healthcare organizations. ML, IT, and consulting companies focus their algorithms on making digital records analysis, and organization operations more efficient and streamlined. This process can save the health organizations millions of dollars, that in turn result in them charging patients less. For insurance companies specifically, their risk analysis systems have been optimized by machine learning algorithms, saving both the patients and the companies money.

Within research, ML is beginning to be commonly used, even by government organizations. Research companies will either develop their own machine learning algorithms, or they will pay for outside experts, the previously mentioned companies, to incorporate machine learning into their research. The use of ML increases the productivity of the lab, as a ML algorithm can easily provide a “high quantity of throughput data” (Horenberg, 2020). This allows researchers to focus on running more experiments, instead of just analyzing numbers.

ML has not begun to support physicians in the clinic. Physicians typically do not trust an algorithm when a human life is at risk. Even the most innovative departments have only began to test out ML with beta systems. Figure 2 shows an ANT diagram for a generalized analysis of ML in healthcare system. This diagram is generalized in the sense that while there may be some cases of machine learning already being used clinically, in general, doctors have not adopted machine learning in the clinic.

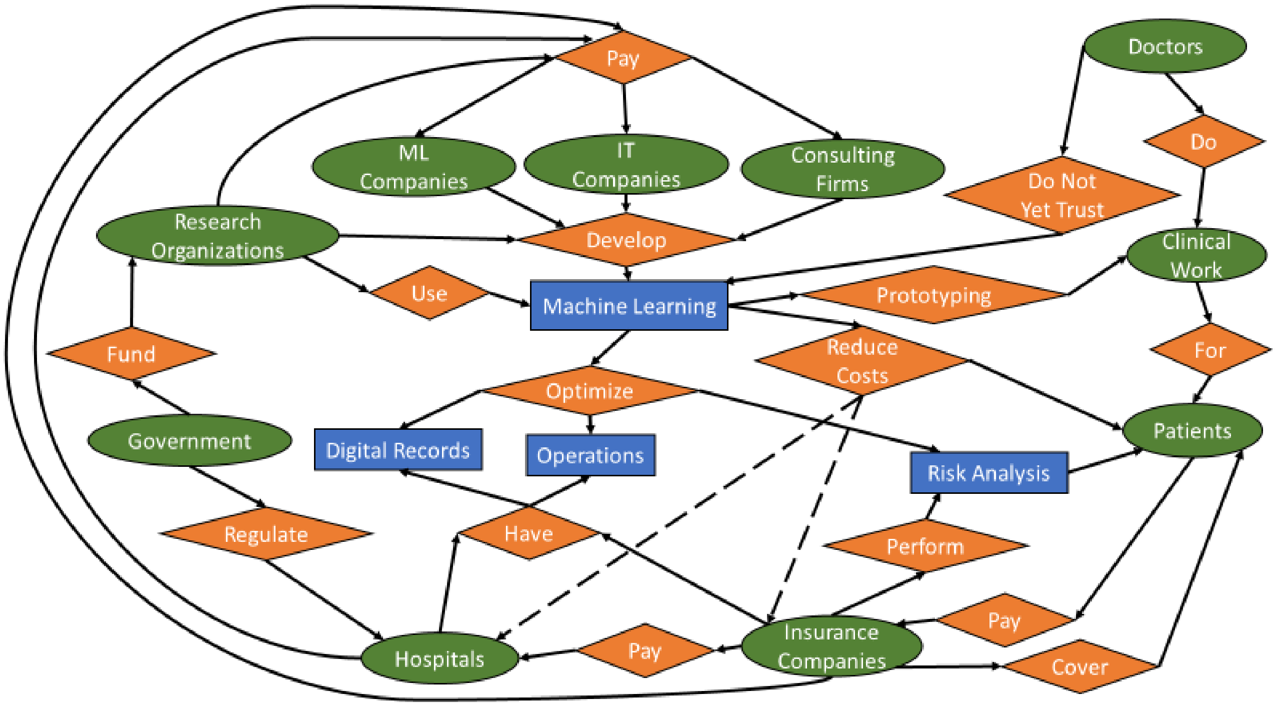

Figure 2. Machine Learning within the Healthcare System. This shows the current healthcare system with the effects of machine learning. The blue boxes denotate technologies, the green ovals denotate actors, and the orange diamonds denotate the relationship. In the case of "Operations", the box refers to the technologies/processes that currently determine the organizations' operations. Dotted arrows were used to assist visualization (Burke, 2020).

A major limitation to my research was the novelty of machine learning being used in the medical field. The Food & Drug Association only provided restrictions on medical machine learning algorithms in December 2017 (Software as a Medical Device Working Group, 2017). There has been lots of research on machine learning, but almost no cases in which it was an entrenched part of a healthcare system. This made it very difficult to analyze what the long-term effects of machine learning could be. Additionally, there were only a few doctors who have been able to use ML in any capacity. Since the purpose of this paper was how ML changed the interactions between stakeholders, the dearth of physicians made my analysis using ANT significantly more difficult.

In the future I will strive to keep my mind more open and free of biases. Since my capstone project dealt with the use of machine learning in a clinical application, I assumed the majority of my findings would be clinically based. It was in fact quite the opposite. Without my predisposition I would have altered many of my interview questions. Also, due to the lack of responsiveness I would reach out to a wider group of interviewees. By contacting more people, I could have found another interviewee with a unique view on the situation.

I will use this research to remember the most important part of the engineering process, is that the final product is used. This has shown me that adoption of new technology is a long and arduous process. If I develop a new technology, I will make sure to spend just as much time promoting adoption of it, as I will designing it. Once the technology begins to be widely used, then I can focus all my efforts on improving my product.

Conclusion

The majority of the data I have identified agrees that machine learning can function effectively to improve the medical field when it is used correctly; however, through the testimony of my interviewees, I found the current set up is not conducive to machine learning. The barriers of low data accessibility and low AI literacy within the health field will continue to limit the potential of ML in the future. The question becomes: how do we efficiently use machine learning within the healthcare system?

First, data needs to be made more accessible for the algorithms. This could happen in three ways. One, overhauling the current system and relaxing HIPAA laws, allowing for centralized and deidentified data. Two, the generating fake patient data through simulations. This is currently happening at IT companies like MITRE to get around data constrictions. By using national averages, it is possible to create a large amount of statistically accurate patient data that can be used to train a ML algorithm. This fake data will contains the biases of averages without the same specificity as real data (Abraham, 2020). Three, taking a non-healthcare approach to solving healthcare's problems. Companies like Apple and Fitbit have already amassed a large amount of patient data that they own the rights to. These companies would be able to run their data through machine learning algorithms and then sell their results to the medical field. This might be the most likely scenario to occur, due to the data's centralized ownership, accuracy, accessibility, and standardized format.

Once data is made available, AI literacy within the healthcare community needs to be improved to productively use ML algorithms. This will allow medical professionals to adopt ML in the scenarios most helpful to them. It will also increase healthcare organizations efficiencies and improve their operations. We need to promote more education such as found in American College of Radiology Bulletin machine learning special section to promote literacy (Jones, 2017). "The technology is not too difficult to understand", there is simply not a strong enough push by the medical field to learn about it (Walton, 2020). If the healthcare system is able to utilize machine learning correctly, there is the potential to drastically reduce healthcare costs while maximizing patient care.

References

Abraham, B. (2020, March 5). *Machine Learning Interview 2* [Personal communication].

Bhardwaj, R., Nambiar, A. R., & Dutta, D. (2017). A Study of Machine Learning in Healthcare. *2017 IEEE 41st Annual Computer Software and Applications Conference (COMPSAC)*, *2*, 236–241. https://doi.org/10.1109/COMPSAC.2017.164

Cabitza, F., Rasoini, R., & Gensini, G. F. (2017). Unintended Consequences of Machine Learning in Medicine. *The Journal of the American Medical Association*, *318*(6), 517–518. https://doi.org/10.1001/jama.2017.7797

Dr. Sana Syed. (2020, March 20). *Machine Learning Interview 4* [Personal communication].

Enlitic, Inc. (2019). *Enlitic*. Enlitic.Com. https://www.enlitic.com

Fumo, D. (2017, August 17). *Types of Machine Learning Algorithms You Should Know*. Medium. https://towardsdatascience.com/types-of-machine-learning-algorithms-you-should-know-953a08248861

Harrison, M. I., Koppel, R., & Bar-Lev, S. (2007). Unintended Consequences of Information Technologies in Health Care—An Interactive Sociotechnical Analysis. *Journal of the American Medical Informatics Association*, *14*(5), 542–549. https://doi.org/10.1197/jamia.M2384

Hecht, J. (2019). The future of electronic health records. *Nature*, *573*(7775), S114–S116. https://doi.org/10.1038/d41586-019-02876-y

Horenberg, A. (2020, March 4). *Machine Learning Interview 1* [Personal communication].

Inovalon. (2020a). *CDEaaS™ Case Study*. https://www.inovalon.com/resourcelibrary/documents/cdeaas-case-study/

Inovalon. (2020b). *NLPaaS^{TM} Case Study*. https://www.inovalon.com/resourcelibrary/documents/nlpaas-case-study/

Jones, J. (2017). Case: Extra Set of Eyes. *Imaging 3.0*. https://www.acr.org/-/media/ACR/Files/Case-Studies/Imaging3_An-Extra-Set-of-Eyes_July2017.pdf

KenSci. (2017). *Transforming the care continuum*. KensSci. https://www.kensci.com/solutions/overview/

Latour, B. (1992). Where are the missing masses? The Sociology of a few mundane artifacts. In D. G. Johnson & J. M. Wetmore (Eds.), *Technology and Society, Building Our Sociotechnical Future* (pp. 151–180). MIT Press.

Salem, O., Guerassimov, A., Mehaoua, A., Marcus, A., & Furht, B. (2014, January 1). *Anomaly Detection in Medical Wireless Sensor Networks using SVM and Linear Regression Models* [Article]. International Journal of E-Health and Medical Communications (IJEHMC). https://doi.org/10.4018/ijehmc.2014010102

SAS Insights. (2019). *Machine Learning: What it is and why it matters*. https://www.sas.com/en_us/insights/analytics/machine-learning.html

Software as a Medical Device Working Group. (2017). *Software as a Medical Device (SAMD): Clinical Evaluation*. Food and Drug Administration. https://www.fda.gov/media/100714/download

Thomas, M. (2019, September 25). *Ultra-modern Medicine: Examples of Machine Learning in Healthcare*. Built In. https://builtin.com/artificial-intelligence/machine-learning-healthcare

U.S. Department of Health & Human Services. (2008, May 7). *The HIPAA Privacy Rule* [Text]. HHS.Gov. https://www.hhs.gov/hipaa/for-professionals/privacy/index.html

Waelbers, K. (2009). Technological Delegation: Responsibility for the Unintended. *Science and Engineering Ethics*, *15*(1), 51–68. https://doi.org/10.1007/s11948-008-9098-x

Walton, A. (2020, March 19). *Machine Learning Interview 3* [Personal communication].

Welcome AI. (2019). *Quotient Health*. Welcome.AI. https://www.welcome.ai/quotient-health