Digitization of Perioperative Medical Records

AI for Social Good: Responsible Use of Artificial Intelligence to Solve Modern Social Issues

A Thesis Prospectus Submitted to the

Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree
Bachelor of Science, School of Engineering

John Radossich
Spring, 2021

Technical Project Team Members
Mary Blankemeier
John Radossich
Sarah Rambo
Charlie Thompson

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

Signature _____ Date _____
John Radossich

Approved _____ Date _____
Donald E. Brown, Department of Data Science

Approved _____ Date _____
Hannah Rogers, Department of Engineering and Society

Introduction

As humanity transitions further into the Age of Information, the dichotomy of engineers will reveal itself as very polarized. With technological capabilities breaking barriers of what was once possible, the engineering field has grown in size and popularity. Accompanying this spike in technical interest is the division of engineers into two groups; those who answer the "how" question, and those who answer the "why" question. The latter are more focused on the underlying reasoning behind their research and development. The prior, those who answer the "how," are more focused on innovation without purpose, whether it be for monetary benefit or a true passion for engineering. As a result, this separation of motivation has led to an increasing ignorance for how this technological revolution can fix some of society's larger issues. Specifically, this trend is highlighted by the emerging technological obsession with artificial intelligence. Endless resources are being put towards machine learning and deep learning algorithms to optimize the functions of society (Pohler, n.d.). However, very few of these solutions extend their capabilities to complex social issues that need to be addressed. With worldwide artificial intelligence spending anticipated to double over the next four years, the application of these methods should be focused on contributing to failing mechanisms of society (Soohoo, 2020).

One application of artificial intelligence that is being researched for social good is the implementation of centralized databases in third-world country healthcare systems. Currently, research is being done to enhance the current healthcare system in Rwanda, Africa. Through the establishment of clean databases and data analysis tools, healthcare decisions in low-income

countries such as Rwanda will shift towards more of a result-based practice (Wyber, 2015). Specifically, natural language processing and image recognition models are in development in order to read and translate written medical records to a centralized database. Key insights will be derived from the centralization of medical data which would eventually lead to enhanced healthcare analytics. The final deliverable, a streamlined process of uploading and translating encrypted written medical records to a secure database to be analyzed.

Technical Topic

Similar to the outbreak of data across most industries, healthcare systems in high-income countries have increasingly utilized big data to identify best modes of practice. Conversely, developing countries have been left behind in this respect, lacking the resources necessary to establish these databases. Rwanda, Africa is an example of one of these developing countries, one with a lack of technical integration in their healthcare system. The consequence of an anecdotally driven system is the exclusion of key findings hidden amongst millions of observations. Specifically, there is an overwhelming need for the digitization of perioperative (data collected during surgery) data flowsheets. These sheets contain data that would contribute to the understanding of perioperative mortality rate, determine the aspects of adverse surgeries, and gauge the overall quality of treatments in the hospital (Feng, 2017). Furthermore, these doctors will understand the correlations between surgical procedures and medicine administered, developing a deeper understanding of trends surrounding the efficacy of certain treatments.

In an effort to alleviate the lack of data availability in Rwanda, the UVA Health System has sponsored a project to construct a system that facilitates the digitization of patient medical

records. Prior to understanding the development and implementation of such a system, it is integral to understand the many constraints associated with the task. To start, as the data involves patients of the Rwandan hospital, it is mandatory to uphold the highest standard of patient confidentiality at all times. Similarly, it is essential to incorporate cryptography practices to mitigate the risk of hackers accessing the data and breaching patient confidentiality. Finally, there is a degree of financial inflexibility of such a project, requiring the team to build a system that costs very little to implement and maintain. In summary, the team hopes to streamline the process of uploading encrypted photographs of written patient medical records, passing these sheets through a natural language processing and image recognition model, and translating this data to a secure database for Rwanda's use.

In order to complete the task at hand, the team will construct a hybrid Android application in JavaScript. The application's main function comes from its integrated scanner which locates the edges of the medical record, takes a picture of the sheet, and converts the image to grayscale. Greyscale is an image processing technique prevalent when paired with image recognition models as it simplifies the algorithm and reduces computational requirements (Kanan, 2012). As the user uploads the picture, several processes take place. First, the application is disconnected from the public server and the image is saved to the phone's hard drive for asymmetric encryption. The encryption process works as follows; the picture is deconstructed into an array that contains every pixel of the image. Each pixel contains four bytes of data known as the RGBA color model – a four-digit number that characterizes the color and contrast of a single pixel. The array is then passed through a hash function which converts these four-digit numbers into an array of strings, the type of data required for image encryption. The

picture will then be sent to UVA and once behind their firewall, the image will be decrypted and reconstructed by reversing the encryption process previously described (Crypto-IT, n.d.).

Then, the images are passed through a natural language processing model to identify the medicine used in that particular surgery. Next, an image processing model is used to identify two items: the status of several patient information checkboxes and the systolic/diastolic blood pressure readings taken in five-minute intervals throughout the surgery. After these pieces of data are extracted via various machine learning models, the data will be stored in a SQL database for the use of Rwandan doctors to identify the most effective practices. With this data, doctors can analyze correlations between the variety of medicines administered and the efficacy of the surgery. Over time, these findings will develop an organized understanding of how patients should be treated under different circumstances, optimizing the efficiency and success of the hospital.

STS Topic

In today's technological revolution, the Age of Information, many companies claim to be establishing systems inclusive of machine learning and artificial intelligence. In fact, 47% of businesses had at least one AI capability in their business processes, and it is anticipated that this number will reach 87% with AI-based cloud services (Matthew Stewart, 2020). While the use of these methods is wide-ranging and certainly applicable, an obsession has stemmed from this technology primarily for social recognition and funding (Torres, 2018). The selfish view surrounding the capabilities of AI has established an extremely shallow exploration of these methods in solving social problems. Artificial intelligence is being used to analyze the stock

market, execute millions of stock trades, and identify valuable sports gambling events (Mills, 2020). Not only are these practices being ignored for social good, but in some cases, they are being used detrimentally. Machine learning models require training data to "teach" the machine the interactions of life. However, this data is often collected, cleaned, manipulated, and processed by humans, adding a layer of inherent bias to the model (Murray, 2019). These methods of data manipulation lead to the ironic addition of human error to many of these machine-driven algorithms.

As previously mentioned, both implicit and explicit biases accidentally find themselves into AI models despite no sign of malicious intent. However, independent of the developer's objective, these models often make decisions that affect the lives of others. Again, the errors in these practices come back to the human element of them. Jason Bloomberg, President of Intellyx, captures this phenomenon when he says, "As human behavior makes up a large part of AI research, bias is a significant problem, data sets about humans are particularly susceptible to bias, while data about the physical world are less susceptible" (Murray, 2019). Thus, artificial intelligence misuses typically stem from the reflection of real-world biases that find their way into the way these models are trained.

One example of how machine learning algorithms are incorrectly being used is the US court and corrections system. Specifically, a machine learning algorithm named COMPAS has been implemented in a Wisconsin court system to gauge the likelihood that a convict will commit another crime if released. While this may seem like an appropriate and socially beneficial place to integrate AI, it was found that the algorithm was biased against African American prisoners. In an investigation by ProPublica, it was found that the model incorrectly

flagged black prisoners as more likely to reoffend by over 20% (Murray, 2019). These biases have found their way into various mechanisms of the United States in similar forms. It has been proven that certain algorithms have biases against race, gender, religion, socioeconomic status, and geography (Müller, 2020). As previously stated, the underlying cause of these biases is a direct reflection of the biases that exist in our lives and translate the data that is collected. Thus, it is recognized that many machine learning processes are being inappropriately applied and require regulation before decisions are to be made from them.

While the irresponsible misconduct surrounding many of these algorithms is an issue in today's society, some would argue that the problem should be addressed at its core. That is, the selection of problems that AI is being used to solve is faulty and these technologies should be used for social good. Specifically, there are many channels of sustainability efforts through which machine learning and artificial intelligence could be integrated. Many sustainability efforts are reliant on some third-party source of energy which drives the collection of this energy. For example, wind farms, water currents, and solar rays are harvested and converted to raw energy. However, the collection process lacks optimization as these methods rely on waiting in one spot. With the investment of resources into AI, these processes would collect energy at maximum efficiency by predicting daily wind, water, and solar maximums and moving these collectors accordingly (Haughey, 2020).

This topic can be most aptly analyzed using the Actor-Network Theory and Technical Momentum STS frameworks. The Technical Momentum framework views technological development as a spectrum, with social construction of technology and technological determinism at each extreme. Social construction of technology is the idea that a society

determines the use and development of technology while technological determinism argues that the introduction of an emerging technology will alter the functions of society (Smith, 1994). Under this scope, products are evaluated as a function of time which connects the two extremes defined previously. At the conception of a technology, society has control over its use and defines how it will be integrated into society. However, as these inventions grow larger, their momentum makes them hard to control and they begin to determine how society functions. While the evolution of technological utilization makes sense inherently, the theory lacks an explanation regarding the purgatory when a technology resides between social construction and a technological determinism.

As previously stated, I will also be analyzing this topic under the Actor-Network Theory (ANT). This framework will ignore the humanity aspect of a system as ANT views human and non-human elements equally as actors within the same network (Cressman, 2009). By diminishing the biases towards a machine or a computer, the topic can be fully analyzed in terms of its effect on a society (or a network). Eliminating the weight of human element, the theory allows for an unbiased evaluation of the interactions between beings of the universe. While this theory is strongly backed and widely popular, its shortcomings reside in the lack of intentionality attributed to non-living actors. This notion contradicts the basis of the theory that all elements of the system are to be treated and evaluated equally, and introduces slight bias. I believe this to be an incredibly appropriate framework as AI algorithms are created to mimic the human decision-making process, making sense that they are equally analyzed as part of our societal network (Reed, n.d.).

Methodologies

Research Question: To what extent should machine learning and artificial intelligence be integrated into all applications of life as opposed to mastering the use of these technologies in socially impactful ways?

In terms of the research that will be conducted to answer this question, there are several sources that must be considered. To start, it will require research into the majority of use cases that currently exist and implement artificial intelligence. Next, it would be beneficial to establish an accurate timeline of the capabilities of AI today, tomorrow, and ten years from now. Next, the introduction of the downsides of AI and the technology being a direct reflection of biases that already exist. Lastly, the conclusion consists of research surrounding the use of AI for social good and thus prove that there is a fundamental issue surrounding the distribution of machine learning applications. In order to most effectively convey these areas of focus, anecdotal sources with statistical justifications will be sought after. These documents will facilitate the understanding of the current system in place, as well as the statistical extent to which this system exists.

Conclusion

The technical paper studies the use of artificial intelligence through the digitization of Rwandan hospital medical records. The team will design, create, and implement a hybrid Android app that streamlines the process of uploading medical records to a SQL database through the use of a machine learning image recognition model. This project will open the door for Rwandan doctors to enhance their healthcare system by implementing a data-driven system that adapts from its mistakes and decreases the mortality rates of patients. This is one example of

how machine learning and artificial intelligence can be used to benefit society and increase the quality of life in low-income countries.

Similar to the capstone research explained above, this paper will analyze the irresponsible misuse of machine learning algorithms. The emergence of these technologies, and their effect of society, will be compared to that of the iPhone – a technology that has forever changed the way humans connect and interact. Furthermore, the paper will analyze the biases found in machine learning algorithms and discuss the interdependencies of humanity and machines. Finally, the paper will conclude by discussing the benevolent possibilities of these algorithms and how they should be used. At the end of this paper, the reader will understand the truth about artificial intelligence and gain a unique perspective about its future.

References

Copeland, M. (2019, November 06). The difference between ai, machine learning, and deep learning?: Nvidia blog. Retrieved March 02, 2021, from https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/

Cressman, D. (2009). A Brief Overview of Actor-Network Theory: Punctualization, Heterogenous Engineering & Translation. https://summit.sfu.ca/item/13593

Crypto-IT. (n.d.). Retrieved November 09, 2020, from http://www.crypto-it.net/eng/theory/modes-of-block-ciphers.html

Feng, Z., Bhat, R., Yuan, X., Freeman, D., Baslanti, T., Bihorac, A., & Li, X. (2017, November). Intelligent Perioperative System: Towards Real-time Big Data Analytics in Surgery Risk Assessment. Retrieved November 09, 2020, from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6157906/

Haughey, C. (2019, June 27). AI for Social Good: 7 Inspiring Examples. Retrieved October 16, 2020, from https://www.springboard.com/blog/ai-for-good/

Kanan, C., & Cottrell, G. (2012, January). Color-to-Grayscale: Does the Method Matter in Image Recognition? Retrieved November 09, 2020, from https://journals.plos.org/plosone/article?id=10.1371%2Fjournal.pone.0029740

Matthew Stewart, P. (2020, July 29). The Limitations of Machine Learning. Retrieved October 16, 2020, from https://towardsdatascience.com/the-limitations-of-machine-learning-a00e0c3040c6

Mills, S., Baltassis, E., Santinelli, M., Carlisi, C., Duranton, S., & Gallego, A. (2020, October 08). Six Steps to Bridge the Responsible AI Gap. Retrieved October 16, 2020, from https://www.bcg.com/en-us/publications/2020/six-steps-for-socially-responsible-artificial-intelligence

Murray, J. (2019, May 08). Racist Data? Human Bias is Infecting AI Development. Retrieved October 16, 2020, from https://towardsdatascience.com/racist-data-human-bias-is-infecting-ai-development-8110c1ec50c

Müller, V. (2020, April 30). Ethics of Artificial Intelligence and Robotics. Retrieved October 15, 2020, from https://plato.stanford.edu/entries/ethics-ai/

Pohler, L. (n.d.). A Technological Perspective on Misuse of Available AI. Retrieved October 15, 2020, from https://www.consciouscoders.io/wp-content/uploads/2018/08/A_Technological-Perspective_on_Misuse_of_Available_AI-1.pdf

Reed, M. (n.d.). The Classification of Artificial Intelligence as "Social Actors". Retrieved October 15, 2020, from https://scholarworks.gsu.edu/rs_theses/58

Torres, B. (2018, July 25). How Artificial Intelligence Can and Will Be Corrupted, and How to Prepare For It. Retrieved October 16, 2020, from https://medium.com/botchain/how-artificial-intelligence-can-and-will-be-corrupted-and-how-to-prepare-for-it-477271598a94

Smith, M.R. (1994). Technological Determinism in American Culture. *Does Technology Drive History?: The Dilemma of Technological Determinism*. (pp. 1-17). Cambridge, Massachusetts. London, England. The MIT Press.

Soohoo, S. (2020, August 25). Worldwide Spending on Artificial Intelligence Is Expected to Double in Four Years, Reaching $110 Billion in 2024, According to New IDC Spending Guide. Retrieved November 09, 2020, from https://www.idc.com/getdoc.jsp?containerId=prUS46794720

Wyber, R., Vaillancourt, S., Perry, W., Mannava, P., Folaranmi, T., & Celi, L. (2015, March 1). Big data in global health: Improving health in low- and middle-income countries. Retrieved November 09, 2020, from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4339829/