The Age of Artificial Intelligence and the Necessity of Intentional Thought: Implications of Machine Learning on Data Privacy, Bias, and Information Consumption

STS 4500 Prospectus

Computer Science Engineering

Bachelor of Science

The University of Virginia, Charlottesville

Name: Mitchell Hong

STS Advisor: Alice Fox

Projected Graduation Date: May 2024

Submission Date: 5/1/2023

How Machine Learning Technology Impacts Society

Overview:

As a prospective computer scientist, analyzing the implications of a groundbreaking technology such as machine learning is not only borne out from an urge of importance but also simply by interest. Machine Learning has recently diffused from the research sector into commercial industry, impacting all parts of society. In more detail, corporations now utilize statistical data produced by machine learning models to encaptivate consumers unlike any time before. Along with this, data collection policies and bias prevention tactics have changed greatly as machine learning progressively integrates into society. Explorations of case studies will examine the transition from the pre-machine learning era to the present day, providing a foundational understanding of the subjects. Utilizing Actor-Network Theory, a comprehensive model of these themes will be developed. Ultimately, this will enhance our understanding of how machine learning affects data privacy, bias, and information consumption within society

Positionality:

My prowess and expertise in technology are the products of the early influences in my life. Growing up in a Korean-immigrant household, I was always pushed to strive for more. Whether it was academic or personal endeavors, the urge to excel was paramount. Additionally, as I was raised in an upper-middle-class area, I was always exposed to the newest and latest technologies. My schools always had access to computers and tablets, enabling students and staff alike to deepen their understanding of modern technologies. I immersed myself in Computer Science classes at every opportunity given and heavily researched growing technological fields in my personal time. The coupling between my household’s Korean academically oriented culture and the ease of accessibility to technology served as a crucible for technology literacy.

During the Summer of my second year in college, while my colleagues were working tech internships at respectable companies, I was stuck at home due to an Air Force Reserve Officer Training Corps obligation. Unwilling to sit idle while my colleagues enriched their knowledge in technology, I began a personal project. The project involved emerging technologies such as machine learning and information scraping off of the internet. I watched countless hours of informative videos and partook in Google’s machine learning course. The weeks spent on this project resulted in a machine learning model capable of predicting the margin a football team might win by, given the week's matchup. I have since then researched other fields in artificial intelligence and plan on taking a machine learning course during my seventh semester in undergrad. Evidently, my aptitude and unique experiences have cultivated my trajectory into artificial intelligence research.

Problematization:

Machine learning, the disruptive technology of the modern era, is being ushered into every facet of society. It is undeniably a powerful tool at the forefront of progress; however, more care and thought are needed before fully integrating it into every sector of society. If not done right, machine learning technology has the potential to infringe on data privacy rights,

further biases unintentionally, and create addictive applications that exploit users. Before society becomes riddled with these problems, it is vastly important to analyze how machine learning technology truly impacts its stakeholders. Machine learning technology is mostly unrestricted in how it is utilized, but should this be the case?

Guiding Question

Underneath the glamour of its capabilities, what are the major ethical implications of machine learning technology on data privacy, bias, and information consumption?

Projected Outcomes:

As Machine Learning technology is integrated more and more into society, it is imperative that clever, intentional thought is laid into the framework upon which the technology is being implemented in. The ethics of institutions implementing M.L. technology should be majorly reformed to better reflect the immense implications of the technology. These reforms should prioritize the well-being of all stakeholders, ensuring that decisions are not driven solely by profit or efficiency, but also by minimizing any potential negative impacts on society and individuals. Additionally, by further researching this disruptive technology, the general public will be empowered, by means of knowledge, and able to protect themselves by avoiding some of the negative aspects of machine learning.

Technical Project Description:

As aforementioned, I constructed a machine learning sports betting model during the summer of my sophomore year in college. During this endeavor, I encountered many unforeseen obstacles that widened my expanse of knowledge. From Google's Machine Learning course, I expected to simply plug a dataset into functions that the TensorFlow library provides and produce a fully functioning model. However, I faced challenges in data collection, cleaning, and application.

My initial challenge was finding usable datasets on the NFL 2021 season. Most datasets were unavailable in CSV form, and some websites required paid subscriptions, which I wanted to avoid. Fortunately, I found team stats such as yards per play, defensive yards per play, etc. in CSV form and stored them for later use. However, team records and game-by-game scores, essential pieces of information for my model, were only displayed on websites in HTML cells. Consequently, I began researching Python libraries that would enable me to scrape HTML off of websites into Pandas data frames. The Beautiful Soup library proved invaluable for this task. When attempting to implement these functions into my program, I once again encountered an unforeseen obstacle: the webpage needed to be rendered before I could scrape and parse the HTML. After more research, I discovered that Selenium was the solution, and with Beautiful Soup, I finally collected the necessary data to train my model.

Once I had collected all the required data, I realized that it was malformed. Critical pieces of data had unidentifiable characters or were of the wrong data type. To remedy this, I wrote a program that removed the erroneous characters and converted the datatypes to a uniform format that the model could use.

Using TensorFlow functions, I finally constructed my sports betting model. However, due to the small dataset I had acquired, the model's efficiency was limited. Despite this, completing the model provided me with valuable insights into designing and constructing a robust model.

Preliminary Literature Review & Findings:

Research on machine learning technology boasts about its sheer power and impact in various sectors of society. However, it is suggested that power and impact alone should not be the only factors considered when introducing new technologies. Ethical and moral aspects of the technology must be analyzed in tandem with its physical capabilities to ensure positive outcomes for the general public.

Machine learning, by nature, is only effective if the training data provided is clean and accurate. Unfortunately, in certain sectors, such data may be challenging to obtain. Both Mhasawade (2021) and Samorani (2021) illustrate this issue in the healthcare sector. Machine learning with clean data, to only name a few use cases, can assist in predicting disease outbreaks, identifying high-risk populations, and optimizing treatment plans. On the contrary, health disparities are reflective of social oppression and its influence on the health of individuals that identify with such marginalized communities (Mhasawade 2021, p.663). When the data is dirty and misrepresentative, machine learning has the potential of perpetuating existing biases and furthering discrimination, lowering the overall level of community health.

Shah (2020) and Lang (2018) discuss the use of machine learning technology in the consumer space, where highly funded corporations can use regression tools and enormous amounts of data to tailor technologies that drive addiction and expose customers to irresistible advertisements, driving corporate profits. Overall, the future of advertising is going to be even better than before as artificial intelligence and machine learning will bring more control of advertising to companies (Shah 2020, p.1).

It is evident that M.L. technology depends on reliable data. Further, corporations and institutions have shifted efforts to mass data collection. Liu (2021) and Strobel (2022) detail various fatal flaws in mass data collection efforts, noting key issues such as data protection, fairness, and transparency. With the large amounts of data collected, data breaches pose a significant threat, furthering the necessity to handle data with the utmost regard.

Given the difficulties in data availability and the relative newness of the field, researchers in machine learning technology face many challenges. Building on their research, this paper aims to create a cohesive data bank that is able to inform the general public about the many implications of machine learning technology on data privacy, bias, and information consumption.

STS Project Proposal:

In essence, STS is the study of the reciprocal relationship between modern science, technology, culture, and society as a whole, examining how they mutually influence and shape one another.. Too often, science and technology are treated as tools that society uses as needed. In reality, science and technology are dynamic entities that greatly impact the culture and flow of

society. Furthermore, the reciprocal relationship between society and technology means that cultural values, social norms, and human needs also play a significant role in shaping the development and application of technology. As machine learning is a technology significantly altering the way society operates, it is paramount that an STS lens is used to analyze its many aspects. Overall, a more encompassing, full-scope point of view must be employed to better understand how science and technology interconnect with society, and how society, in turn, influences the evolution of science and technology.

Mainly, the ethics associated with machine learning will be discussed in great detail. The articles written by Samorani and Blount (2020), Mhasawade, Zhao, and Chunara (2021), Kaplan (2016), and Vesnic-Alujevic, Nascimento, and Polvora (2020) will be invaluable in this endeavor. While other researchers delve into various important topics related to machine learning, these authors particularly emphasize the ethical impacts of machine learning concerning data privacy, bias, and information consumption. This focus aligns with the STS perspective on societal and ethical implications. Samorani and Blount (2020) highlight the inequalities that are potentially perpetuated when implementing machine learning in medical appointment scheduling. Mhasawade, Zhao, and Chunara (2021) provide insights into algorithmic fairness in machine learning and the challenges associated with it. Kaplan (2016) discusses ethical and societal issues related to artificial intelligence. Vesnic-Alujevic, Nascimento, and Polvora (2020) describe the societal and ethical impacts of artificial intelligence by examining European policy frameworks. These authors' analyses shed light on the complex challenges introduced by data-driven systems, including issues of fairness and ethics.

In order to get a full-scope understanding of how machine learning technology shapes society, Actor-Network Theory (ANT) will be employed. This framework in particular will be vital as it assists in the understanding of how human and non-human entities interact. In addition, ANT emphasizes the importance of the relationships between actors rather than the characteristics of the individual actors themselves. In the case of machine learning, by understanding the complex relationships between actors such as algorithms, data, humans, and institutions, a comprehensive understanding of the socio-technical implications of machine learning can be gained. According to ANT, these relationships are constructed through micro negotiations, further enforcing that technology shapes the social context in which it is developed and used. This approach aligns with the definition of STS as it emphasizes the need to understand how science and technology are embedded in social processes and institutions.

For my research, as primary data collection means, I intend on using both literature analysis and case studies. The literature analysis will aid in my ability to identify not only points that the authors reach a consensus on but also what topics are more hypothesized. In general, I will gain more comprehensive knowledge of the ethics of machine learning while also discovering what is missing in the field. Through the use of case studies, I will be able to examine real-world examples that highlight the practical implications and challenges associated with machine learning, providing valuable context and insight into the ethical landscape. Ultimately, a sophisticated mapping of the ethical and societal issues associated with machine learning will be constructed.

Barriers & Boons

Having researched the field of machine learning, I acknowledge that multiple potential limitations could impact my work. My first limitation stems from the extent of my research. Although I have dedicated a considerable amount of time to this topic, a more definitive, robust understanding could be obtained if I had more time to dedicate to it. Additionally, the availability of data and case studies are limited as machine learning is a relatively young field. As a result, it is difficult to find relevant and reliable information, which may limit the scope of my research.

Another potential limitation is my inexperience in machine learning. Although I have gained some experience in this field over the past year or two, there is still much to learn. This lack of experience may hinder my ability to employ more advanced techniques, skewing the accuracy and effectiveness of my research.

To address these potential limitations, I plan to develop my skills and knowledge through various means. This includes but is not limited to, reading more literature on the subject, taking classes, and seeking out collaborations with experts in the field. Additionally, I intend on engaging more within the community. By doing so, I will gain insight into current research and potential avenues for exploration.

Despite the challenges and limitations, I remain confident that with dedication and a commitment to learning, I can overcome these potential blind spots and produce valuable contributions to the field of machine learning. Innovative approaches and a better understanding of the field and its potential application will be gained with a continuous learning mentality.

References

Lang, T. (2018). Understanding Consumer Behavior with Recurrent Neural Networks. Journal of Consumer Research, 45(5), 956-976. doi: 10.1093/jcr/ucy056

Liu, B., Ding, M., Shaham, S., Rahayu, W., Farokhi, F., & Lin, Z. (2021). When machine learning meets privacy. *ACM Computing Surveys*, *54*(2), 1–36. https://doi.org/10.1145/3436755

Strobel, M., & Shokri, R. (2022). Data Privacy and trustworthy machine learning. *IEEE Security & Privacy*, *20*(5), 44–49. https://doi.org/10.1109/msec.2022.3178187

Samorani, M., & Blount, L. G. (2020). Machine learning and medical appointment scheduling: Creating and perpetuating inequalities in access to health care. *American Journal of Public Health*, *110*(4), 440–441. https://doi.org/10.2105/ajph.2020.305570

Mhasawade, V., Zhao, Y., & Chunara, R. (2021). Machine Learning and Algorithmic Fairness in public and Population Health. *Nature Machine Intelligence*, *3*(8), 659–666. https://doi.org/10.1038/s42256-021-00373-4

Shah, N., Engineer, S., Bhagat, N., Chauhan, H., & Shah, M. (2020). Research trends on the usage of machine learning and artificial intelligence in advertising. *Augmented Human Research*, *5*(1). https://doi.org/10.1007/s41133-020-00038-8

Kaplan, J. (2016). Artificial Intelligence. https://doi.org/10.1093/wentk/9780190602383.001.0001

Vesnic-Alujevic, L., Nascimento, S., & Pólvora, A. (2020). Societal and ethical impacts of Artificial Intelligence: Critical Notes on european policy frameworks. *Telecommunications Policy*, *44*(6), 101961. https://doi.org/10.1016/j.telpol.2020.101961