Thesis Project Portfolio

Machine Learning: Reasons Behind Average Lifespan

(Technical Report)

Seeing Eye to AI: A Comparison on the ethics of AI services and African American Enslavement as Tools for Humanity

(STS Research Paper)

An Undergraduate Thesis

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Rishal Jagtap

Fall, 2022

Department of Computer Science

Table of Contents

Sociotechnial Synthesis

My technical project and STS research paper are closely intertwined with the driving factor being the implementation of AI in today's world and its use in the future. My technical paper takes the more expected approach to how engineers view AI, using machine learning algorithms to predict life expectancy in various states across the US. In contrast, I wanted my STS research paper to focus on the ethics related to artificially intelligent beings, how we treat them, what we can expect for their future, and how they relate to humans unlike normal tools. Part of what made machine learning so interesting was that for the first time, not only was I able to learn from the project, but the project itself was able to *learn* from me. In this sense, my technical project reminded me of how humans act, it felt like we were learning together. Knowing nothing at first, but with the right data, being able to predict outcomes is much like a toddler learning to swim. If you taught the algorithm incorrectly, you'd get bad results, just as now a toddler could drown if not taught to paddle.

My technical report is an extension of a prior project in my Machine Learning course where our team used Machine Learning-based models to determine the average Life expectancy of Virginians. The primary objective of our project was to create a tool that made it easy to visualize the various hurdles citizens in different parts of Virginia face that lead to a lower life-span. In my technical report, I proposed using the project on a wider scale, shifting outside of Virginia borders to gather data from all across the United States. The project had two major phases and was based on the Virginia Census on life expectancy. The first phase would use the K-means clustering model and create three distinct centroids. All the data points would then be plotted around the three centroids and would be checked to see which feature caused them to appear so close to their associated centroid. In the next step, we used T-distributed Stochastic

Neighbor Embedding (T-SNE) to determine how many of the 23 possible features actually had an impact on the data. The user would have the opportunity to set their own confidence intervals, say 95%, to remove any feature that does not aid in reaching that threshold.

From the experience of teaching my algorithm stemmed the basis of my STS research paper, which started as a slew of questions surrounding the sociotechnical impact of artificially intelligent beings on a predominantly human-led society. If AI really do have the ability to become like humans, would it really be correct for us to exploit them as a resource? More importantly, was it even realistic to view AI in the same light as humans, are they even close enough in nature? To answer these questions, I referred to "Engineering as a Social Experiment" by Martin and Schinzinger, a framework that uses past human experiences as a form of "experiment" in order to better understand how we can change future implementations of similar design. While we are still far from actually employing AI, we have seen the exploitation of beings similar to us in the form of slavery, which was built on the fundamentals that some humans were lesser than others. Through my STS research using "Engineering as a Social Experiment", I explore how similar AI are to humans in terms of consciousness and sentience, then transition to debating the type of world we could expect if AI are treated like slaves.

Through my STS research paper, I found that the ethics behind AI are much more complex than any other tool we've had to deal with as a society. The stark contrast between my projects exemplifies the duality in engineers with wanting to produce a solution right away versus taking an approach that was well thought out. While writing my technical report, I looked to numbers as a means of determining right from wrong. When I eventually moved onto my STS research paper, I noticed that my subjective questions on AI behavior weren't being asked until

after the technical portion where ethics did not factor into success. The major takeaway from my research was that it is much more beneficial to ask such questions before, during, and after the project has come to completion. Ultimately this process would allow engineers to take the initiative and formulate the consequences of their creations.