Identyti: Digital Storage of Personal Identification Documents
(Technical Report)

An Analysis of Kantian Ethics on the Usage of Artificial Intelligence in Legal Sentencing
(STS Research Paper)

An Undergraduate Thesis Portfolio

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

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By

Sridhiraj Jayakumar

May 1, 2020

Table of Contents

Socio-technical Synthesis: Digital Storage and Legal Sentencing

My technical work and STS research are loosely connected by the idea of using machine learning systems in critical environments, and its implications. Machine learning is a broad topic, but both my research and work focus on systems that determine the classification of certain objects. My technical work revolves around using this technology for a small role in a whole system, while my research highlights the moral consequences of using this technology in legal sentencing. So, while machine learning may be a small part in my technical work, it is still an essential part of my project and relates to my STS research.

My technical work revolves around a proof of concept regarding a solution for the digital storage of personal identification documents. In our prototype, we use machine learning to classify documents that a user uploads. Specifically, a document can be classified as six different types in our project: Identification, Insurance, Medical, Financial, Work/Academic, and Legal. To assist the user in figuring out the category in which the uploaded document falls under, we use the machine learning model that I built, to classify it. To build this model, I used around one hundred documents per category to train the model based on text information and distinguish between the language used in these forms. However, there may be documents in which these categories overlap, which is why we allow the user to change the category, if they disagree with the result. This leads to my STS research, which explores the use of machine learning in critical situations, mainly the moral consequences of using technology that might be biased.

My STS research revolves around the use of machine learning in legal court proceedings. Particularly, I explore the use of machine learning in courts to determine recidivism risk in individuals. Recidivism risk is the likelihood that an individual will recommit a crime. In my research, I highlight the moral consequences, under Immanuel Kant's theory, of using this

technology to assist in legal sentencing. I find that, with the basis of Kant's theory, the usage of risk assessment technology with the explicit use to assist in sentencing is immoral mainly because it violates Kant's categorical imperative and it relinquishes the judging party's moral autonomy. The goal of my research is to explore the dangers of using machine learning and to assist further researchers in developing safe measures to use this technology in critical settings.

Performing my technical work in conjunction with my research has helped me implement machine learning in my project in a safe and secure way. Through my research, I came to understand the importance of understanding the technology before applying it and to implement safe guards that allow the user to refute against the decision of the system. In the court cases I researched, I found that it was important for the defendant to understand the system's assessment of their risk, and to be able to refute any accusations made by the system. In a similar fashion for my work, I wanted the user to be able to change the category in which a document was placed if they disagreed with the machine's classification. By conducting my research and my work using similar topics, it allowed me to thoroughly explore the different implications of using machine learning in different settings, which ultimately helped refine each other.