

Thesis Project Portfolio

Power of Difference Assessment System

(Technical Report)

Ambiguous yet Polarizing: An Ideographic Approach to Analyzing the Variation in Usages and Conflicts of Values That Have Arisen From the Term “Algorithmic Fairness”

(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

Nuzaba Nuzhat

Spring, 2020

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

Power of Difference Assessment System

Ambiguous yet Polarizing: An ideographic Approach to Analyzing the Variation in Usages and Conflicts of Values That have Arisen from the Term “Algorithmic Fairness”

Prospectus

Sociotechnical Synthesis

With the passage of time, advancements in technology, and changes in society, we are interacting much more often with people whose life experiences, ideas, and beliefs are significantly different from ours, whether that be because of their gender, race, culture, sexual orientation, socioeconomic status, physical abilities, or religion. The goal of my technical project was to rebuild an assessment system that helps people deal better with the challenges these dissimilarities engender. My STS project, on the other hand, is centered around technology (mostly machine learning algorithms) that perform almost the opposite function: they make high-risk decisions that often disproportionately affect one section of the population over the other, perpetuating unfair social divides. The objective was to understand how, in the struggle against such discriminatory machine behaviors, “algorithmic fairness” as a term wields its rhetorical power and shapes non-technical human actors’ beliefs and courses of action.

For the technical project, my team and I developed a web application called “The Power of Difference Assessment” that lets users register with their email, fill up demographic information, then respond to 70 difficult and/or controversial statements across various sociocultural locations. For example, in the sociocultural location of religion, they are asked how much they agree to the statement “I love different religions but don’t know how to reconcile this with what I feel is the truth of my faith”. After the user has responded to all the statements, they are emailed the results, which explains through a series of graphs where they excel and where they can improve with regards to their role in ensuring a just, equitable and inclusive society. The system also comes with a range of administrative functionalities that allows our client to manipulate the data, enroll new staff members, and visualize the information collected. Together,

these functionalities have made the system drastically more secure, functional, and scalable, and have paved the way for the assessment to amplify its societal impact beyond its founding years.

For my STS project I explored another scenario where social concerns related to justice, equality, diversity, intersect with software systems – the movement for “algorithmic fairness”. Although computer scientists and legal scholars are struggling to come up with precise definitions of fair algorithms, it is not stopping other stakeholders from using the term in ways that influence the entire sociotechnical system. In this project, I employed the rhetorical tool of ideographs to look at how government agencies, and the general public use the term “algorithmic fairness” to mean different things in different contexts and how the term’s usage both conflicts with and strengthens other goals and values individuals and groups hold dear. The analysis should help stakeholders become more aware of the implications of using such an ambiguous yet powerful term, and when engaging in conversations, encourage them to be more specific about what they mean by “algorithmic fairness”, and why and how much they think it matters.

The projects together depict both sides to the same coin: the technical project showcases how machines can influence the way we address humanity’s struggle with the inter-human variability of physical traits, backgrounds, and power, whereas the STS project emphasizes how human beings, even those without technical knowledge, can guide the approach to machine weaknesses regarding the very same issues. Working on them concurrently has given me an appreciation for the bidirectional nature of human software interaction that is prevalent in every sociotechnical system. In addition, it has highlighted to me, as I hope it will to other software engineers, how quickly the moral debates surrounding fairness, equality and justice are permeating into the software industry, and why we, as the experts, need to be better versed about

all relevant arguments, and discuss as a moral community how we want it to impact what we code.