

Thesis Project Portfolio

Cloud System Reliability: Expanding the Infrastructure to Support a Commercial Login Service

(Technical Report)

Large Language Models: A Case Study on Corporate Values and Transformative Technologies

(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

Sidhardh Burre

Spring, 2024

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

Cloud System Reliability: Expanding the Infrastructure to Support a Commercial Login Service

Large Language Models: A Case Study on Corporate Values and Transformative Technologies

Prospectus

Sociotechnical Synthesis

Algorithms make decisions that directly harm our happiness (Zhang et al., 2022), freedoms (Benjamin, 2020), and livelihoods (Eubanks, 2018). Their design shapes the way that we interact with one another and the world around us. Combined with their far-reaching nature and rapid propagation, algorithms are already becoming a keystone of our society. Therefore, it is paramount to find a balance between developing algorithms with innovative technical capabilities and societal benefits. To find this balance, software development, the process responsible for the design, development, and maintenance, must be closely aligned with core human values. The driving mechanism of this study was gaining understanding of the inherent values of software development and the origin of those values so that future technologies are functionally efficient and uphold ethical standards to benefit society at large. The technical portion of this study explores software development hands-on through a commercial banking internship. My STS thesis analyzes the values of software development through Value Sensitive Design to address ethical concerns in LLMs.

My technical research centered on enhancing the reliability of cloud-based systems for commercial banking. In the first phase of my internship, I employed tools such as Splunk and New Relic, to expand existing monitoring capabilities. This resulted in the halving of incident resolution times. In the second phase, I used the Spring WebFlux framework, to develop a prototype that significantly improved system performance, yielding a more than sevenfold increase in concurrent user support. These improvements underscore the vital role of sophisticated software development practices in ensuring the operational integrity of critical financial services. Further, the internship provided insight into the values that software development teams adopt in the face of stakeholder pressure to deliver products on a rapid timeline as well as how those values can be reversed as needed to rectify prior mistakes.

I chose to examine the ethical dimensions of Large Language Models (LLMs), particularly their propensity to perpetuate societal biases. To determine the underlying values of design decisions that resulted in ethical failures, Value Sensitive Design (VSD) was utilized to break down LLM design processes. Through a detailed analysis of the data selection processes and model training methodologies, as well as their resultant ethical failings, the research highlighted how LLMs could unintentionally reinforce existing social problems. By comparing the development landscape of LLMs with Open-Source Technologies (OSTs) through Value Sensitive Design, the study advocated for the integration of ethical considerations early in the software development cycle of LLMs. By integrating ethical considerations into the development process, it is possible to foster technologies that genuinely reflect diverse human values.

The synthesis of findings from both the technical project and STS research illustrates a comprehensive approach to software development that prioritizes both technical robustness and ethical responsibility. My work contributes to the broader dialogue on how to develop technology that is both technically advanced and socially responsible. These contributions are crucial for guiding future developers, ensuring that they construct tools for positive societal impact rather than merely commercial gains. Future research should explore the implementation of dynamic ethical frameworks in the early stages of software development processes. This includes developing more sophisticated ethical audits and adding continuous monitoring systems early in the development process. Additionally, the development community should focus on creating more transparent systems that allow engineers and users to understand and control the data used by these technologies.

I am grateful for the guidance and insights provided by my advisor Professor Caitlin D. Wylie, whose expertise was instrumental in shaping the scope and depth of my research. My colleagues during

my summer internship provided invaluable feedback that enriched my technical work. Finally, I appreciate the support from my family and friends, whose encouragement was crucial throughout my academic journey.