

Sociotechnical Synthesis

STS 4600
Spring 2024

Benny Bigler-Wang
B.S. Computer Science

Artificial intelligence (AI) is a rapidly growing field, and the integration of Behavioral Detection Systems across various domains has sparked both excitement and concern. There are several different multifaceted implications for AI-driven Behavioral Detection Systems, that when further inspected may outweigh the understanding and utility that these systems provide. Behavioral Detection Systems analyze certain human behaviors in various forms (image, speech, text, etc.) and try to recognize abnormal behavior or make predictions based off of an input.

This thesis is composed of three distinct components: a technical project evaluating the CS 3240 curriculum at the University of Virginia, a Science, Technology, and Society (STS) paper that critically examines the social implications and ethical dimensions of Behavioral Detection Systems, and a Prospectus that focuses on the integration of Behavioral Detection Systems in Cybersecurity in order to detect malware more effectively.

The technical project delves into the current computer science learning experience within UVA's CS 3240 Advanced Software Development course, and addresses the challenges of preparing students for software development careers in the real world. Comparing what was taught in UVA's CS3240 Advanced Software Development course to relevant technology now used in the industry, highlights both the strengths and weaknesses of the current curriculum and hints at potential adjustments that will better equip students to face the challenges in the software development field. By the end of the course, students have learned how to use many different third-party technologies and frameworks but have not learned ones that may be more relevant to today's standards. An ongoing assessment to reevaluate the efficacy of the revised CS 3240 curriculum is imperative to ensure that it is preparing students to use the ever evolving technology being introduced into the industry.

The STS paper defines a social integration spectrum that AI-driven Behavioral Detection Systems can fall on and explores the degree to which potential socio-ethical consequences such as discrimination and privacy loss can occur based on where they fall on the spectrum. Through the analysis of several case studies and literature reviews, the relationship between how broadly connected a certain application of a Behavioral Detection System is with social-interactive data or society is established and shows the ethical considerations, privacy concerns, and potential discriminatory practices inherent in the deployment of these systems across different domains. This relationship that is present in the deployment of Behavioral Detection Systems can be attributed to an oversight in the current capabilities of these types of AI models and calls for more comprehensive regulation regarding fairness and the protection of data privacy depending on one's place on the social integration spectrum.

Both main components, the technical paper and STS paper, are divorced from one another but they do share a central theme of evaluating a balance between technological or curriculum advancement and the implications behind them. Without hindering the growth or effectiveness of new systems, it is imperative to interrogate the reasoning and ethical dimensions

of its applications. For the evaluation of Behavioral Detection Systems as they are utilized for various social-integrated applications, from cybersecurity to government surveillance, it is important to investigate the ethical risks and societal harm that can be brought about in the real world. Rollout of actual Behavioral Detection Systems needs thoughtful deliberation, the creation of a robust regulation framework worked on by domain experts, and interdisciplinary collaboration between those experts and AI developers to ensure safety and security. The STS paper and Prospectus paper are more tightly connected as the Prospectus exhibits the potential benefits of using Behavioral Detection Systems while mitigating the ethical risks and societal harms and specifically investigates its application in the cybersecurity field as a better alternative to the outdated signature-based detection methods.

In conclusion, this thesis mainly contributes to the ongoing discussion of AI's place in the world and the ethics and governance that come along with it, offering insights for future development, deployment, and preparedness to the societal consequences of Behavioral Detection Systems. By acknowledging the potential socio-ethical risks associated with these systems, policymakers and AI developers can do a better job implementing responsible AI technologies that are able to provide utility while decreasing the odds for unintended consequences.