Thesis Project Portfolio

Spectre and Meltdown: A Look into Mitigation and Detection

(Technical Report)

Analysis of AI Governance

(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

Caroline Hickey

Spring, 2021

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

For my STS research, I chose to explore AI governance, the approach taken to regulate and govern new AI technologies. I initially began by looking into bias in AI, and solutions to mitigate this bias. I discovered that a very effective strategy would be to create standards for AI technologies and regulate their use. This led me to research different strategies that countries, organizations, and companies have taken to govern AI systems. My ultimate goal was to find the most cohesive and comprehensive AI governance strategy through examination of various governance frameworks. I chose three entities, the US government, a non-partisan organization, and the Singaporean government. After analysis and research, I discovered that there wasn't a framework that provided sufficient governance of AI. I suggested that through a combination of strategies, especially pulling from Singapore's governance, a complete AI governance framework was feasible.

For my technical project, I explored the Spectre and Meltdown vulnerabilities. These two hardware vulnerabilities exploit speculative execution in order to leak sensitive information that was formerly thought to be impossible to access. They affect nearly all processors, meaning every computing system, including every computer, smart phone, and cloud application, is affected. Since Spectre and Meltdown are hardware vulnerabilities, they are not easily fixed by software patches and require extensive solutions. I began by researching how these attacks were executed, and what current mitigations companies have already provided. I discovered that the solutions already published do not fix the vulnerabilities completely, causing attackers to still be able to execute the attacks. Full solutions to the vulnerabilities cause performance to slow by as much as 30%, making them difficult to justify.

Both research topics allowed me to gain insight into important topics in the field of Computer Science. Through my STS research, I was able to look into the intricacies of governing a new technology, especially one as powerful as AI. My results were helpful to show what needs to be done, but also uncovered how unregulated AI is, considering how extensively artificial intelligence is already being used. It also showed that the motivations of those who have begun to govern AI motivated by power and money, rather than by the dangers and implications of an unregulated powerful system. For future work, I would recommend researching the feasibility of implementing an "ideal" governance structure in various countries. This might reveal if, and when, a governance structure will be put into place. My technical project made me aware of the dangers of prioritizing performance over security. I discovered feasible solutions that might help prevent against attacks in the short-term, and discovered the impact of the vulnerabilities of the future of hardware production. In the future, I recommend other researchers find the solution with the best security versus performance trade-off. This would allow us to choose the most feasible and comprehensive solution to protect against these attacks.