Thesis Project Portfolio

Architecting Computer Vision Workloads on the Cloud

(Technical Report)

The Emergent Technopolitics of Big Tech Domination

(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

Arvind Anand

Spring, 2022

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

Commented [SF1]: Insert your Sociotechnical Synthesis below this title. Don't change the title.

Finally, if you have changed any of the fonts, make sure they are at least consistent across all titles and text.

A significant portion of the waste created in operating rooms comes from the wasteful use of single-use sterilized tools, which cannot be used again once opened. The high-intensity nature of the operating room often forces surgeons to react and adapt uniquely to each patient and each operation, requiring all possible single-use tools to be open and present at all times. All of the opened single-use tools are thrown away regardless of whether they have been used or not – this generates a huge amount of waste that could otherwise be avoided if the correct tools can be predicted ahead of time. PeriopGreen is a pre-seed startup based out of UVA that leverages computer vision to track and eliminate wasted single-use tools in the OR. The following technical white paper breaks down how my senior capstone team and I architected PeriopGreen's system for running computer vision workloads at scale, and the following STS thesis is a broad commentary on large scale centralized compute infrastructure, the problems that arise as a result, and the proposed solutions to fix them along with their drawbacks.

The technical white paper breaks down the problem statement, technical constraints, design principles, MVP, and final product along with its deliverables: we've dubbed this deliverable the Model Engine. While the name suggests that the Model Engine solves a fundamental deep learning problem, it solves the problem of scale. It's relatively trivial to run footage through a neural network if it's for recreational use, but running machine learning workloads on terabytes of data in the cloud requires deliberate engineering for high resiliency and availability - the system is expected to handle high volumes of video footage from multiple hospital ORs without any downtime to the core business. As a result, the Model Engine runs on AWS and leverages its technologies and other technologies like Docker to operate at a large scale. The current product is complete and ready to serve customers at PeriopGreen.

The STS thesis analyses the technopolitics behind centralized compute platforms like AWS. This is done by analyzing case studies where AWS used its position of power to influence other actors and the open internet itself. In addition, the paper dives deeper into the mechanisms of power and the main driving forces of the tech industry itself. Notions of data inequity reveal troubling patterns that emerge from such data consolidation. It turns out that data asymmetries give rise to economics that ultimately disadvantages average users. The paper compares decentralized technologies with centralized computing and discusses the pros and cons of future decentralized solutions.

I want to end by thanking my capstone partner and co-lead engineer at PeriopGreen, Dylan Fernandes – without him, I wouldn't have had the fantastic opportunity to work with this company. I'd also like to thank Dr. Meyer and the rest of the PeriopGreen team for supporting this incredible project.