Thesis Portfolio

A Robust Deep Learning Approach to Pedestrian Detection via Thermal Imaging
(Technical Paper)

When Cost-Benefits Go Wrong:
Differences in Decision-making between Physical Product and Software Defects
(STS 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

Navya Annapareddy
Spring, 2021

Department of Engineering Systems and Environment

Table of Contents

Sociotechnical Thesis

The total amount of data in the world is projected to grow at an exponential rate in the past few decades. Consequently, much of this data is typically generated, collected, manipulated, or otherwise utilized via software like machine learning models. Despite this, no machine learning model or piece of software is perfect with regard to transparency, degrees of uncertainty, false positives or negatives, and/or bias. How a defect in a given product is handled, regardless of whether it is a physical product or a virtual one, has great implications on how much impact the defect has post-production. While physical products currently have extensive regulations for widely-accepted and standardized metrics like deviance and error tolerance, non-physical products like software and machine learning models do not have the same level of regulation from developers or third-party regulators. Notably, the decision making processes and developmental cycle of physical and non-physical products greatly differ. This also manifests into physical and non-physical software defects being handled differently. These differences in defect responses between physical products and software are explored in the STS paper.

The technical project involves building a deep learning model to identify pedestrians given inputs of thermal imaging. This is different from traditional pedestrian detection models given that the only data inputted will be grayscale image data of pedestrians rather than full color RGB data. Thermal imaging has advantages over color data in that it is not affected to the same extent by weather conditions and features of the pedestrians themselves such as skin color. Metrics that the model will be evaluated on include how accurate predictions are as well as false positives.

This STS research paper and Capstone project thus share their focus on addressing defects and biases in software and machine learning. Both of them involve examining the subjective nature of reporting process information and what constitutes an appropriate margin of error. While the STS research paper provides discourse on how decision making around defects in physical and non-physical products are handled differently, the technical project involved the creation and testing of a deep machine learning model to detect pedestrians and cyclists in real time using thermal imaging. While the deliverables for the technical and STS components were different, both were meaningful. The Capstone project taught me how to build an end to end solution encompassing both existing and novel machine learning and data augmentation techniques. The STS research paper, on the other hand, taught me how to employ and integrate STS frameworks and ethical perspectives to explore a given topic.