Thesis Project Portfolio

Adaptive Synaptogenesis Neural Networks: Creating Factorial Encodings Using Inhibitory Neurons

(Technical Report)

A Deontological Analysis of the Amazon AI Recruitment Tool

(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

Cooper Scher

Spring, 2023

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

My STS research and technical work are related through an exploration of the evolving methods of performing machine learning (ML) and artificial intelligence (AI) tasks. AI is the field of computationally creating human-level intelligence while ML is a subfield of AI that uses data to learn patterns and relationships between items of interest. Each work investigates ML from a different perspective. The STS research focuses on the ethical implications and issues of the implementation of AI as a replacement for human functions such as human resources (HR) while the technical portion looks at improving an existing ML algorithm, the adaptive synaptogenesis neural network model. Even though each project approaches the topic from a different vantage point, the notion of the quickly changing nature of AI algorithms is explored in both projects.

In my technical report, I modify the adaptive synaptogenesis neural network model to achieve factorial encoding. Adaptive synaptogenesis is an unsupervised algorithm that attempts to break down a dataset into categories in an efficient manner. Factorial encoding is the property that each input is encoded in such a way that all of the output codes are conditionally independent with respect to the category of input, producing space-efficient and simple representations or encodings of an input dataset. The neural network was modified to include inhibitory feedback, which was simulated on datasets using MATLAB. The simulation results were compared to previous models to demonstrate the new factorial encoding property.

In my STS research paper, I also look at the use of ML algorithms, but in this case it is the ethics of a hiring AI previously used by Amazon to automate the task of finding good job candidates by their resume. I explore how Amazon's AI was profoundly sexist, and I discuss how solutions that just focus on "non-biased" algorithms are insufficient. Instead, I claim Daniel Dennett's standard of AI that can rise to a moral agent as a standard of sufficiency to replace the complex moral decision-making capability of human HR. To support this, I look at the circumstances of the Amazon case study and explore recent literature on how AI and ML can and should respond to issues of bias and interpretability. The goal is to bring attention to the nuanced nature of algorithmic bias, which is often unintentional and not easily solved directly.

Working on these projects together, I have learned not only about why there is a need to continue developing ML algorithms that are more equipped to handle issues like biased input data, but also I have learned directly about the process of research and developing those kinds of improvements. In particular, the STS research helped contextualize the value of my technical project. While I did gain much understanding and experience from working on modifying the adaptive synaptogenesis neural network model, I was approaching it from a sense of curiosity instead of understanding its societal impact before undertaking the STS research project. Similarly, I was able to develop a better intuition about what kind of research efforts are needed to achieve the goals in my STS claim after doing my technical project. Thus, I believe that working on both projects concurrently was beneficial with a synergistic increase in understanding about the impact and context of each project.