Personal Statement – Ahad Jawaid

Introduction: *python inputBankingInfo.py Running... Finished*. And with that, it was done. The perpetual twelve hours of weekly data entry into Excel spreadsheets had forever ended. As an assistant real estate manager, I was tasked with repetitive work, like data entry. Every day, I slogged through inputting tenants' accounts payable into an Excel sheet, thinking there had to be a better way. That's when I discovered *Automate the Boring Stuff with Python*, my first programming book, and automated the data entry process—marking the beginning of my journey as a programmer. Four short years later, I have written code for Amazon, one of the largest companies on the planet; taught hundreds of students about machine learning (ML) and artificial intelligence (AI); participated in hackathons across the country; researched synthetic human-like speech generation; now I'm working towards creating robots that can execute tasks defined in natural language, while also exploring ways to communicate with computers directly through our brains. My journey into research wasn't straightforward. I didn't even know what it meant to do research until the summer before my third year of university, and I didn't stick to a single lab throughout my undergraduate years. Although this may have resulted in fewer publications, it helped me discover my interests, become more independent as a scientist, and strengthened my conviction that I want to pursue a career in research.

Intellectual merit

System Security Lab: During the summer before my junior year, I collaborated with Dr. Kangkook Jee in his System Security Lab at the University of Texas at Dallas (UTD). My role was to extend the Graph Neural Network (GNN) Explainer algorithm to work with heterogeneous graphs, enabling us to generate explanations about why the GNN classifier detected a program as malicious based on the corresponding heterogeneous resource interaction graph. After programming the algorithm, I validated its accuracy by applying it to a classifier trained on the MUTAG (Mutagenic Chemicals) dataset. I confirmed the reliability of the algorithm's explanations by comparing the molecules it identified as most important for classification as a mutagenic compound to established scientific ground truths, finding a strong match with existing research. This work led to my **open-source contribution of the algorithm to the Deep Graph Library**, the **second-largest GNN library** and the algorithm was later used in the lab's research publication.

Independent Study in Speech Synthesis: After working in the System Security Lab, I became more interested in machine learning and the idea of doing research. In the spring semester of my junior year, unable to join a machine learning lab, I pursued an independent study with Dr. Doug Degroot, the Director of the Center of Applied ML & AI at UTD. Inspired by my love of audiobooks and at Dr. Degroot's suggestion, I began exploring expressive speech synthesis. I had three goals for this independent study: (1) read at least 50 papers, (2) replicate a popular speech synthesis model from scratch, and (3) write a literature survey. Balancing this independent study with my seven other classes was challenging, but by the end, I had read about 78 papers, **replicated two architectures (WaveNet and FastSpeech) from scratch**, and produced speech—albeit of poor quality—which taught me the challenges of building a real working system. Writing the literature survey on my own was challenging since I didn't fully grasp the intricacies of the subject, but it greatly improved my scientific writing skills. Later, I shared what I learned about implementing machine learning algorithms and reading research papers on my website.

Amazon: In the summer of my junior year, I interned at Amazon as a Software Development Engineer on the Natural Language Understanding (NLU) Team at Alexa. I learned about the inner workings of their NLU system, the "brain" behind Alexa (the voice assistant). My role involved streamlining the process of correcting translation errors, such as when saying "Turn off the lights" in Spanish is misinterpreted as "Turn off the security system." I developed a tool for linguists to update the system's knowledge base, potentially **saving $12,000 annually in labor costs**. Through this work, I realized the importance of continuously incorporating new knowledge into decision-making systems, which led me to self-study reinforcement learning (RL). I shared what I learned about RL by writing blog posts and creating open-source algorithms to make it more accessible to others. By observing the dichotomy between my team

in software engineering and the sister team in applied science, I realized that engineers implement ideas while researchers develop them, which solidified my belief that research was the right path for me.

Speech and Machine Learning Lab: After returning from my internship at the start of my senior year, I reached out to Dr. Berrak Sisman, who led the Speech and Machine Learning Lab at UTD, to continue my research into expressive speech synthesis. She offered me an REU (Research Experience for Undergraduates) position. My goals were to (1) develop my own novel research and (2) write at least one research paper. I began by revisiting some code I had written called Bayesian Inference, which ensembled the outputs of a speech synthesis model while applying dropout resulting in improved speech quality. Discussions with PhD students suggested it could be novel, but further investigation revealed the technique had already been published. However, this inspired me to pursue a new project: applying the mixture-of-experts technique to the style encoders of text-to-speech (TTS) models for style transfer. After implementing my method on an existing TTS model, I observed a 10% improvement in speech quality and style similarity, based on subjective evaluations of the generated speech samples with out-of-domain reference speech. This work was later **published at the NeurIPS 2024 Workshop on AI-Driven Speech, Music, and Sound Generation**, where I am listed as the **first author**. I will be attending the workshop in December 2024 to present my work. **I received $800 for this research as part of the 2024 Jonsson School Undergraduate Research Award**, which is granted to only two computer science students.

Graduation: After graduating, I received a full-time offer from Amazon for a software engineering position. However, after considering my desire to create AIs that are not currently possible, I turned down the offer and chose to continue my research through a master's degree, with the intention of pursuing a PhD.

Intelligent Robotics and Vision Lab: Starting my master's program in computer science, I decided to pursue a topic more closely aligned with my long-term interest: intelligent decision-making. After exploring the major fields in decision-making, I saw potential for impact in robotics. So, I reached out to Dr. Yu Xiang, who accepted me into his Intelligent Robotics and Vision Lab at UTD. Under his guidance, I began researching autonomous robotic object manipulation. I proposed extending an existing method that trained a policy—a mapping from observations to actions—based on tasks defined in natural language (e.g., "Move the teacup onto the plate"). The policy used 2D image observations to predict the robot arm's movements by forecasting the future trajectories of a subset of pixels in the image observations to guide the predicted action. To address the complexity of mapping 2D observations to actions in a 3D space, I incorporated depth information into the policy's inputs. Since the dataset didn't include depth information, I re-rendered the robotic manipulation dataset using the MuJoCo simulator and validated the trajectories by visualizing them as point clouds. Next, I modified the model to use 3D trajectories and 3D images. Our initial experiments showed slight improvement on some tasks in the LIBERO benchmark compared to the baseline, we are now testing the model on a more challenging benchmark. I developed a new metric to evaluate the difference between predicted trajectories and the robot's actions to identify whether to focus on improving the trajectory prediction or the policy guided by it. The goal of this project is to move toward robots that can be given tasks by humans using everyday language.

Human Computer Interaction: Currently, I am also enrolled in a research course on Human-Computer Interaction. Motivated by this opportunity, I wanted to explore brain-computer interfaces (BCIs), which I believe is the future of human-computer interaction. Inspired by recent advances in recreating images from brain signals during visual imagination, I began investigating how this technology could be used to create art directly from brain signals. Further research revealed a major limitation in these recreated images—they fail to accurately reproduce colors. This led me to hypothesize that we might only need the outline of a shape to generate art, which could then be input into a sketch-to-art model. However, I encountered two challenges: I needed guidance on conducting experiments with human participants and required access to an electroencephalogram (EEG) for the experiments. To address these issues, I wrote a research proposal and discussed the project with Dr. Rawan Alghofaili, the professor leading the course. After she expressed interest, we began collaborating on the project. Currently, together with two PhD students from Dr. Alghofaili's lab, I am working on the application for institutional review board (IRB) approval to recruit participants for the study. I personally purchased an EEG to conduct the experiments.

So, far we have implemented the method using an open-sourced EEG dataset and will be testing it on real individuals after we get IRB approval. My hope is that this project will inspire others to rethink the potential of BCIs, moving beyond brain signal classification towards higher bandwidth applications of brain-computer communication.

Broader Impact

My primary goal is to make people's lives easier using AI. Outside of research, I pursue this goal by mentoring, teaching, and helping students learn about AI, particularly encouraging them to enter the tech field. I believe this is especially important for individuals from diverse fields, as their domain-specific knowledge can lead to more impactful applications of AI. By equipping those familiar with specific problems in their fields with AI tools, we can drive innovation more effectively than the other way around. Therefore, I have tried to accomplish this in two main ways: by teaching people about AI and by helping students get jobs in the tech field.

Teaching: When I first began learning about machine learning, I found the mathematical concepts intimidating. This inspired me to help others by explaining these ideas in a more accessible way, only requiring a basic understanding of mathematics (such as the understanding of a slope of a line). To that end, I became the workshop lead for the Artificial Intelligence Society, a campus organization, where I developed and **presented eleven workshops**, and led a team of four to create nine of those workshops. I particularly encouraged students from various backgrounds to attend my workshops, even if they didn't have a background in computer science. These workshops covered areas such as Deep Learning, Computer Vision, Natural Language Processing, Model Deployment, and Reinforcement Learning. The **largest event had around 300 participants**, while on average, each workshop attracted 57 students. Our feedback data showed an average **engagement rating of 4.1/5.0**, with **96% of participants indicating they learned something new**. In addition to the in-person workshops, my website, where I share what I've learned about AI and research over the years, has gained a small following, with some posts receiving a few hundred visits. Moving forward, I plan to continue educating students through workshops, and YouTube videos, presenting AI techniques in a simple and accessible manner to foster a diverse and vibrant AI community.

Resume Consultations: After struggling with creating a compelling resume for two years and receiving over 600 rejections during my internship search, I gained valuable insight into what makes a strong resume. Eventually, I secured five offers. Wanting to help others facing similar challenges, I offered free resume consultations to students via LinkedIn and campus flyers. Out of 78 sign-ups, I **selected eight disadvantaged students** who could benefit the most from my guidance. For each student, I provided a one-hour consultation, offering feedback on their resumes, followed by a detailed three-page report with suggestions on wording, organization, content, and formatting. I also followed up with additional support after they revised their resume. As a result, **five of the students were able to secure jobs**. In the future, I plan to continue offering resume reviews and career guidance, with a focus on supporting disadvantaged students from diverse backgrounds in entering the technology field to bring diverse perspectives to many of the key challenges that will be present in the future.

Future Goals

After completing my master's, I plan to continue my research in robotics by pursuing a PhD in Computer Science and later becoming a research scientist. In my research career, my primary goal is to develop autonomous robots capable of reliably performing a wide range of tasks defined by humans, using natural inputs like text. My experiences in both research and industry have equipped me with the technical and interpersonal skills necessary for a successful research career. Additionally, my experiences running workshops and helping students find jobs has instilled in me a deep commitment to educating and assisting others. Receiving the NSF Graduate Research Fellowship would enhance my future research by giving me the flexibility to explore areas I find both interesting and impactful, as well as the freedom to continue to help other students. Awarding the fellowship to me would enable the NSF to further its mission of advancing artificial intelligence in robotics and broadening participation in the tech field.