Personal, Relevant Background and Future Goals Statement

Personal Statement

In the summer before my senior year of high school, I got the chance to participate in the Hampshire College Summer Studies in Mathematics. HCSSiM is a six week program for gifted students who want to learn more about how math is done in the real world. One of the most salient characteristics of HCSSiM for me was that we weren't taught mathematical techniques directly. Rather, we were guided towards discovering these ideas ourselves through trial and error and focused discussion. The first night at the program, all students are asked a benign question: "how many pieces can you cut a watermelon into with n cuts?". By working together, and extrapolating the problem to higher dimensions, we were able to see that this problem had a beautiful combinatorial structure that linked it to Pascal's triangle. The time I spent in this program was the first exposure I had to how math is done in the real world. I began to understand for the first time that math was about more than solving problems, but about noticing patterns and developing new ideas through experimentation. I internalized a conception of math as a dynamic research field that I could contribute to, rather than a fixed set of formulas and algorithms to be memorized.

In college I began to gravitate more towards research. I was approved for a Summer Undergraduate Research Fellowship, studying audiovisual illusions. This particular area fascinated me because of its accessibility - illusions are a universally fascinating phenomenon, and it is a thrill to be able to share your work in a way that people can actually see and hear themselves. I recruited many of my friends to be test subjects in my experiments that summer, and it was always a pleasure to explain afterward the underlying reasons why they had seen what they thought they had seen. Illusions also fascinate me because of their connection to cognition and the structure of the brain. Even simple illusions can reveal inaccuracies in the brain's processing mechanisms that underlie fundamental structural features. The idea of being able to study thought through the brain's ability to recognize patterns was intriguing, since I see pattern recognition as a main characteristic of intelligence. Over the course of the summer, I got to work with postdocs and professors on my project. Together, we discussed and refined different versions of our experiment to make our illusion as strong and as measurable as possible. My mentors on the project always valued my input, and I learned about what it meant to work in a research group. I continued working on the project for the better part of my senior year, and with an award from the George W. Housner fund, I eventually traveled to a vision sciences conference to present my work. Getting to go to the VSS conference made me feel like an active member of the scientific community. I explained my work to interested scientists from across the country, and I got to attend talks and ask poster presenters questions as a peer.

The summer after I graduated, I was emailed about an open position as a Junior staff member for HCSSiM. Remembering the a great experience I had had when I was a student there, I volunteered for the position. I ended up in a position to pass on the lessons I was taught those many years ago: The value of teamwork, the importance of recognizing the structure of problems, and the ability to generate new ideas. I am proud of what my students were able to accomplish at HCSSiM, and I developed an aspiration to not only continue to learn, but to continue spreading my ideas to others.

Relevant Background/Intellectual Merit

At the end of my undergraduate career, I had graduated with a double major in Computer science and math, and I had done research in psychology. I wanted to move forward with research in what I see as the intersection of these fields: Theoretical Machine Learning. Machine Learning uses Computer Science to model and understand the capability that we as humans have to learn from our previous experiences. My math background made me interested in the theoretical aspect of Machine Learning: not only would I be studying a computational version of intelligence, but I would be giving it a rigorous mathematical treatment. I was particularly interested in neural networks, a technique in ML that has recently exploded in popularity, and which is based on parallels with the brain. In 2018, I joined the Ph.D. program at the University of Illinois, Urbana-Champaign, advised by Prof. Matus Telgarsky. Over the past year, we have been studying the essential limits on the computational power of neural networks. One thread of our research studied the massive advantages in power that deep neural networks seem to have over shallow ones. We managed to prove theorems that show how even networks with multidimensional inputs have this advantage in power. Although this work has met with success, it still doesn't give the complete picture of what functions neural nets can represent. To improve our understanding of this area, Prof. Telgarsky and I are now studying single-layer networks, the structure of which seems to have much to say about the capabilities of networks with more layers. Besides my research on neural networks, I have also sought a more broad understanding of the ML field as a whole. As a graduate student, I have taken/audited classes on Statistical Learning Theory and Reinforcement Learning in order to study other perspectives on the field.

Broader Impacts

One thread of my research on neural networks has turned into a paper "Size-Noise Tradeoffs in Generative Networks", which has been accepted for spotlight presentation at this year's Conference of Neural Information Processing Systems. Generative networks apply the ideas of neural networks in a way that allows a computer to be creative by creating new examples of data. These generative models have a wide arrays of applications, even in fields far from computer science such as chemistry and medicine, where generative nets can hypothesize new chemical structures or drugs, or in the arts, where generative nets can create original music or imitate the artistic style of famous painters. What my work shows is a fundamental fact about the structure of these models: that it is possible to compensate for a lack of randomness in the initial stages of the network processing pipeline by increasing the size of the network just a small amount. It reveals a beautiful interplay between dimensionality and complexity in neural nets. The upshot of this is that many of the design decisions made about input format in the implementation of these networks matter very little when compared to decisions about the structure of the network itself. It is my hope that resolving this question will lead to research focused more on these important structural questions without having to worry about design details. I am excited to present my findings at this year's NIPS conference, because I feel I can give the other attendees a very clear picture of why these design tradeoffs work the way they do. I have an opportunity to teach others fundamental facts about generative networks, which is incredibly exciting. In a wonderful coincidence, part of the work is based on the "watermelon problem" that I had first learned so many years ago at HCSSiM — It just goes to show how important it is in math and science to share your discoveries, so that others can build off them even more.

Future Goals

I want to continue with my mission to describe universal truths about how knowledge and understanding develops in artificial systems. Eventually, I hope to have a complete view of what neural networks can do and why they are so well suited to answering questions that humans are interested in about data. As a graduate student at UIUC, I am required to be a Teaching Assistant during my time as a Ph.D. student, and I hope that I get the chance to do it more than once. Prof. Telgarsky has suggested that I take a position as a TA for one of his courses on the theory of Machine Learning, and I am eager to do so. As a TA, I will get the chance to apply my understanding to show others how to approach new problems. I look forward to passing on my view of both the advances we have made in this field, as well as the limits of what we can achieve.