Personal, Background, and Future Goals Statement

Introduction

My interest in robotics and machine intelligence began with witnessing my grandmother's mobility challenges as she aged. There is something agonizingly unjust about losing your physical agency, and seeing my grandmother experience the loss of hers, including her reluctantly giving up hosting family holiday celebrations, has been a perpetual source of motivation that drives me to make assistive, intelligent robots a reality. I am compelled to work on the fundamental problems of robotic control and learning because solving them has the potential to break down physical human limitations.

Working with robotics and computer systems has always seemed inevitable to me. I began programming modest websites when I was 10, attended a STEM magnet high school with a focus on computer science and electrical engineering, and through college and beyond have embraced opportunities to work in numerous challenging research and software engineering roles in academic and industrial environments. While studying abroad in London, I got my first taste of machine learning research, and received a Best Paper award at an NLP conference workshop for my project. In the past year, I have contributed to three papers that were accepted to major robotics and machine learning conferences and workshops. I want to spend the next stage of my life diving deep into these fields, and continuing to push forward the frontiers for what robots can achieve in the real world.

Relevant Background

In the spring of my Junior year (2016), I conducted a research internship at the Machine Reading Lab at University College London under the supervision of Sebastian Riedel. Our project focused on machine learning for language understanding, specifically **learning embedding vectors for new tokens** based solely on their descriptions. We chose a set of tokens, Emojis, and demonstrated that we could use neural nets to quickly learn embeddings that were useful in downstream sentiment analysis tasks. Over 10% of Twitter posts and 50% of Instagram posts contain Emojis, but at the time most NLP models simply ignored them when trying to understand posts. We showed that using the new embeddings alongside traditional embeddings, we could increase absolute sentiment classification accuracy of posts containing Emojis by 8.4%. We published our results as a **conference paper at the EMNLP 2016 Workshop on Social NLP, where I was honored with a Best Paper award** and, as first author on the paper, I was asked to give a 30-minute oral presentation on the work. This experience was formative for my understanding of how to train neural networks and conduct impactful research.

This initial experience convinced me of how powerful machine learning could be in solving real problems, and led me to explore how learning could be used in computer vision tasks. As a Princeton Senior, I conducted my thesis work at the Princeton Neuroscience Institute with Professor Sebastian Seung (letter writer). Under his supervision, I **explored 3D reconstruction of neural tissue from a series of 2D electron microscope images**. I experimented with a diverse set of deep convolutional neural networks, including 3D U-Net and Residual architectures, achieving results approaching state-of-the-art performance on various benchmark tasks. My approach was ranked 14th out of hundreds on the ISBI global leader board for 2D segmentation of neural tissue. My thesis project led to recognition in the Princeton Computer Science department, and I graduated magna cum laude. The techniques we employed are ones used in state-of-the-art real-world vision systems, giving me important insights into how learning might be useful for robotic vision.

Through these experiences, I realized that many of the problems facing robotics can potentially be ameliorated with learning methods. Over the past year, I have conducted fundamental research at the intersection of machine learning and robotics as a Machine Learning Research Engineer at the Samsung Artificial Intelligence Center in New York, led by Professor Daniel D. Lee (letter writer, faculty at Cornell University). Earlier this year I contributed to a project that sought to build a system that could **teach a robotic arm to manipulate household objects using a passive end-effector**. Our approach revolved around the insight that, given a large amount of high-quality geometric information about a scene at training time, we could use traditional robotic planning algorithms like RRT to generate accurate ground-truth manipulation trajectories. We then used these high-quality trajectories as a supervision signal to train a neural network to predict manipulation trajectories directly from noisy RGB-D images. Our method achieved a roughly 90% manipulation success rate on the class of objects it was trained on (shoes), and generalized to new class instances not seen during training. The results of this project are to be **published as a full conference paper at IROS 2019**.

At the AI Center, I have also investigated several research problems in deep reinforcement learning that are relevant to robotic control. In early 2019, I led a project that sought to improve the training and runtime dynamics of Deep Q-Learning for continuous control tasks. The work introduced a novel training setup, whereby we trained a Q-network off-policy using the cross-entropy method (CEM) to induce a policy, and then independently trained a policy to mimic the behavior of this CEM policy. Because CEM has major computational bottlenecks, the policy we trained to imitate it achieved a 120x savings in computational complexity, while maintaining the high performance of CEM-based Q-learning. We **published our results at the ICML 2019 Workshop on Reinforcement Learning for Real Life**, where I presented the work as co-first-author. More recently, our group sought to understand how deep reinforcement learning agents can learn to explore sparse reward environments. This problem is crucial for robotic control in the real world, where tasks like grasping are most-naturally phrased with sparse reward functions (i.e. getting reward only when you complete a task). Our group developed a method, QXplore, that uses an adversarial Q-function to maximize the value-prediction error (TD-error) of a standard Q-function. We demonstrated that we were able to learn high-performing policies on a number of simulated robotics tasks with 40% fewer training steps than competing exploration methods. Our work will be **published as a paper at the NeurIPS 2019 Workshop on Deep Reinforcement Learning**.

I believe that to be successful in a highly-technical and complex field like intelligent robotics, one must excel in the engineering of systems as well as in making theoretical contributions. Through a diverse set of software engineering experiences, including at Google and Microsoft, I have proven myself to be a very capable software engineer and system-builder. While working at Google immediately after graduation, I was promoted in my first year, when the median time to first promotion is roughly 2 years. I also proved my ability to work on complex, large-scale problems when I **led an organization-wide effort to understand how geographical data changes flowed through dozens of Google's Maps API services**. This involved coordinating a large amount of work between multiple teams in 5 different time zones, and eventually enabled the safe release of several country-level data changes.

As I move forward into my graduate education, I intend to draw on my expertise in machine learning, robotics, and software engineering to investigate how our existing understanding of robotic control and design can be synthesized with powerful learning algorithms.

Intellectual Merit

My research and industry experiences demonstrate both my aptitude for and commitment to sustained research efforts in the fields of machine learning and robotics. I have made key contributions to **four papers published at major venues** in Machine Learning (NeurIPS 2019, ICML 2019), Robotics (IROS 2019), and NLP (EMNLP 2016), and won Best Paper at one of them. I derive a substantial amount of pleasure from digging into difficult theoretical problems that also have practical applications, and have a strong desire to continue working on such problems at the intersection of machine learning and robotics.

My liberal arts education has also shaped my research. As a Princeton undergraduate I was exposed to many different fields outside my computer science major, including taking several electrical engineering, mechanical engineering, and entrepreneurship classes, as well as Chinese language and History courses. Additionally, I completed several graduate-level courses in artificial intelligence while studying abroad at University College London. I am particularly proud that I was able to achieve such breadth while diving deep into the subfields of artificial intelligence and machine learning.

Broader Impacts

I am a firm believer in the power of enabling those around you to achieve their potential, and to make STEM fields more inclusive. I spent several years in **leadership roles in the Princeton Entrepreneurship Club**, where I helped organize Princeton's biannual 500-person hackathon, HackPrinceton, and led the Careers team, which offered interview-prep workshops and resume review services for the student body.

I also realized a **passion for academic mentorship as a Lab TA** for several Princeton computer science courses. As a Lab TA, I worked to solidify students' understanding of fundamental computer science concepts, and helped them debug and complete programming assignments. I worked hard to understand others' thought processes, and to find the root of their misunderstanding when they were stuck. This experience was intensely rewarding, because I had a major positive impact on how students thought about computer science and programming.

A core component of my worldview is a commitment to the democratization of technology. Technological progress in the last few decades has brought about unprecedented changes to how we lead our lives. But the benefits of technology have not been felt evenly - society has yet to combat the problems of inequality and privacy that have arisen because of technological progress. As a researcher in these fields, these issues are near and dear to my heart. I want my work and the work of my peers to benefit all of humanity, not just a select few. I intend to **advocate for just applications of AI and robotics**, and aspire to be a technical resource for policy-makers as society continues to grapple with the implications of emerging technologies, perhaps by serving on local or federal technology panels in an advisory capacity.

Future Goals

I aim to place useful robots into the homes of average people. After graduating from a Ph.D. program, I intend to continue working directly on the technical challenges of intelligent robotics, either in an industrial or academic research lab. In the long term, I want to start a consumer robotics company and bring robotic appliances that perform useful tasks in the home (i.e. cooking, cleaning, organizing) to market. In particular, I want to focus on creating robotic assistants to empower people with physical limitations (due to age or disability) to maintain independence and agency in their own homes.