Introduction
I want to create robots that positively impact people's everyday lives on a personalized level within homes and workplaces. These robots will not be limited to one task or a set of tasks, but rather will adapt to new and challenging environments. We need algorithms that allow for adaptation and autonomous robot learning in the real world, and self-supervision is a key part of enabling robots to learn dynamically without requiring humans to teach them. My experiences working on machine learning applications coupled with my robotics background position me well for my academic career in robot learning. My aim is to become a professor, making an impact not only through this work but also through my teaching and mentorship of future generations of researchers, who will further expand the impact of robot learning.

Intellectual Merit
On my all-girls high school FIRST Robotics Competition team, I discovered the excitement of using code to make a robot aim and shoot basketballs. I enjoyed extending my robot's capabilities by implementing vision systems and PID controllers. However, I learned that even with these capabilities, the robot struggled to generalize–it was brittle in the face of uncertainty, failing under lighting differences and sensor errors. Before matches, I had to go out onto the field with the robot to reprogram thresholds, hardcoding the specifics of that environment. Experiencing first-hand the promises and challenges of robotics inspired me to study computer science at MIT. Since then, I have seen and contributed to real successes in machine learning, as it generalizes from data to make systems more resilient.

As an undergraduate, I spent a summer doing research with Professor Brendan Frey at his University of Toronto startup, Deep Genomics, where I applied machine learning techniques to problems in genetics. The model I developed achieved state-of-the-art results on an RNA pattern recognition task and could demonstrate which nucleotide causes a genetic disorder called contractural arachnodactyly. My paper [1] describing these results was accepted for an oral presentation at the 2016 Machine Learning for Computational Biology workshop at NeurIPS. My experience at Deep Genomics helped me realize the importance of a solid academic foundation, including a PhD in machine learning, to pursue a career in research. This goal inspired me to explore areas where machine learning research and model innovation are even more central: computer vision and robotics.

Back at MIT, my research in Professor Antonio Torralba's group kindled my interest in computer vision, specifically in unsupervised learning and generative models. For example, I worked on "predicting the future" by generating frames in a video. Through this project, I learned how leveraging large amounts of unlabeled data can build better models. In video prediction, the data is self-supervised as future frames serve as annotations for the past. Deep neural networks currently work best with large amounts of labeled examples, but if humans can achieve high-quality perception and understanding with few labeled examples, machines should be able to as well.

I tackled the issue of limited labeled data working with Professor Anca Dragan on self-driving cars at Waymo. My team had trained a driving model to match past human driving segments, providing a supervisory signal for a task that is difficult to specify. However, models learned from demonstrations have two challenges: it is hard to impose safety constraints, and the learned system cannot exceed the performance of the expert (human driver). I improved upon the imitation-learned model by fine-tuning with reinforcement learning, using a reward function that penalizes behaviors such as exceeding the speed limit or being too close to other vehicles. My framework leveraged available data while allowed the car to not only match, but exceed the performance of human demonstrations.

Through my computer vision and computational biology research, I discovered how large-scale machine learning can make rapid progress on important applications: data abundance leads to machine learning success. At the same time, I realized machine learning's shortcomings; the dependence on large quantities of labeled data has prevented machine learning from having real-world success in robotics, where physical limitations make data collection expensive. I am designing algorithms that use self-supervised learning, similar to my undergraduate computer vision research, but in robotics settings like those I worked with at Waymo.

Broader Impact

Teaching: After graduate school, I aim to become a professor, and one aspect of being a professor that particularly excites me is the ability to mentor and teach students. I have taught computer science to middle school girls in Boston, high school students in Italy, and undergraduates and professors at MIT. My passion for teaching began in high school, when I was a teaching assistant for a math club, led programming workshops for my robotics team, and taught at a Khan Academy STEM summer camp. These experiences excited me about the power of one-on-one teaching: I fell in love with those moments when you see a lightbulb go off as a student makes a connection. I continued this form of teaching as a freshman at MIT, where I was a lab assistant for Introduction to Electrical Engineering and Computer Science (6.01).

In January 2016, I was selected as an instructor for MIT Global Teaching Labs. Through this program, I taught 4th- and 5th-year computer science students at a technical school in Italy. This experience allowed me to hone my teaching skills in the presence of a partial language barrier and learn how to balance individual student needs with overall class progress. The following year, I co-developed and lectured for MIT's first Deep Learning course (6.S191), which was taken by three hundred MIT undergraduates, PhD students, post-docs, and even full professors. This was my first full-length lecture in front of a large audience, and I learned how to carefully optimize the material to make it engaging and understandable. My lecture on multi-modal learning from the course has over 11,000 views on YouTube. I am eager to continue teaching at all scales as my academic career progresses. I am passionate about access to education, so I plan to make my course materials available online to reduce the barrier to entry in machine learning.

Outreach: As a woman in machine learning, I strongly advocate for increasing diversity, especially gender and racial representation, in computer science. Through the MIT Society of Women Engineers, I developed #HelloWorld, a semester-long program that teaches web development to middle school girls and encourages them to pursue computer science. I created and taught the curriculum for the program, which included HTML, CSS, and JavaScript. As the director, I recruited other MIT students to act as mentors for the girls, allowing us to pair 30 girls with 10 mentors each semester. Starting this program helped me discover my passion for sharing computer science with students and helping youth from underrepresented groups fulfill their potential. During my sophomore year, I founded Code for Good, an MIT course (6.S187) and student group that brings students together with local nonprofits to foster direct, positive impact using computer science. Now in Code for Good's sixth year, hundreds of MIT students have done projects with 40+ nonprofits. At a time when computer science is driving commercial success in the world, I am proud to have started this program which encourages students to use their skills altruistically and improves the reputation of computer science. I have helped students start similar programs at three other universities.

I have continued investing in outreach programs as a graduate student. I was invited to join the committee for CMU's OurCS conference for undergraduate women in research, where I am working to increase student attendee diversity. I found that some outreach programs (including #HelloWorld) attract students who already have access to opportunities. I led an effort to fundraise and provide travel scholarships for lower-income students to attend OurCS, including international students from Ghana, Ethiopia, and Mexico. By removing cost barriers, I am optimistic that OurCS can empower a diverse group of students to pursue research dreams that would otherwise be out of reach.

To increase the diversity of applicants to PhD programs, I mentor nine undergraduates in AI and help them with first paper submissions to graduate school applications. I am one of 16 women out of 113 PhD students in my department (CMU's Robotics Institute). I started an AI Research Mentorship program with Professor David Held to scale this beyond what I can achieve on my own. In the first iteration, I paired 37 underrepresented undergraduates with 37 PhD student mentors, and in the second year, I doubled these numbers. Our program is in a growing network of mentorship programs across universities: I collaborate with Professor Sergey Levine at UC Berkeley and Professor Chelsea Finn at Stanford and have recently helped start an MIT chapter with Professor Tamara Broderick. In graduate school and beyond, I hope to make an impact by mentoring younger researchers. Teaching and outreach complement the impact of my research. Diversity in robotics means new perspectives, ideas, and applications, all of which contribute to the success of robotics in diverse communities.

Future Goals

During my PhD, I intend to make improvements in sample efficiency so robots can learn from less data and use self-supervision so robots can learn from their own activity rather than from human-constructed labels or rewards. My long-term research goal is to allow multi-purpose robots to move beyond specialized roles in factories and other controlled environments and into unstructured environments like homes and workplaces, interacting with and helping humans on a daily basis. I plan to pursue a career in academia, building a group doing groundbreaking machine learning research with close ties to these applications. The NSF GRFP would give me the freedom to wholeheartedly pursue robotics research and volunteer my free time towards making my field more diverse.

References

[1] **V. Dean**, A. Delong, B.J. Frey. *Deep Learning for Branch Point Selection in RNA Splicing*. MLCB 2016.