Personal Background: One of my most cherished memories from childhood is that of building my first robot: *Cyber*. Like many salient memories, this one was unexpected: a Professor from the US had coincidentally decided to visit my hometown of Coimbatore, India to conduct a LEGO Robotics workshop that my parents enrolled me in. Over a long weekend, I learnt to design, build and program a sumo-wrestling LEGO robot and participated in a short but fierce competition that Cyber narrowly won. Watching Cyber - something that I had created and programmed - behave semi-intelligently in the real-world during that competition sparked a deep excitement and curiosity for intelligent machines that lasted long after Cyber had to be dismantled and returned.

The opportunity to build Cyber inspired me to pursue similar STEM projects over the ensuing years via robotics competitions and science fairs. Many of these competitions encouraged participants to find real-world applications for their robots or science projects, and it was during the course of this that I was introduced to Coimbatore's differently-abled community. While I was surprised to learn first-hand about the unfortunate struggles these people and their caregivers face in so many aspects of daily life, I was more surprised to learn about how technology was helping change their circumstances. I discovered how text-to-speech apps on smartphones were enabling blind people to avail ridesharing services, and how assistive software allowed people with limb disabilities to use computers and even learn to code. Inspired by this, I worked on a successful science fair project to create an ultra low-cost "hands-free" computer mouse. Seeing how such new technology can enable people sparked a passion for using technology to change lives that has guided my decisions ever since.

I got involved with robotics research in college because it seemed like an unprecedented opportunity to explore my twin passions for learning more about intelligent machines and using them to impact lives. Initially, I focused on software systems and hardware because I thought that these were the main aspects that needed to be improved for robots to be more useful to humans. However, I soon realized that a deeper issue is intelligence and developed a passionate curiosity to understand and computationally replicate embodied intelligence. As a result, my research projects have spanned two different categories: leveraging human-collaboration for robotic tasks, and improving autonomy capabilities.

Intellectual Merit

Relevant background: Humans 2 Robots Lab (H2R): since September 2017, Within my first week at Brown University, I eagerly joined Prof. Stefanie Tellex's 'Humans 2 Robots Lab' and began introductory projects under the 'Virtual and Augmented Reality' (VR/AR) sub-group applications of VR/AR to human-robot collaboration. After completing Prof. Tellex's Graduate Student level seminar and learning about the research process, I accepted a summer research internship with the lab. As part of this, I successfully integrated a new mobile manipulator robot with the lab's AR/VR software stack, learnt concepts behind computer networking for robotics, picked up Python and C# programming skills, learnt how to use the Robot Operating System (ROS) and contributing to research codebases. My efforts enabled various lab members to successfully use this robot for experiments. I worked closely with one such colleague to design and conduct experiments that used Object-Oriented Partially Observable Markov Decision Processes (OO-POMDP's) to allow our robot to find various household items given some initial human language input as to their locations. This work soon led to my first conference publication at ICRA 2019.

When asked to come up with an independent project idea, I realized wanted to enable our mobile manipulator to solve complex, real-world tasks like multi-step pick-and-place or even cooking.

Some discussion with a Ph.D. student got me interested in exploring AR as an interface to enable humans to directly specify information relevant to a robot completing a task. I wrote an extended abstract that was accepted to a workshop at HRI 2019 and soon began implementing a full pipeline that would enable humans and robots to communicate via holograms while performing specific tasks. During the course of my this implementation, the same Ph.D. student realized that such a system could be used to extend Semantic Mapping — an existing approach to constructing maps for robots to navigate environments — to include manipulation actions (such as opening doors or turning off switches). I adapted my existing pipeline to this use case, implemented scripts that would actuate the robot based on commands from a MagicLeap One AR headset and helped design and run experiments to validate the new mapping approach. I also contributed to writing a full conference submission that was accepted at IROS 2021.

Intelligent Robot Lab (IRL): since October 2018, Having worked on two different projects that leveraged human collaboration for robots to perform tasks, I came to realize that it is infeasible to have a human specify *all* the information a robot will need to perform useful tasks. I understood that Machine Learning (ML) and other aspects of AI are necessary to build truly robust real-world autonomy. To this end, I got more involved with the 'Intelligent Robot Lab' led by Prof. George Konidaris. Motivated by research topics Prof. Konidaris raised, I proposed and initiated a project titled 'Task Scoping' to develop algorithms that could enable agents to automatically prune irrelevant objects and actions given a specific goal without having to resort to classical AI planning approaches. I initially struggled with this project because I had little background in AI planning or knowledge of how to construct theoretical proofs. However, with persistent effort and the advice of Professors Konidaris and Michael Littman, I devised a new approach to this problem, proved interesting theoretical properties of this approach and validated its usefulness on a suite of domains from the literature. This work led to a recent conference paper submission to AAAI 2021.

In parallel to this work, I got involved with a project aiming to improve the autonomy capabilities of robots. This project aimed to develop a novel Imitation Learning algorithm that could learn a goal-conditioned control policy various tasks like button-pressing and peg-insertion within a regular grid. I helped design and implement a deep neural-network based Behavior Cloning algorithm that could generalize to novel goals at test time. I also collaborated with researchers from Mitsubishi Electric Research Labs (MERL) to design and conduct experiments to verify our algorithm's efficacy on a peg-insertion task. After completing this project and submitting a conference paper, I began ongoing follow-up work to investigate using similar techniques for more sophisticated Imitation Learning algorithms.

Reinforcement Learning and Behavior Lab (RLAB): since February 2019, In order to further pursue improved autonomy for robotics, I joined Prof. Michael Littman's RLAB group and took an independent study course with Prof. Littman to deeply understand the basics of Reinforcement Learning (RL). I made significant contributions to the lab's open-source RL library 'simple-rl', which led to my becoming the Head Teaching Assistant for Prof. Littman's Graduate-Level RL course the following semester. In this role, I helped oversee a Final Project that culminated in 18 different groups submitting short papers to the 2019 NeurIPS Reproducibility Challenge Workshop that were all subsequently accepted. After this, I began actively collaborating with Prof. Littman for the 'Task Scoping' project. His advice and suggestions would become an integral part of the theoretical proofs that were core to the conference paper submission.

Uber Advanced Technologies Group (ATG) Research: May 2020 - August 2020, During the course of my research at Brown, I became curious and motivated to learn about how ML meth-

ods for robotics were being used to solve real-world problems at scale in industry. In an effort to pursue this curiosity further, I joined Prof. Raquel Urtasun's lab at Uber ATG as a Summer Research Intern to work on research for Self-Driving Vehicles (SDV's). There, I became interested in the sub-field of Active Learning and explored how recently published Active Learning techniques could simultaneously improve sample-efficiency and reduce labelling costs for a neural-network model in use at ATG. While the specific details of my project are currently under NDA, a conference paper is currently in preparation.

Future Goals: My research experiences thus far have crystallized my academic research goal: I want to improve the cognitive and collaborative capabilities of robots so that they can serve everyone, including the differently-abled, in their daily lives at both work and home. One promising direction to achieve this is Neurosymbolic: combining modern deep learning algorithms with more classical AI approaches in order to leverage the ability of the former to generalize to variations of tasks given data, and that of the latter to produce explainable, robust plans to accomplish any number of useful tasks. I hope to advance our understanding of Neurosymbolic AI by attempting to build more classically-inspired structure — such as hierarchy and the notion of objects — into modern learning algorithms and also exploring how these algorithms can connect with more classical modules to yield integrated real-world systems. Over the long term, my goal is to enable intelligent, collaborative robots to become as useful and ubiquitous as computers.

Broader Impacts

My research journey thus far has benefited tremendously from the support of various graduate students and Professors. Consequently, I've always felt a strong responsibility to pay this forward. To this end, I've volunteered to give talks and presentations about my undergraduate research experiences at 4 different events for prospective undergrad researchers within Brown CS. Additionally, I recently became one of three Meta-Undergraduate Research Assistant (MURA) within the Brown CS Department to facilitate more undergrad research by holding weekly advising hours, answering student questions and liaising with faculty to create novel research opportunities. Within the first few months of beginning this role, I've helped catalogue research opportunities that have been viewed by over 400 students, which has so far led to undergrads becoming involved in more than 5 different projects. I've also been interviewed on a YouTube channel with over 300 subscribers that seeks to promote diversity within CS Research and generated over 150 views [1]. Through these efforts, I hope to increase diversity within computing research and any provide capable, motivated student with the support they need to tackle the hard problems in AI and CS more broadly. I believe this can mutually benefit both the research community and the students themselves.

It is very important to me that I continue to pay opportunity forward for those that will stand on our shoulders. I will strive to ensure the products of my research are always open-source, clearly communicated and applicable to real-world scenarios. Additionally, I plan to conduct outreach and mentor students who might come from economic or geographical backgrounds that lack the opportunity for them to engage in STEM research. Given my own story, I understand the importance of involving such students in globally-applicable research, and I hope to inspire in them the sheer sense of possibility and wonder that hit me all those years ago when I saw Cyber work. I believe the NSF GRFP would be a tremendous new opportunity that will propel me to my academic and personal goals, and am grateful for the committee's consideration.

[1] https://youtu.be/OOAPni0ZUW8