Statement of Objectives - MIT EECS

Nishanth Jay Kumar (njk@mit.edu)

My research goal is to develop generally intelligent robots capable of solving a wide variety of tasks under the complex and dynamic conditions of the real world. To this end, I am specifically interested in enabling agents to competently solve multiple tasks via **continual or lifelong learning** and **combining learning with planning** by learning models from low-level data that support high-level planning to a variety of different goals. I am also interested in improving the generalization capabilities of current learning algorithms by leveraging **hierarchical structure or informative priors**. Having worked with both modern machine learning techniques and classical AI algorithms, I am excited by the prospect of **integrating Neuro-Inspired and Symbolic AI**, thus my interests also span formal methods, control theory, deep learning and cognitive science. These interests were shaped by my undergraduate research experiences, which have themselves spanned a variety of topics but have ultimately sought to improve the intelligent capabilities of real-world agents.

I first became interested in research when I joined Prof. Stefanie Tellex's "Humans To Robots Lab" within my first week at Brown University. Having never had formal Computer Science experience, I spent my first few months developing programming skills and learning about computer networking, the Robot Operating System (ROS), and other aspects of the lab's software stack. Through exposure to different research projects in the lab, I grew interested in enabling robots to solve household tasks. I joined a project that was developing a novel algorithm for fast multi-object search in household environments using Partially Observable Markov Decision Processes (POMDPs). I contributed to both the implementation and conceptualization of this project: I wrote a series of ROS nodes to connect a physical robot's sensors and actuators to the states and actions specified in the POMDP, and designed a realistic experiment involving the robot searching for 3 objects within a 5-room environment in our lab. The novel algorithm enabled the robot to solve this challenging POMDP and find all 3 objects in only 6 minutes, including the time to plan and execute all actions in the real world. This work led to my first publication [1] at ICRA 2019.

One limitation of this multi-object search work is that it required hand-specified detectors for different objects. After some discussion with colleagues, I realized that holograms from Mixed Reality Head-Mounted Displays (MR-HMDs) could obviate the need for such detectors. I presented a first-author extended abstract based on this idea at a workshop at HRI 2019 and implemented a system whereby a human wearing an MR-HMD could use holograms to indicate an item's initial and goal locations for pick-and-place tasks. I wrote scripts to transform these holograms to the robot's coordinate frame and apply an off-the-shelf planner to have the robot move near the object and attempt to grasp it, then either displace it to the goal location or request corrective input from the human upon failure. While I did this, a PhD student in the lab realized that my pipeline could be used as part of his effort to generalize semantic maps to manipulation tasks. Together, we formulated '*Action-Oriented* Semantic Maps' by borrowing formalisms from Object-Oriented MDP's. Additionally, we conducted experiments that demonstrated novice users were able to utilize our system to solve common household tasks like throwing away trash or turning off a light switch. This work led to a second-author conference publication [2] at IROS 2020.

While my work with MR-HMDs enabled our robot to perform a few impressive tasks, it leaned heavily on human intervention. I found this dissatisfying and joined the "Intelligent Robot Lab" led by Prof. George Konidaris to work on enabling greater autonomy. After several research discussions, I became interested in enabling AI agents to plan to solve multiple tasks within very large and rich domains. This is a challenging problem because planning over such models' state-action spaces is prohibitively expensive and inefficient since solving any particular task often requires only a small, goal-relevant subset of states and actions. Motivated by the intuition that backwards search can help discern what states and actions might be goal-relevant, my co-authors and I devised an algorithm we called 'Task Scoping' that was capable of producing goal-based abstractions for a few problems of interest. We submitted a student-abstract based on this work to AAAI 2020 that was accepted, but subsequently realized our proposed abstraction was unsound (i.e, it might prune states and actions that are actually necessary for planning). After much literature review and helpful guidance from Professors Konidaris, Tellex and Michael Littman, I was able to refine our algorithm, prove that the resulting abstraction is sound and complete, and validate its utility on a suite of benchmark domains from the classical planning literature and even a novel planning domain from the popular videogame Minecraft. This work led to a joint first-author submission [3] to AAAI 2021.

In parallel to these efforts, I developed a curiosity for machine learning (ML) methods because of their ability to automatically learn policies and models from data instead of relying upon human expert modelers. I joined Prof. Michael Littman's "Reinforcement Learning and Behavior" lab and took an independent study course to deeply understand the basics of Reinforcement Learning (RL). I made significant contributions to the lab's "simple-rl" library, which led to my becoming the Head Teaching Assistant for Prof. Littman's Graduate-Level RL course. In this role, I supervised 2 final project groups' attempts to reproduce results from the recently-published GoalGAIL [4] paper and 1 group's attempts to re-implement a supervised learning approach to playing 'Diplomacy' [5]. I also helped oversee 16 other submissions to the 2019 NeurIPS Reproducibility Workshop, all of which were accepted.

In order to combine my interests in ML and robotics, I joined a project supervised by Professors Tellex and Konidaris. This work aimed to develop a novel imitation learning algorithm capable of better generalization than existing approaches by learning a goal-conditioned control policy for different tasks. We focused on the tasks of button-pressing and peg-insertion within a regular grid, attempting to have a real robot generalize to the entire grid after learning the task on only a few grid locations. I co-led the project by helping to devise a novel behavior cloning algorithm to accomplish our objective and designing experiments to validate the algorithm's utility in simulation and in the real world. I implemented a data collection and evaluation pipeline on a KUKA robot arm for experiments on the button-pressing task and collaborated with researchers from Mitsubishi Electric Research Labs (MERL) to implement a similar pipeline on a MELFA arm for the peg-insertion task. This work resulted in a first-author submission [6] to RSS 2020 that was subsequently rejected due to a lack of comparison to state-of-the-art methods like GoalGAIL. To address this, I am currently implementing various state-of-the-art robot learning algorithms in simulation and studying the benefits and drawbacks of our approach with respect to these.

To explore how ML and robotics are helping to solve hard problems in industry, I spent this past summer as a Research Intern at Prof. Raquel Urtasun's lab at Uber ATG. I became interested in the sub-field of Active Learning (AL) for self-driving vehicles and worked on devising a novel method inspired by recently-published AL techniques to simultaneously improve sample-efficiency and reduce labelling costs for a deep neural-network model in use at ATG. Additionally, I gained valuable software engineering skills by learning to design and manage a large codebase in collaboration with a team of researchers as well as write documentation to enable others to build upon my code after my internship. While the specific details of my project are currently under NDA, a conference paper is in preparation for submission to ICCV 2021, on which I expect to be co-first author.

Graduate school is the best next step for me to continue to develop as a researcher and pursue my thirst for scientific discovery, and the Embodied Intelligence group within MIT's world-renowned CSAIL is the premier place to pursue my research goals. My specific research interests align well with **Professors Leslie Kaelbling** and **Tomás Lozano-Pérez's** work on integrating learning and planning via meta-planning and their focus on augmenting current learning algorithms with structures that can be exploited for computational leverage. My interests also align with **Prof. Pulkit Agrawal's** work on hierarchical reinforcement learning and his recent interest in investigating how to specify the goals and constraints of a task to prevent an agent from narrow overfitting or other unwanted tendencies. Additionally, my research goal is well-represented by the work of **Prof. Nick Roy's** 'Robust Robotics Group'. Finally, my interest in combining Neuro-inspired and Symbolic AI is reflected by recent work from **Prof. Josh Tenenbaum's** 'Computational Cognitive Science Group'. With so many groups working on a variety of important problems towards the shared dream of understanding and advancing embodied intelligence, I am beyond excited for the collaborations, challenges and opportunity that being a graduate student at MIT will bring.

References

[1] Arthur Wandzel, Yoonseon Oh, Michael Fishman, **Nishanth Kumar**, Wong Lawson LS, and Stefanie Tellex. Multi-object search using object-oriented pomdps. 2019 International Conference on Robotics and Automation (ICRA), pages 7194–7200. IEEE, 2019.

[2] Eric Rosen, **Nishanth Kumar**, Nakul Gopalan, Daniel Ullman, George Konidaris, and Stefanie Tellex. Building plannable representations with mixed reality. Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2020.

[3] **Nishanth Kumar***, Michael Fishman*, Natasha Danas, Michael Littman, Stefanie Tellex, and George Konidaris. Task scoping: Building goal-specific abstractions for planning in complex domains. arXiv preprint arXiv:2010.08869, in preparation, 2020.

[4] Yiming Ding*, Carlos Florensa*, Mariano Phielipp, and Pieter Abbeel. Goal-conditioned imitation learning. Advances in Neural Information Processing Systems, 2019.

[5] Philip Paquette, Yuchen Lu, Seton Steven Bocco, Max Smith, Satya O-G, Jonathan K Kummerfeld, Joelle Pineau, Satinder Singh, and Aaron C Courville. No-press diplomacy: Modeling multi-agent gameplay. Advances in Neural Information Processing Systems, 32:4474–4485, 2019.

[6] **Nishanth Kumar***, Jonathan Chang*, Sean Hastings, Aaron Gokaslan, Diego Romeres, Devesh Jha, Daniel Nikovski, George Konidaris, and Stefanie Tellex. Learning deep parameterized skills from demonstration for re-targetable visuomotor control. arXiv preprint arXiv:1910.10628, in preparation, 2019.