Research Statement

Lawson L.S. Wong, Khoury College of Computer Sciences, Northeastern University

We are now at a new golden age of robotics and artificial intelligence. Autonomous driving, once an academic dream, is quickly becoming a reality, and promises to be the first of many waves of robots that will penetrate everyday lives. Robots can revolutionize home/office services, eldercare, agriculture, flexible manufacturing, search and rescue, exploration, and more. Given the imminent success of autonomous vehicles, we may expect these other domains to be solved soon as well.

However, there is a fundamental difference that separates autonomous driving from the rest. Driving is a structured task in an engineered environment with clear objectives. Service robots need to complete a wide range of tasks in complex human environments, often with little or abstract guidance. In short, we need *generalization*: to new tasks, to new scenarios, to new forms of user specification. The gap between the two is therefore a machine learning problem.

What are we trying to learn? Most generally, a robot is a function mapping sensor inputs to actuator outputs – "pixels to torques". The high-dimensional nature of such a function clearly presents an intractable learning problem. To support effective robot learning, and therefore robot generalization, we need to decompose the learning problem: make it modular, compositional, and significantly more low-dimensional. **My research agenda is to identify and learn intermediate state representations that enable effective robot learning and planning, and therefore enable robot generalization.**

At Northeastern, I lead the Generalizable Robotics and Artificial Intelligence Laboratory (GRAIL), where we make progress on this agenda from multiple angles. In particular, we LEARN:

- Language: Use natural language as an interpretable intermediate state representation
- Estimation: Develop structured representations and models using the notion of objects
- Abstraction: Understand fundamental methods for state abstraction in reinforcement learning
- Robotics: Motivate and apply our approaches to challenging real-life tasks and settings
- Navigation: Validate our ideas and generate insights using a particularly intuitive problem

Abstraction

At its core, the question of finding good state representations relates to finding a good *abstraction*: What are the things that matter, and more importantly, what details are irrelevant? With sufficient detail, every environment and task is unique. Generalization is only achievable when the similarities between situations are exposed, and irrelevances hidden, at appropriate levels of abstraction.

To study this problem formally, we consider the problem of state abstraction in Markov decision processes (MDPs) and reinforcement learning (RL), where the set of possible states (and actions) is typically too large to enumerate and learn individually from; they must be aggregated in some way. We are currently studying criteria for good state aggregations from an information-theoretic perspective. Using notions from rate-distortion theory and information-bottleneck methods, we first considered what it means to minimally represent MDPs and policies. Similar to lossy compression in communication, which aims to maximally preserve information content while transmitting the fewest number of bits, in an MDP the state space should be kept as small as possible while preserving the ability to take actions with high-value outcomes. By adapting information-theoretic bounds and algorithms to the MDP setting, we proposed new algorithms and bounds for state abstraction [1].

We are currently addressing two key limitations in this promising direction. First, the work described above was limited to small discrete state spaces. Recently, deep reinforcement learning methods have demonstrated remarkable success on image-based domains, where states are represented by images. To scale up our state abstraction approach, we developed a new variational information bottleneck method that learns a low-dimensional continuous encoding of image-based states, similar to variational autoencoders. This encoding is further used to extract a small *discrete* abstract MDP that is an approximate bisimulation of the original MDP [2]. This allows us to apply dynamic programming algorithms on the abstract MDP to obtain policies for new tasks, which can then be efficiently translated into the original MDP via the learned bisimulation.

Second, both works described above rely on having access to near-optimal policies in the original MDP. Ideally, agents should learn good representations of the environment while they are also learning about how to act in the environment itself. A primary challenge is that abstractions learned in this manner are imperfect due to inevitable errors encountered during learning, and these errors are amplified in the planning process, leading to suboptimal behavior. We are currently investigating theories to characterize the effects of such errors, criteria to identify when abstraction errors have occurred, and methods to fix these errors when detected.

We were recently awarded an NSF grant (Award #2107256) [3] to continue our investigation.

Navigation

Apart from developing the foundations of state abstraction, I am equally interested in starting from applications, finding representations for specific domains, and gradually deriving more general abstraction principles. Specifically, I have been focusing on the class of *abstract navigation* problems, where an agent is situated within an environment whose layout has never been seen before, and the agent is expected to navigate to a goal without first training on or even exploring this domain. This task may appear impossible without further guidance, but we provide the agent with additional information: a rough 2-D map illustrating the rough layout of the environment, as well as indicators of its start and goal locations. This is akin to a tourist attempting to find a landmark in a new city: without any further help, this would be very challenging; but when equipped with a 2-D map with a "you are here" symbol and an indicator of the landmark, the tourist can easily plan a path to reach the landmark without needing to explore or train excessively. The abstract navigation problem is an instance of a more general question: **If we provide a rough solution of a problem to an agent, can the agent learn to follow the solution effectively?**

Although the solution is technically accessible via the abstract 2-D map, many challenges need to be overcome to use it effectively. The visual appearance, scale, and perspective in the map is completely different from the agent's. The map lacks much of detail in the real world and may even be incorrect. The agent needs to localize itself and learn how its actions affect its abstract position on the map. We have recently proposed two approaches to addressing this problem. The first approach assumes the rough map is fairly accurate, allowing us to directly plan on it and obtain sequences of subgoals. Each subgoal is then given to a low-level controller that is trained to robustly reach nearby landmarks in the agent's visually rich and egocentric perspective [4]. Zero-shot navigation in novel environments is achieved in this hierarchical fashion. The second approach relaxes this assumption by treating the abstract map as an unstructured image, adopting an end-to-end approach that learns to plan using the supplied map [5]. The correspondence between the abstract map and the agent is implicitly learned via a map-conditioned transition function.

We are currently improving the robustness and efficiency of our approaches, with the objective of scaling up to visually realistic environments and applying on physical mobile robots. We were recently awarded a Northeastern TIER 1 seed grant to continue our work.

Language

Natural language is a particularly compelling form of abstraction. Consider following a recipe or a sequence of instructions; language is expressive yet compact, and relies on intelligent users to fill in the details to complete novel tasks. It is also an effective medium for agents communicating and collaborating with humans, since it does not require additional training for the user, and helps make agents' decisions interpretable. I have been investigating using language for both directions of communication, viewing the robot as a follower and as a speaker. In postdoctoral work, I developed methods for grounding natural language instructions to robot navigation and manipulation behaviors [6, 7, 8]. We are currently extending these approaches to visually realistic environments and complex tasks, and proposing ways to combine language grounding with reinforcement learning.

In parallel, effective collaboration requires *two-way* communication, and therefore I have been pursuing ways of applying natural language *generation* to robotics. For example, natural language can be used to describe a policy learned through demonstration, allowing the human teacher to check whether the correct policy was learned. Natural language can also be used to provide instructions to human users entering a new environment or encountering a new task. For this latter application, we have proposed a "rendezvous" problem, where an agent needs to instruct a user about how to navigate from their current location to an intended destination, in an environment that the agent is familiar with but the user is not, similar to providing directions to a tourist. This problem has generated a surprisingly rich set of research questions, including localizing the user via natural language, planning a path that is convenient to describe to the user, describing the path in a succinct yet accurate way, and tracking and correcting the user via further natural language interaction.

Estimation

In language, navigation, and other robotics applications, the notion of objects is inescapable. We instinctively model the world in terms of these abstract entities and reason about their dynamics and interactions. In doctoral work, I proposed approaches for modeling the world in terms of objects, detecting them from robot perception, estimating their states across space and time, and integrating them with non-object representations of the world [9, 10, 11, 12, 13, 14]. Recently, there has been a resurgence of interest in object-based reinforcement learning, and we are currently investigating methods for discovering objects via interaction, and using the abstraction of objects to learn and act more efficiently. More generally, I plan to investigate how to represent, learn, and use structured world models in the context of model-based reinforcement learning. A core question in this direction is how to design flexible estimators for learned representations and abstractions.

Robotics

As the name of my group suggests, we are ultimately interested in representations that facilitate robot generalization, hence we maintain many active projects in robotics to validate and further inspire our approaches. My work on estimating object-based models and language grounding has been applied to object search [15, 16, 17]. I have been involved in the abstract Markov decision process framework, with applications to robot navigation [18, 19]. Recently, with support from a Khoury seed grant, we have proposed new imitation learning methods for performing challenging bimanual robot manipulation tasks [20]. We are also interested in transfer learning; we recently proposed an approach for using knowledge in related tasks to form action priors for new tasks [21]. A MathWorks microgrant further supports my teaching and research efforts in robotics.

My long-term objective is to identify and learn state representations and abstractions that are relevant for these and more areas of robotics and artificial intelligence.

References

[1] David Abel, Dilip Arumugam, Kavosh Asadi, Yuu Jinnai, Michael L. Littman, and Lawson L.S. Wong. State abstraction as compression in apprenticeship learning. In *AAAI Conference on Artificial Intelligence*, 2019.

[2] Ondřej Bíža, Robert Platt, Jan-Willem van de Meent, and Lawson L.S. Wong. Learning discrete state abstractions with deep variational inference. In *Symposium on Advances in Approximate Bayesian Inference*, 2021.

[3] NSF: RI: Medium: Learning MDP abstractions for autonomous systems using variational methods and symmetry groups. `https://www.nsf.gov/awardsearch/showAward?AWD_ID=2107256`.

[4] Chengguang Xu, Christopher Amato, and Lawson L.S. Wong. Hierarchical robot navigation in novel environments using rough 2-D maps. In *Conference on Robot Learning*, 2020.

[5] Linfeng Zhao and Lawson L.S. Wong. Model-based navigation in environments with novel layouts using abstract 2-D maps. In *Neural Information Processing Systems Workshop on Deep Reinforcement Learning*, 2020.

[6] David Whitney, Eric Rosen, James MacGlashan, Lawson L.S. Wong, and Stefanie Tellex. Reducing errors in object-fetching interactions through social feedback. In *IEEE International Conference on Robotics and Automation*, 2017.

[7] Nakul Gopalan, Dilip Arumugam, Lawson L.S. Wong, and Stefanie Tellex. Sequence-to-sequence language grounding of non-Markovian task specifications. In *Robotics: Science and Systems*, 2018.

[8] Dilip Arumugam, Siddarth Karamcheti, Nakul Gopalan, Edward C. Williams, Mina Rhee, Lawson L.S. Wong, and Stefanie Tellex. Grounding natural language instructions to semantic goal representations for abstraction and generalization. *Autonomous Robots*, 43(2):449–468, 2019.

[9] Lawson L.S. Wong, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Collision-free state estimation. In *IEEE International Conference on Robotics and Automation*, 2012.

[10] Lawson L.S. Wong, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Data association for semantic world modeling from partial views. In *International Symposium of Robotics Research*, 2013.

[11] Lawson L.S. Wong, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Not seeing is also believing: Combining object and metric spatial information. In *IEEE International Conference on Robotics and Automation*, 2014.

[12] Lawson L.S. Wong, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Data association for semantic world modeling from partial views. *International Journal of Robotics Research*, 34(7):1064–1082, 2015.

[13] Lawson L.S. Wong. *Object-based World Modeling for Mobile-Manipulation Robots*. PhD thesis, Massachusetts Institute of Technology, 2016.

[14] Lawson L.S. Wong, Thanard Kurutach, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Object-based world modeling in semi-static environments with dependent Dirichlet-process mixtures. In *International Joint Conference on Artificial Intelligence*, 2016.

[15] Lawson L.S. Wong, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Manipulation-based active search for occluded objects. In *IEEE International Conference on Robotics and Automation*, 2013.

[16] Xinkun Nie, Lawson L.S. Wong, and Leslie Pack Kaelbling. Searching for physical objects in partially-known environments. In *IEEE International Conference on Robotics and Automation*, 2016.

[17] Arthur Wandzel, Yoonseon Oh, Michael Fishman, Nishanth Kumar, Lawson L.S. Wong, and Stefanie Tellex. Multi-object search using object-oriented POMDPs. In *IEEE International Conference on Robotics and Automation*, 2019.

[18] Nakul Gopalan, Marie desJardins, Michael L. Littman, James MacGlashan, Shawn Squire, Stefanie Tellex, John Winder, and Lawson L.S. Wong. Planning with abstract Markov decision processes. In *International Conference on Automated Planning and Scheduling*, 2017.

[19] Dilip Arumugam, Siddarth Karamcheti, Nakul Gopalan, Lawson L.S. Wong, and Stefanie Tellex. Accurately and efficiently interpreting human-robot instructions of varying granularities. In *Robotics: Science and Systems*, 2017.

[20] Fan Xie, Alexander Chowdhury, M. Clara De Paolis Kaluza, Linfeng Zhao, Lawson L.S. Wong, and Rose Yu. Deep imitation learning for bimanual robotic manipulation. In *Neural Information Processing Systems*, 2020.

[21] Ondřej Bíža, Dian Wang, Robert Platt, Jan-Willem van de Meent, and Lawson L.S. Wong. Action priors for large action spaces in robotics. In *International Conference on Autonomous Agents and Multiagent Systems*, 2021.