

Learning Robot State Representations and World Models

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.**

In my thesis, I focused on identifying useful state representations, specifically within the context of mobile-manipulation robots. World models in mobile robots traditionally only represent occupancy of the world, because navigation was the primary task. For mobile manipulation robots, *objects*, the primary target of the manipulator, are essential. I proposed treating objects as first-class entities in world models, and introduced methods for acquiring and tracking information about objects. Crucially, because of partial observability in the world, I developed approaches for aggregating object information from perception across space [ISRR 2013, IJRR 2015], time [IJCAI 2016], and sensing modalities [ICRA 2014]. This work provided essential knowledge about objects for mobile-manipulation, enabling robots to perform inherently object-based manipulation tasks. For example, by considering long-term co-occurrences of objects within a home environment, I showed that mobile-manipulation robots can more effectively search for items, finding items using fewer costly manipulation actions [ICRA 2013, ICRA 2016].

In current work as a postdoctoral fellow, I have sought to understand the implications of having such an object-based world model, assuming it has already been acquired. In collaboration with students and colleagues at Brown, we have developed methods for planning in abstract state spaces using abstract models, while being able to ground the results to the real world within the same framework. We have looked at how knowledge of objects and other abstract concepts such as rooms and places enables more effective human-robot communication and collaboration. I have also started to consider how robots may *learn* appropriate state representations and abstractions, first by considering abstractions that humans commonly use. **Ultimately, I want robots to understand user-specified tasks, automatically identify relevant state representations, and acquire abstract models that specifically solve the given tasks.**

Current research

My dissertation work only considered how to *estimate* an abstract state, but did not consider how it could be *used efficiently*, nor did it touch upon the more fundamental questions of representation *identification* and *learning*. My current work therefore addresses these themes:

How should agents reason with objects and other abstract entities?

Objects provide a compressed, low-dimensional, and tractable way to describe the world. However, simply describing is insufficient – we need to reason about the world as well, in a similarly abstract fashion. One class of tools for principled decision making under uncertainty is Markov decision processes (MDPs), a fundamental model for reinforcement learning. Although MDPs are highly versatile models, the world is typically considered *only* at the finest level of resolution, which means that every pixel, every speck of dust, every robot joint angle defines a different state, and we must consider all possible states during decision making.

Together with colleagues at Brown and University of Maryland, Baltimore County, I introduced the notion of *abstract* MDPs, which allows us to construct and reason about a *hierarchy* of MDPs [ICAPS 2017]. At the lowest level in this hierarchy, the world is modeled in full resolution; however, most of the crucial decision making occurs in higher-level MDPs, which are significantly smaller and tractable to plan in. Abstract MDPs provide a unified way to reason at the abstract level, such as with objects, *and* seamlessly ground to the full complexity of pixels and continuous torques.

Why objects? What determines a good state representation?

A good representation should maximally compress the state space while maintaining sufficient fidelity about the world to complete the user’s tasks. But there are other important criteria when interacting with human users: the right state representation should allow humans to talk to robots. In work with students at Brown, I considered how to use knowledge of objects when grounding natural language instructions to robot actions [ICRA 2017, RSS 2017, RSS 2018, AURO 2019]. We found that objects and other abstractions provided an effective way for humans to specify *goals*, instead of explicit robot actions. Indeed, it would be inconceivable for a human to ask a robot to fetch an object without being able to refer to the object itself! This highlights yet another way in which objects are crucial to the world model for mobile-manipulation robots.

We have also studied how to ground language to alternative specifications, such as linear temporal logic [RSS 2018]. On a more theoretical side, I have also analyzed representations from an information-theoretic perspective, using the information bottleneck framework [AAAI 2019].

How will state representations be acquired or learned?

The ultimate question in state representations is how the abstraction itself will be obtained. Certainly understanding objective measures for scoring representations, as in the previous point, will be helpful. However, we must also face the facts: empirically, learning representations from scratch is hard. Are there alternative approaches for acquiring good representations?

One source of abstractions is humans. Humans, in particular professional designers, have devised beautiful icons and abstract ways of representing the world. Some of these designs also serve as task instructions, such as a manual for furniture assembly, or a route on a two-dimensional map for navigation. I am in the early stages of studying how to have *robots* understand these abstractions. Humans are very capable in transferring knowledge in these scenarios, for example, successfully assembling a bookcase they have never handled before simply by following the manual. My objective

is to endow robots with this capability too. Specifically, I have been focusing on the class of *abstract navigation* problems, where the user provides a vague map for a novel environment with a route to follow, and the mobile robot should be able to follow the given instructions [AAAI SS 2018]. To achieve this, the agent first needs to learn the connection between real-world first-person sensory information and the abstract state on the map; this is framed as a supervised learning problem. Once the connection is established, the appropriate action to take can be read off the abstract map. The goal of this study is to see whether robots can learn to make use of *existing* abstractions; certainly, this is an easier task than coming up with the representation itself.

Future research directions

My recent work is only the beginning of a long endeavor to have robots automatically identify and learn state representations. I firmly believe that this is a necessary step on the path to general-purpose intelligent robots. I expect my upcoming research efforts to revolve around the following related themes, listed roughly in order of increasing open-endedness:

Learn and extend existing abstractions. (Continuation of final item above.)

While the abstract navigation task mentioned above is a reasonable and familiar first domain, I plan to expand the approach to more complex tasks, in particular one involving manipulation of objects. Furniture assembly, and other dexterous skills, are challenging multi-step manipulation tasks that *we do not know how to solve yet*. I would further argue that learning to *use* existing abstractions is insufficient; an agent has not truly understood the abstraction until it can *produce new instances* of the abstraction. Instead of abstract *navigation*, I want abstract *mapping*: after a robot roams around extensively and builds a detailed map using simultaneous localization and mapping (SLAM), can it then produce a compact *abstract* map, such that humans and agents new to the environment can easily follow? Considering natural language as an abstraction, can we use the same framework to produce natural language instructions for novel tasks? For example, once an agent manages to achieve super-human level performance on a video game using reinforcement learning, can it then abstract its learned behavior and produce an interpretable walkthrough for humans? Here we see that an added benefit of learning to use existing human abstractions is that robots can then use the abstraction, such as natural language, to communicate with humans.

Focus on and track the task-relevant dimensions.

Even if we have good abstractions for every entity in the world (e.g., all the objects), at any point in spacetime, a robot can only manipulate at most a few of these entities. For example, we would not care about the arrangement of kitchen objects until we need to cook, and even then, the precise pose of a utensil is irrelevant until we want to physically use it. Thus robots should learn to *ignore most* of the (abstract) world, *most* of the time. The challenge is that what should be “most” changes significantly between tasks. To achieve *robot generalization*, we need automated ways of determining which dimensions of the state space are immediately relevant for the task, and have flexible mechanisms for switching between dimensions of focus. One idea is to use hints from the task specification: if a user specifically mentions certain objects or landmarks, then those should clearly be represented, sought after, and tracked. Experience from previous plans on similar tasks can also be leveraged, both for determining what is useful, and what is *not* useful. Once relevant state dimensions have been determined, we also need automatic ways to estimate unknown parts of the relevant state, a subject of my ongoing work on flexible estimation [In preparation].

Be robust to inevitable model errors and misspecification.

The above approach for finding relevant state dimensions is quite *aggressive*, as opposed to the *conservative* approach of tracking the full state of the world. This means that acquired models will inevitably be wrong, especially when the agent is too aggressive in ignoring parts of the world, or when learned abstractions do not align well with the world’s dynamics. Rather than revert to conservatism, robots should strive to be efficient and actively embrace the possibility of making errors. The problem of detecting and correcting for model misspecification is precisely within the domain of statistics and machine learning, and I hope to foster collaborations in these areas, to apply and adapt their techniques to robotics. Indeed, one of the big challenges data scientists face is model selection, and advances on this front will benefit all data scientists, human *and* robot.

Learn state representations for novel robotic applications.

So far, mainly within reinforcement learning, state abstraction has been applied to navigation tasks. This is perhaps due to the familiarity of the task and the existence of clear abstractions. However, we need to progress to tasks that *we do not know how to solve yet*. Abstractions clearly exist in other robot domains: the notions of “pre-grasp pose” and “caging” for grasping, the various “gaits” in locomotion. Can we discover these and other abstractions, use them to decompose difficult tasks, and generalize to new tasks? We severely lack representations for tactile sensing, soft/deformable robots, manipulation, and more. These areas in robotics likely need significant advances before **robot generalization** is possible, and **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] L.L.S. Wong, L.P. Kaelbling, and T. Lozano-Pérez. Manipulation-based active search for occluded objects. In *IEEE International Conference on Robotics and Automation*, 2013.
- [2] L.L.S. Wong, L.P. Kaelbling, and T. Lozano-Pérez. Data association for semantic world modeling from partial views. In *International Symposium of Robotics Research*, 2013.
- [3] L.L.S. Wong, L.P. Kaelbling, and T. Lozano-Pérez. Not seeing is also believing: Combining object and metric spatial information. In *IEEE International Conference on Robotics and Automation*, 2014.
- [4] L.L.S. Wong, L.P. Kaelbling, and T. Lozano-Pérez. Data association for semantic world modeling from partial views. *International Journal of Robotics Research*, 34(7):1064–1082, 2015.
- [5] X. Nie, L.L.S. Wong, and L.P. Kaelbling. Searching for physical objects in partially-known environments. In *IEEE International Conference on Robotics and Automation*, 2016.
- [6] L.L.S. Wong, T. Kurutach, L.P. Kaelbling, and T. Lozano-Pérez. Object-based world modeling in semi-static environments with dependent Dirichlet-process mixtures. In *International Joint Conference on Artificial Intelligence*, 2016.
- [7] D. Whitney, E. Rosen, J. MacGlashan, L.L.S. Wong, and S. Tellex. Reducing errors in object-fetching interactions through social feedback. In *IEEE International Conference on Robotics and Automation*, 2017.
- [8] N. Gopalan, M. desJardins, M.L. Littman, J. MacGlashan, S. Squire, S. Tellex, J. Winder, and L.L.S. Wong. Planning with abstract Markov decision processes. In *International Conference on Automated Planning and Scheduling*, 2017.
- [9] D. Arumugam, S. Karamcheti, N. Gopalan, L.L.S. Wong, and S. Tellex. Accurately and efficiently interpreting human-robot instructions of varying granularities. In *Robotics: Science and Systems*, 2017.
- [10] L.L.S. Wong. Learning abstractions by transferring abstract policies to grounded state spaces. In *AAAI Spring Symposium on Integrating Representation, Reasoning, Learning, and Execution for Goal Directed Autonomy*, 2018.
- [11] N. Gopalan, D. Arumugam, L.L.S. Wong, and S. Tellex. Sequence-to-sequence language grounding of non-Markovian task specifications. In *Robotics: Science and Systems*, 2018.
- [12] D. Abel, D. Arumugam, K. Asadi, Y. Jinnai, M.L. Littman, and L.L.S. Wong. State abstraction as compression in apprenticeship learning. In *AAAI Conference on Artificial Intelligence*, 2019.
- [13] D. Arumugam, S. Karamcheti, N. Gopalan, E.C. Williams, M. Rhee, L.L.S. Wong, and S. Tellex. Grounding natural language instructions to semantic goal representations for abstraction and generalization. *Autonomous Robots*, 43(2):449–468, 2019.
- [14] L.L.S. Wong. Refresh: The factor-fuse-forget filtering framework. In preparation.