

Transfer Learning in RL: Some Thoughts on States and Representation

Lutz Frommberger*

Transfer Learning and State Mappings

The term "transfer learning" is a fairly sophisticated term for something that can be considered a core component of any learning effort of a human or animal: to base the solution to a new problem on experience and learning success of prior learning tasks. This is something that a learning organism does implicitly from birth on: no task is ever isolated, but embedded in a common surrounding or history.

In contrary to this lifelong learning type setting, transfer learning in RL [6] assumes two different MDPs \mathcal{M} and \mathcal{M}' that have something "in common". This commonality is most likely given in a task mapping function that maps states and actions from \mathcal{M} to \mathcal{M}' as a basis for reusing learned policies. Task mappings can be given by human supervisors or learned, but mostly there is some instance telling the learning agent what to do to benefit from its experience. In very common words: Here is task \mathcal{M} , there is task \mathcal{M}' , and this is how you can bridge between them. This is a fairly narrow view on information reuse. More organic and autonomous variants of knowledge transfer are desirable.

The Crucial Role of State Variables

Knowledge transfer, may it be in-task (i.e., generalization) or cross-task, exploits similarity between tasks. By task mapping functions, information on similarity is brought into the learning process from outside. This also holds for approaches that do not require an explicit state mapping [3, 5, e.g.], that exploit external concepts such as relations or agent spaces that are defined a-priori. What is mostly lacking so far is the agent's ability to recognize similarities on its own and/or seamlessly benefit from prior experiences as an integral part of the new learning effort. An intelligent learning agent should easily notice if certain parts of the current task are identical or similar to

*University of Bremen, email: lutz@informatik.uni-bremen.de

an earlier learning task, for example, general movement skills that remain constant over many specialized learning tasks.

In prior work [1], I proposed generalization approaches such as task space tile coding [2] that allow to reuse knowledge of the actual learning task if certain state variables are identical. This works if structural information is made part of the state space and does not require a mapping function. However, it needs a-priori knowledge of which state variables are critical for action selection in a structural way. Recent approaches foster the hope that such knowledge can be retrieved by the agent itself: e.g., [4] allows for identification of state variables that have a generally high impact on action selection over one or several tasks. But even if we can identify and exploit certain state variables that encode structural information and have this generalizing impact, these features must at least exist. If they do not exist in the state representations, such approaches fail. For example, if the distance to the next obstacle in front of an agent is this critical value, it does not help if the state representation consists of the agent’s position and the position of the obstacle, because then the critical piece of information is hidden and only given implicitly.

Thus, again, the question of state space design becomes evident. How can we ensure that relevant information is encoded on the level of features? Or how can we exploit information that is critical, but not explicitly given in the state representation? Answering these questions will be necessary to take the next step into autonomous knowledge reuse for RL agents.

References

- [1] Lutz Frommberger. *Qualitative Spatial Abstraction in Reinforcement Learning*. Cognitive Technologies. Springer, Berlin Heidelberg, November 2010.
- [2] Lutz Frommberger. Task space tile coding: In-task and cross-task generalization in reinforcement learning. In *9th European Workshop on Reinforcement Learning (EWRL9)*, Athens, Greece, September 2011.
- [3] George D. Konidaris and Andrew G. Barto. Autonomous shaping: Knowledge transfer in reinforcement learning. In *Proceedings of the Twenty Third International Conference on Machine Learning (ICML 2006)*, pages 489–49, Pittsburgh, PA, June 2006.
- [4] Matthijs Snel and Shimon Whiteson. Multi-task reinforcement learning: shaping and feature selection. In *Recent Advances in Reinforcement Learning*, pages 237–248. Springer, 2012.
- [5] Prasad Tadepalli, Robert Givan, and Kurt Driessens. Relational reinforcement learning: An overview. In *Proceedings of the ICML-2004 Workshop on Relational Reinforcement Learning*, pages 1–9, 2004.
- [6] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *The Journal of Machine Learning Research*, 10:1633–1685, 2009.