Learning Compact Relational Models for the Exploration of Articulated Objects

Sebastian Höfer Oliver Brock
Robotics and Biology Laboratory
Technische Universität Berlin

Perception and manipulation in unstructured environments remain a challenge for robots. The main challenge is the high variability of such environments. Consider the task of opening a door: door handles vary significantly in shape and appearance but they share the same working principle—they need to be pressed or pushed to open the door.

How can the robot deal with the problem of high variability? One might imagine two approaches to this problem. On one hand, the robot can memorize all experiences. The robot then decides how to act by looking up all of its previous experiences and selecting the most similar one. We call this the nearest-neighbor strategy. The obvious drawback of this strategy is that look-up time and memory requirements increase with the number of experiences. On the other hand, the robot can try to learn a compact, task-specific representation from its experiences. Although the learning procedure is computationally expensive the representation has low space-requirements and fast look-up time.

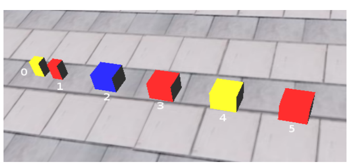

Fig. 1. A simple articulated object. Prismatic joints are located between yellow and red parts.

In this work, we compare these two approaches in the setting of an interactive perception task, namely exposing the kinematic structure of unknown articulated objects [1]. The agent's task is to detect the kinematic structure of an object as fast as possible. The objects vary in size and joint locations, but exhibit a consistent regularity: certain colors encode where a joint is located (see Figure 1). We investigate which of the two approaches mentioned above is more suitable for exploiting this regularity.

In our experiments we model the state space of the agent using relational representations [2]. These representations allow us to equip the agent with background knowledge by defining and grounding task-relevant relations. However, even when properly defining these relations, there still exists a vast number of possibilities how to combine them. Therefore, the agent has to use its experiences to decide how to act. Either it uses the nearest neighbor strategy by performing subgraph matching to compare the current state with its experiences [1], or it uses a learning strategy where it feeds the experiences to a statistical relational rule learner [3]. The rule learner tries to learn a rule set which models the data as precisely as possible while penalizing model complexity.

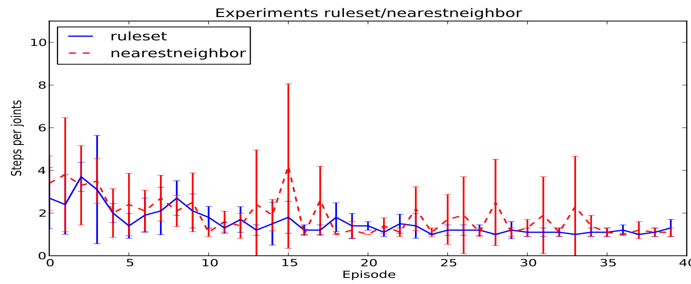

Fig. 2. Number of actions to discover all joints of an articulated objects.

Our results show that the relational rule learner obtains a significantly more compact model without any loss of performance. Figure 2 shows the performance of both strategies by measuring the number of actions required to discover all joints of an object when an ε-greedy strategy with simulated annealing is used. Both strategies converge to an optimal number of actions. However, in terms of compactness the learning strategy outperforms the nearest-neighbor strategy significantly. Whereas the the nearest-neighbor strategy uses an average number of 189.6 experiences (standard deviation 11.5) the rule learner extracts eight successful rules (standard deviation 4.0). This corresponds to a compression factor of about 23. The compact model has additional benefits: The relational rules are human readable and show that the regularities encoded in the world can be learned correctly:

$$\begin{aligned}\texttt{pushLeft}(X):\texttt{prismatic}(X,Y),\texttt{color}(X)=\text{yellow},\\ \texttt{color}(Y)=\text{red},\texttt{rightOf}(X,Y)\\ \Rightarrow\texttt{jointDetected},\texttt{prismatic}(X,Y),\neg\texttt{rigid}(X,Y)\end{aligned}$$

In future work we hope to show that compact models are also more robust to noise and provide better generalization. To do so, we will investigate more complex worlds with more complex regularities.

Acknowledgments

We thank Tobias Lang for providing us the code of the rule learner as well as for helpful discussions.

References

[1] D. Katz, Y. Pyuro, and O. Brock, "Learning to manipulate articulated objects in unstructured environments using a grounded relational representation," in *Proceedings of Robotics: Science and Systems IV*, 2008, pp. 254–261.

[2] S. Dzeroski, L. De Raedt, and K. Driessens, "Relational reinforcement learning," *Machine Learning*, vol. 43, no. 1-2, pp. 7–52, 2001.

[3] H. M. Pasula, L. S. Zettlemoyer, and L. P. Kaelbling, "Learning symbolic models of stochastic domains," *Journal of Artificial Intelligence Research*, vol. 29, no. 1, pp. 309–352, 2007.