Robot Skill Learning based on Interactively Acquired Knowledge-based Models*

Alexander Dürr and Elin A. Topp

Department of Computer Science, Lund University, Lund 221 00, Sweden
{alexander.durr,elin_anna.topp}@cs.lth.se
https://www.lunduniversity.lu.se

Keywords: Ontology-based Knowledge Representation · Hierarchical Deep Reinforcement Learning · Human-Robot Collaboration · Industrial Robot Skill.

1 Problem Statement

Industrial robots are fenced for safety reasons because of hard-programmed robot behavior, containing implicit assumptions about the environment. A knowledge base (KB), containing objects and robot skills, can be built to make assumptions explicit [1] and classic AI search-, reason- and planning-algorithms can be deployed. Advances in robotic behavior representation via Behavior Trees (BT), skills with pre- and post-conditions [2] and Intuitive Programming [3] result in human understandable, explainable and therefore trustworthy task executions facilitating Human-Robot Collaboration. Considering a task to navigate around an obstacle, non-expert robot programmers tend to set a way-point at the start- and end-point, not considering the robot calculating a direct trajectory, resulting in a crash. This shows that despite the above-mentioned efforts, humans still imply assumptions, resulting in incomplete, imperfect and incorrect task descriptions, demanding autonomy from the robotic agent.

Autonomy can be reached with a learning algorithm exploring the continuous robotic action spaces efficiently. Reinforcement Learning (RL) algorithms utilize (raw) data stored in databases to learn a behavior that solves a specific problem under uncertainties [4]. Off-policy learning can be used to learn from human demonstration examples [5]. Finding solutions in high-dimensions combined with reward-sparsity is computationally expensive. If learning is realized in an end-to-end fashion, the found solutions can be non-intuitive, impairing the possibility for Human-Robot Collaboration.

Related work, combining robot behavior representation through BT and learning the structure of the BT, can be found in [6]. Furthermore, the leaves of the BT can be learned, which represent, but are not limited to, robot skills [7].

* This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.

2 Approach and Research Questions

The objective of the project is to combine efficient capabilities in classic AI with RL. A mix utilizing fast, logically proven planning on transferable knowledge and autonomy by learning is proposed.

In a manufacturing setting, it is relatively easy to store a task in a KB with preconditions, describing the start state on a high-level, giving the opportunity for several low-level start states. The post-conditions describe the goal state on a high-level as well. The policy search is driven only by terminal rewards for reaching the goal state from semantically annotated experience and a semantic world model, enabling Actor-Critic RL. On high-level semantic objects a policy can be found, making the policy transferable. The skills are parametrized to enhance flexibility, re-usability, transferability, specifiability, and refinability. Hierarchical RL is used to train a manager, creating a BT, and a worker to find a good parameterization of the skill in the action space.

We want to explore the following research questions: How to represent different skill classes to make them transferable? How to get from high-level task description back to low-level robot instructions? In a later stage: How to autonomously and correctly extend an already populated KB for incremental learning? How to extract relevant high-level knowledge out of low-level data?

References

1. Topp, E. A., Stenmark, M., Ganslandt, A., Svensson, A., Haage, M., Malec, J. : Ontology-Based Knowledge Representation for Increased Skill Reusability in Industrial Robots. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5672-567, IEEE–Institute of Electrical and Electronics Engineers Inc. https://doi.org/https://doi.org/10.1109/IROS.2018.8593566
2. Rovida, F., Grossmann, B., Krueger, V.: Extended behavior trees for quick definition of flexible robotic tasks. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).
3. Stenmark, M., Haage, M., Topp, E. A., Malec, J.: Supporting Semantic Capture during Kinesthetic Teaching of Collaborative Industrial Robots. In: International Journal of Semantic Computing 2018, vol. 12, no. 1, pp. 167–186. https://doi.org/https://doi.org/10.1142/S1793351X18400093
4. Ghadirzadeh, A., Maki, A., Kragic, D., Björkman, M.: Deep predictive policy training using reinforcement learning. In: IEEE/RSJ International Conference on Intelligent Robots and Systems 2017, IROS 2017, pp. 2351–2358. https://doi.org/10.1109/IROS.2017.8206046
5. Ghadirzadeh, A., Bütepage, J., Maki, A., Kragic, D., Björkman, M.: A sensorimotor reinforcement learning framework for physical human-robot interaction. In: IEEE International Conference on Intelligent Robots and Systems 2016, pp. 2682–2688. https://doi.org/10.1109/IROS.2016.7759417
6. Colledanchise, M., Nattanmai Parasuraman, R., Ögren, P.: Learning of Behavior Trees for Autonomous Agents. In: IEEE Transactions on Games (2018). https://doi.org/doi: 10.1109/TG.2018.2816806
7. Sprague, C. I., Ögren, P.: Adding Neural Network Controllers to Behavior Trees without Destroying Performance Guarantees. CoRR, abs/1809.10283 (2018)