Abstract

Reinforcement learning is a class of powerful machine learning methods capable of learning by direct interaction with an environment instead of pre-collected datasets. At the same time, the nature of many tasks with an inner hierarchical structure has evoked interest in hierarchical RL approaches that introduced the two-level decomposition directly into computational models. These methods are usually composed of lower-level controllers – *skills* – providing simple behaviors, and high-level controller which uses the skills to solve the overall task.

While various models of hierarchical reinforcement learning use different architectures, the skill discovery and acquisition remains the principal challenge of this field. Most of the relevant research is focused on resolving this issue, using a broad spectrum of automated methods. Majority of them produce skills that are pre-trained and fixed before the main learning process starts, which may lead to suboptimal skill set, and thus inefficient solution of the overall task.

In this thesis we propose the *Adaptive Skill Acquisition* framework (ASA) aimed to resolve the problem of inefficient hierarchy. It is designed as a universal pluggable component capable of augmenting the existing solutions by new functionality. ASA can observe the high-level controller during its training and identify skills that it lacks to successfully learn the task. These missing skills are subsequently trained and integrated into the hierarchy, enabling better performance of the overall architecture. Besides our new approach, we also provide a review and analysis of available methods for both traditional and hierarchical reinforcement learning.

The conducted experiments on two fundamentally different environments demonstrate the broad applicability of ASA. Embedding the new skills into the hierarchy significantly improves the performance of the overall model, and the ASA-enabled agents exhibit consistent advantage to the baseline. Further ablation tests reveal that the identification of a missing skill is exceptionally robust even with imperfect data, but on the other hand, the elaborate strategies for skill integration do not outperform the baseline ones. A comparative study also confirms that ASA can surpass the previous similarly-oriented model.

Keywords: hierarchical reinforcement learning, skills, adaptive skill acquisition