

Efficient Subgoal Discovery for Hierarchical Reinforcement Learning Using Local Computations

Authors : Adrian Millea

Abstract : In hierarchical reinforcement learning, one of the main issues encountered is the discovery of subgoal states or options (which are policies reaching subgoal states) by partitioning the environment in a meaningful way. This partitioning usually requires an expensive global clustering operation or eigendecomposition of the Laplacian of the states graph. We propose a local solution to this issue, much more efficient than algorithms using global information, which successfully discovers subgoal states by computing a simple function, which we call heterogeneity for each state as a function of its neighbors. Moreover, we construct a value function using the difference in heterogeneity from one step to the next, as reward, such that we are able to explore the state space much more efficiently than say epsilon-greedy. The same principle can then be applied to higher level of the hierarchy, where now states are subgoals discovered at the level below.

Keywords : exploration, hierarchical reinforcement learning, locality, options, value functions

Conference Title : ICITSRL 2021 : International Conference on Intelligent Tutoring Systems and Reinforcement Learning

Conference Location : Barcelona, Spain

Conference Dates : February 11-12, 2021