

Transfer of Samples from Multiple MDPs

Alessandro Lazaric (INRIA Lille-Nord Europe)

October 21, 2011

The objective of transfer in reinforcement learning (RL) is to speed-up RL algorithms by reusing knowledge (e.g., samples, value function, features, parameters) obtained from a set of source tasks. The underlying assumption of transfer methods is that the source tasks (or a suitable combination of these) are somehow similar to the target task, so that the transferred knowledge can be useful in learning its solution. A wide range of scenarios and methods for transfer in RL have been studied in the last decade. In this talk, we focus on the simple transfer approach where trajectory samples are transferred from source MDPs to increase the size of the training set used to solve the target MDP. This approach is particularly suited in problems (e.g., robotics, applications involving human interaction) where it is not possible to interact with the environment long enough to collect samples to solve the task at hand. If samples are available from other sources (e.g., simulators in case of robotic applications), the solution of the target task can benefit from a larger training set that includes also some source samples. In this talk we present some transfer algorithms, their theoretical properties, and preliminary experiments.