

TRANSFER IN DEEP REINFORCEMENT LEARNING: HOW AN AGENT CAN
LEVERAGE KNOWLEDGE FROM ANOTHER AGENT,
A HUMAN, OR ITSELF

Abstract

by Yunshu Du, Ph.D.
Washington State University
May 2021

Co-Chairs: Matthew E. Taylor and Assefaw Gebremedhin

While capable of achieving state-of-the-art performance in complex sequential tasks, deep reinforcement learning (deep RL) remains extremely data inefficient and slow to train. This slow learning speed poses challenges for applying deep RL to real-world situations, especially when poor initial performance is unacceptable or even dangerous. Many approaches have been studied to tackle this problem and one widely used method is transfer learning (TL). The principle of TL is that knowledge acquired from a source agent can be leveraged to assist learning in a different but related target task. This dissertation proposes three types of TL techniques to speed up the learning of a deep RL agent. Specifically, we demonstrate that knowledge can be transferred agent-to-agent, human-to-agent, and self-to-agent.

First, we show that positive transfer can be achieved between two cross-domain agents via direct weight copying if they share visual similarities. Second, we study various pre-training methods using a set of human demonstrations to perform the human-to-agent transfer. Pre-training significantly speeds up the agent’s learning. Third, we explore knowledge transfer from the agent to itself via a novel experience replay framework, namely Lucid Dreaming for Experience Replay (LiDER), in which past experiences are constantly refreshed. Results suggest that the agent can achieve much better performance within the same amount of training data compared to the case without replaying refreshed experiences. Two extensions of the LiDER framework also enable agent-to-agent and human-to-agent transfer, making it a powerful tool to perform all three types of transfer.