Automatic Curriculum Graph Generation for Reinforcement Learning Agents

Maxwell Svetlik, Matteo Leonetti, Jivko Sinapov, Rishi Shah, Nick Walker, and Peter Stone

{maxwell,jsinapov,rishi,nwalker,pstone}@cs.utexas.edu m.leonetti@leeds.ac.uk

University of Texas at Austin University of Leeds

Abstract

Research has shown that transfer learning methods can be leveraged to construct curricula that sequence a series of simpler tasks such that performance on a final target task is improved. Existing approaches typically use handcrafted curricula made by domain-expert humans. To address this limitation, we introduce a method to automatically generate curriculum based on task descriptors and a novel metric of transfer potential. We show that our method produces curricula that improve the agent's learning performance.

Motivation

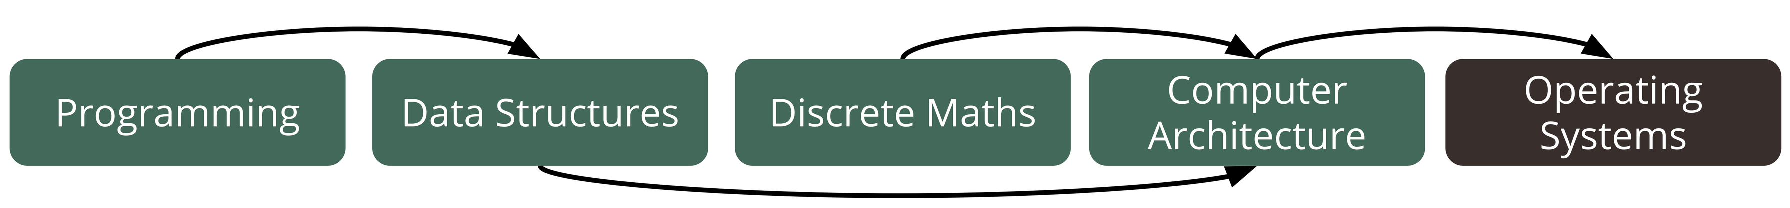

An example computer science curriculum that has Operating Systems as the final target. Many natural curricula repeatedly transfer from multiple sources to a target.

- *Curriculum learning*: accelerate learning a difficult *target* task by learning easier *source* tasks and transferring knowledge
- Existing approaches
  - rely on human task sequencing
  - don't allow transferring from multiple sources

Problem Formulation

- Goal: **reduce the total *experience* required to reach a performance threshold** in the final target task
  - Experience is defined as the total number of *actions* taken
- A set of source tasks and a target task are given
  - Tasks come from a single domain

Algorithm

Inputs

- Pool of source tasks
- Final target task
- Method for assigning task descriptors
- Method for assessing transfer potential

1. Source Task Elimination and Grouping

- Evaluate the transfer potential of each source to the final target, eliminating those with potential less than ϵ
- Assign task descriptors to each task
- Group tasks based on the similarity of their task descriptors

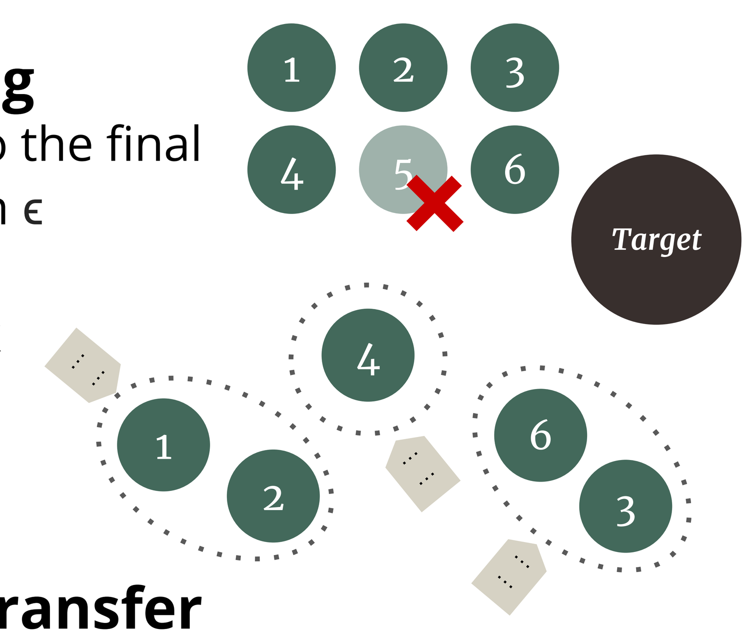

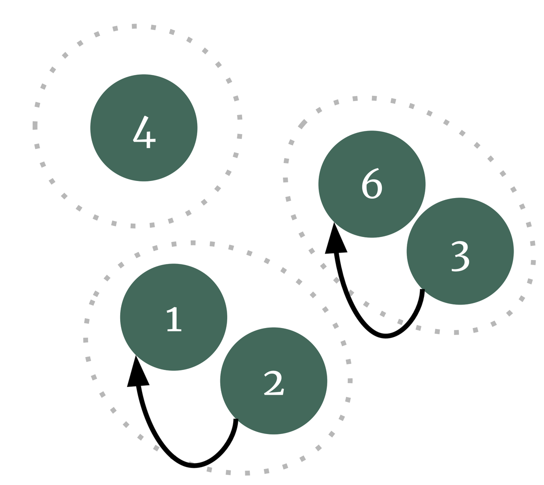

2. Intra-group Transfer

- Evaluate transfer potentials within each group. Assign a transfer edge from each group member to the group mate with the highest potential, subject to the potential being greater than ϵ.

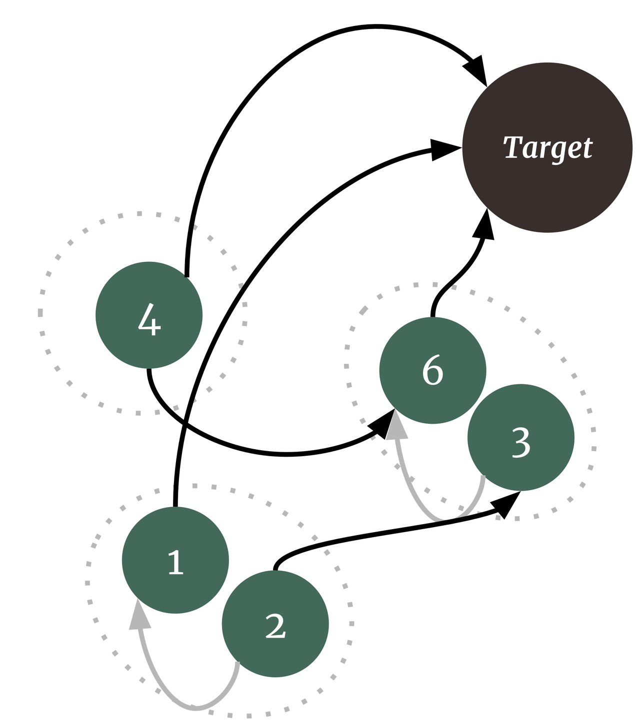

3. Inter-group Transfer

- Assign an edge from each group member with out-degree 0 to the final target task.
- Pick pairs of groups with the most similar task descriptors. Consider the group with the lowest magnitude task descriptor as a source pool and evaluate the transfer potential from each of its members to each target in the second group. Assign an edge from each source to its best target. If the potential is beneath ϵ, do not add the edge.

Output

- Directed acyclic graph
- Topological ordering gives an execution sequence

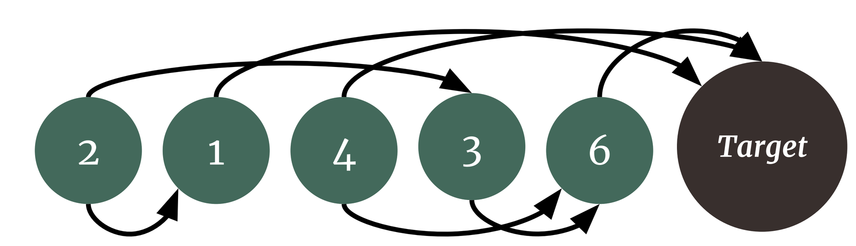

Results

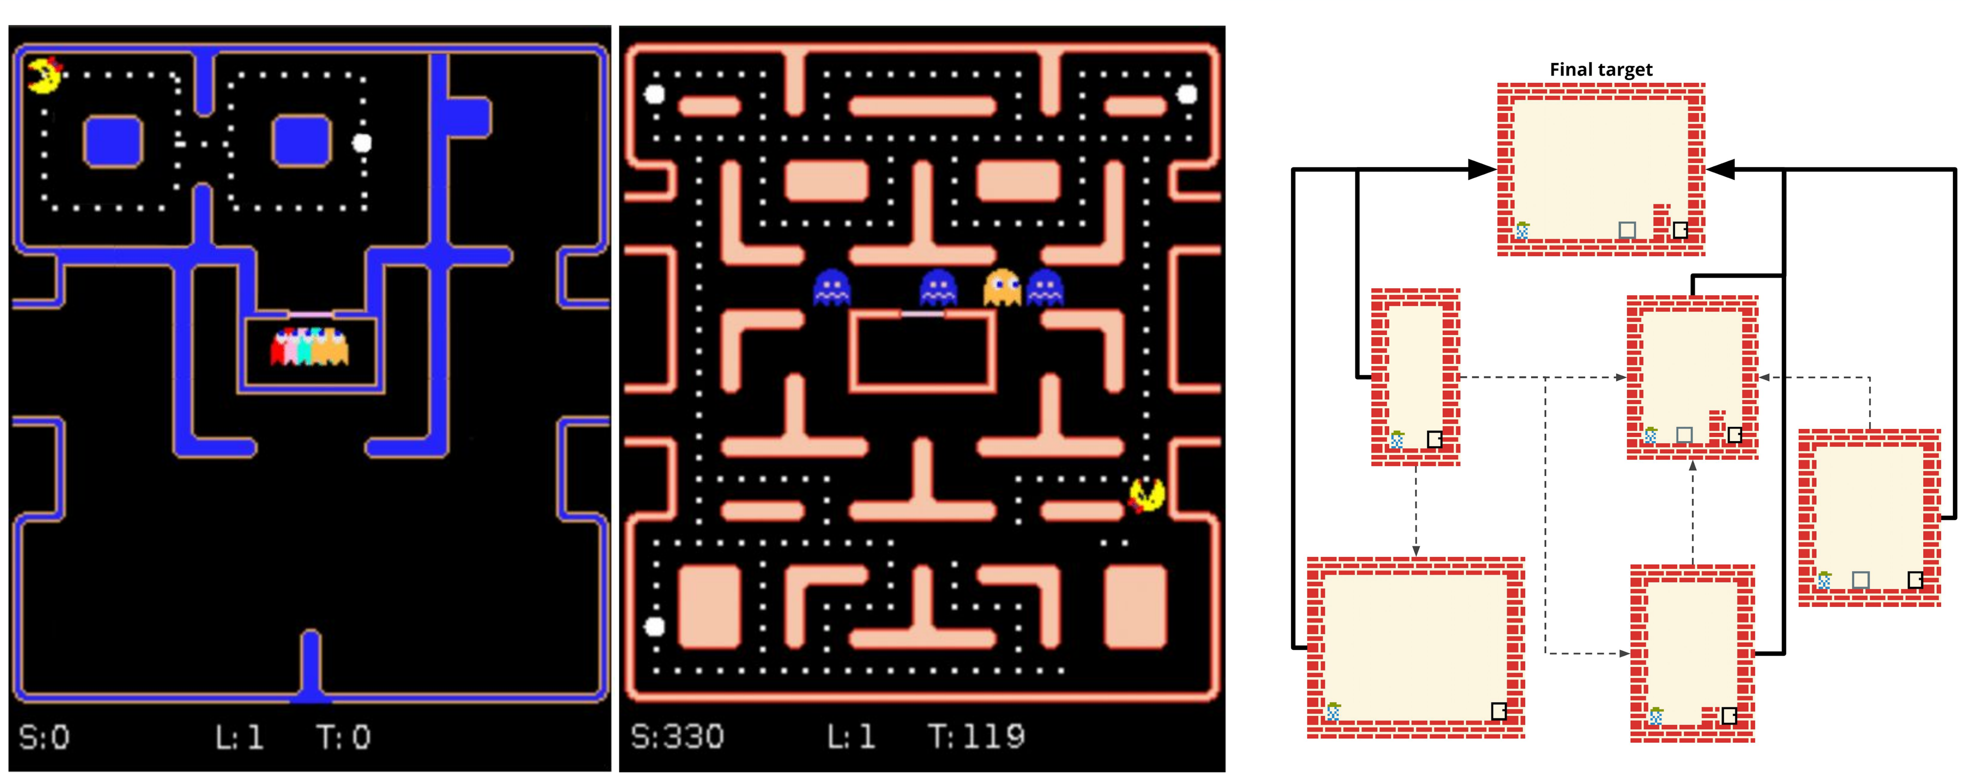

Left: Tasks from the Ms. Pacman domain. *Right:* A curriculum generated from 11 input tasks in the Block Dude domain.

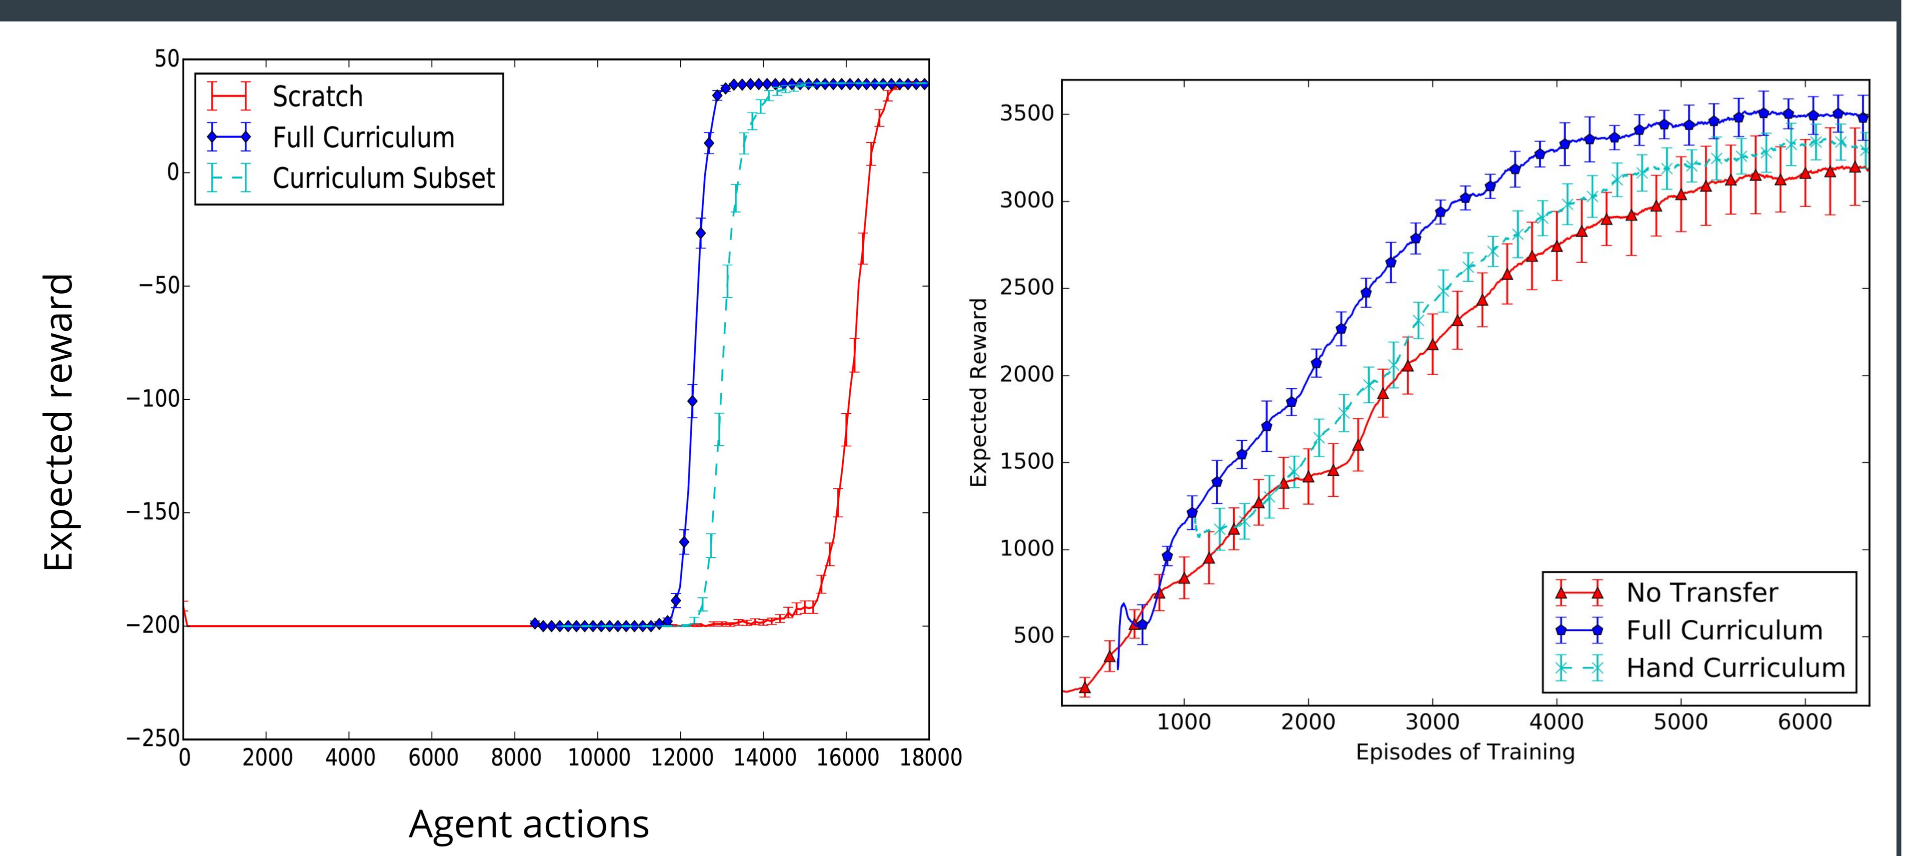

Left: Comparison of learning on the Block Dude domain with different levels of curriculum usage. *Right:* Similar comparison of learning on the Ms. Pacman domain.

Get the paper for more technical details and results
bit.ly/2j55EDr

The University of Texas at Austin Computer Science

This research is supported in part by NSF (CNS-1330072, CNS-1305287, IIS-1637736, IIS-1651089), ONR (21C184-01), and AFOSR (FA9550-14-1-0087).

LARG Learning Agents Research Group The University of Texas at Austin