

Task Grouping for Multi-Task Reinforcement Learning

Measuring relationships between tasks enables disentangling knowledge transfer in multi-task RL.

Geoffrey Harrison

Florian Shkurti

ACADEMIC SUPERVISOR

Mehdi Saeedi

INDUSTRY SUPERVISOR

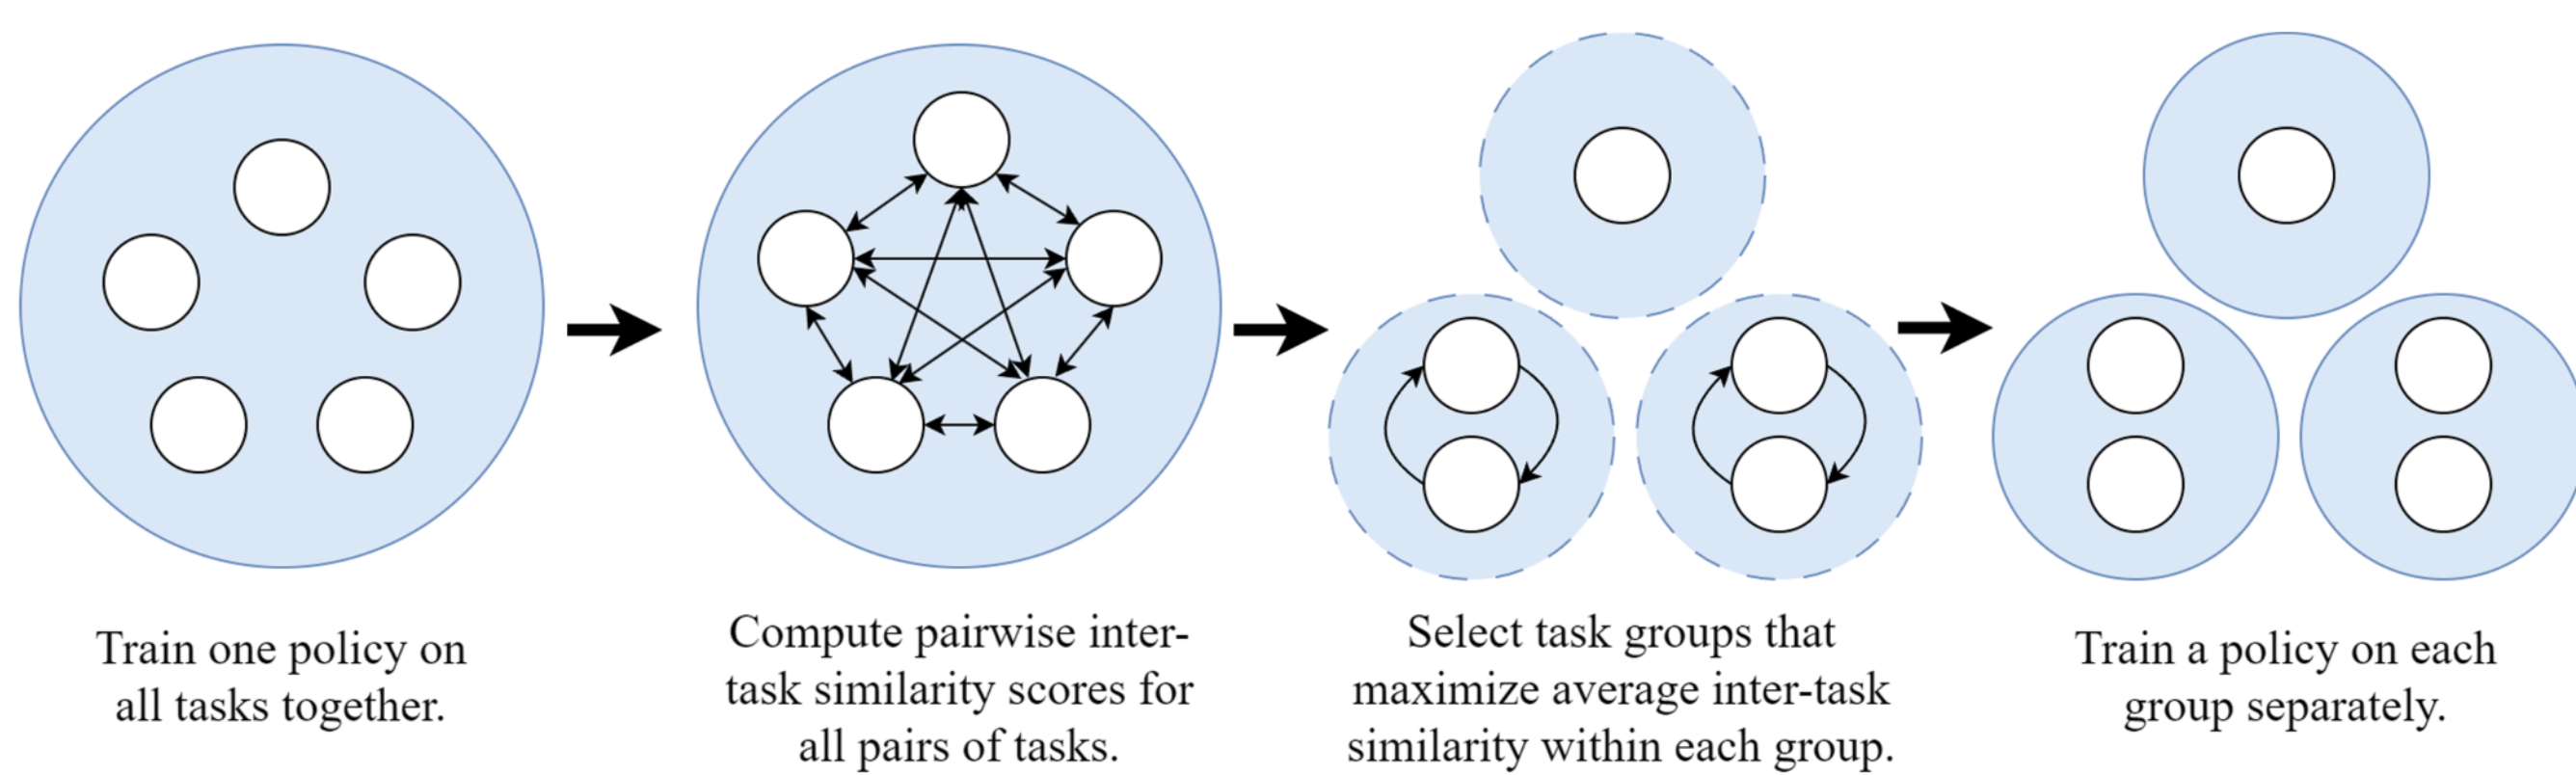


Figure 2: A visualization of our task group selection approach. Each task is represented by a small circle. In this example, 5 tasks are assigned into 3 groups.




Figure 1: An illustration of the 10 tasks of the Meta-World MT10 benchmark.

PROJECT SUMMARY

Despite their impressive accomplishments, reinforcement learning (RL) systems still fall short of human performance when learning a single decision policy for a diverse set of tasks. In fact, naively training a policy on a group of different tasks often results in suboptimal success rates due to the phenomenon of “negative transfer,” where training on dissimilar tasks results in reduced overall success.

A common solution is to first partition the set of given tasks into groups based on a notion of similarity and then learn over each subset separately. However, determining the optimal groupings remains an open problem. In this project, we proposed a strategy to identify meaningful task groupings for multi-task RL (MTRL) based on measures of similarity between tasks. Validating our strategy on the well-known Meta-World reinforcement learning benchmark, we observe that a good choice of task similarity measure results in groups that considerably improve the agent’s ability to solve the benchmarks.

