Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)

Transfer of Knowledge through Collective Learning

Mohammad Rostami
University of Pennsylvania
Philadelphia, PA
mrostami@seas.upenn.edu

Abstract

Learning fast and efficiently using minimal data has been consistently a challenge in machine learning. In my thesis, I explore this problem for knowledge transfer for multi-agent multi-task learning in a life-long learning paradigm. My goal is to demonstrate that by sharing knowledge between agents and similar tasks, efficient algorithms can be designed that can increase the speed of learning as well as improve performance. Moreover, this would allow for handling hard tasks through collective learning of multiple agents that share knowledge. As an initial step, I study the problem of incorporating task descriptors into lifelong learning of related tasks to perform zero-shot knowledge transfer. Zero-shot learning is highly desirable because it leads to considerable speedup in handling similar sequential tasks. Then I focus on a multi-agent learning setting, where related tasks are learned collectively and/or address privacy concerns.

Introduction

The most common setting for a machine learning problem is to learn a single task in isolation. Regression, classification, and reinforcement learning (RL) are standard task examples in machine learning literature. However, learning these tasks is mostly a dynamic, time-dependent process and might involve multiple tasks. As lifelong learners, humans tend not to learn tasks in isolation. Through evolution, humans have developed skills to extract information from learned tasks, store knowledge, and reuse this learned knowledge on future tasks to avoid redundant learning and improve performance.

In the machine learning literature, the above strategy has been investigated through the study of transfer learning (TL) and multi-task learning (MTL). The goal in TL is to use what has been learned from source tasks to improve performance on a target task. Although transfer learning can increase performance on a target task, the learning is unidirectional from source tasks to the target task. In contrast, MTL focuses on learning a group of related tasks more efficiently. The goal of MTL is to optimize a shared knowledge repository to transfer knowledge and jointly to learn task models. MTL can also be performed in an online setting, where the tasks arrive sequentially (Ruvolo & Eaton 2013). In some cases, zero-shot transfer is feasible. That is, after learning few tasks, it is possible to learn a new task *without* data.

Most MTL algorithms are designed to learn task models simultaneously in a batch by a single agent. Only recently, suitable TL and MTL algorithm have been designed for lifelong learning of multiple consecutive tasks (Ruvolo & Eaton 2013), combining the goals of TL and MTL. The current task is learned efficiently with the help of knowledge transfer from previous learned tasks. Then the new obtained knowledge is used to update the knowledge repository. My goal is to study and explore lifelong learning in a multi-agent scenario, where multiple MTL learners try to learn collectively to improve their performance. Despite practical importance, this scenario has not been well explored in the literature.

Multi-Agent Multi-Task Lifelong Learning

MTL hypothesizes that relationships between the tasks can be used to improve generalization performance and increase learning speed, relative to learning each task in isolation. MTL has been been mostly studied in a batch learning settings. Only recently MTL has been extended to a lifelong learning setting, in which tasks arrive consecutively and learned online (Ruvolo & Eaton 2013). However, most existing approaches consider only a single agent and require that the task data are processed at centrally processing unit. However, given the huge modern datasets and data centers, there exist cases that involve multiple distributed heterogeneous agents, each learning multiple tasks e.g., a number of different service robots operating in different locations. For such cases, we can develop mechanisms to share knowledge between agents without sharing the local data e.g., to preserve privacy, and benefit from knowledge transfer between the agents to improve aggregate performance. Only a few previous works have investigated this scenario for MTL (Wang *et al.* 2005). This problem has not yet been explored in a lifelong learning setting.

To address the above challenges, we can formulate an MTL problem over a network of agents. The topology of this network can describe the communication modes of the agents and their relatedness. We can define two sets of global and local knowledge bases which are learned in an iterative alternating procedure, using techniques from distributed optimization e.g., ADMM, to solve the resulting problem. Our problem would involve a composite "sum of convex" objective function subject to local constraints, enforced by the network topology. Several distributed optimization algorithms have been developed to deal with this type of problems which can handle both online and batch settings. Local knowledge bases are private to each agent, and their corresponding optimization problems are updated individually. Upon updating these local bases, global knowledge bases are updated, sharing knowledge between the local bases, allowing for collaborative learning. Solving this problem in this general form is intractable. My research follows a bottom-up approach by considering a special case first and then generalizing to decentralized multi-agent collective learning.

Previous Work: Zero-Shot Learning Using Task Descriptors

Many machine learning tasks can be described by high level descriptors. Consider a cart pole system to be controlled by an agent using reinforcement learning. This system can be described by high level descriptor such cart mass, pole length, and pole mass. Task descriptors can be used along with the tasks' training data to model the inter-task relationships in MTL and transfer learning settings. This is helpful when dealing with multiple agents, because agents might have joint high-level task descriptions that are common between similar tasks and which can facilitate inter-agent knowledge transfer. More importantly, zero-shot learning (ZSL) would be feasible because the agent can learn to perform new tasks using solely high level descriptions. To model task relations, we follow Ruvolo & Eaton (Ruvolo & Eaton 2013) and assume task parameters can be represented sparsely in a shared low dimensional basis. This basis is refined over time as the system learns more tasks and allows knowledge transfer. To incorporate high level descriptors, the descriptor features are assumed to be linearly factorized using a second latent basis over the descriptor space. The second basis captures relationships among the descriptors, with coefficients that similarly embed tasks based on commonalities in descriptions.

Following the above approach, we developed a zero-shot learning method based on this coupled dictionary formulation that incorporates high-level task descriptors to model the inter-task relationships for a single agent , in both lifelong learning and MTL settings (Isele *et al.* 2005). We demonstrated that using task descriptors improves the performance of the learned task policies, providing both theoretical justification for the benefit and empirical demonstration of the improvement across a variety of dynamical control problems. We applied our approach to the domain of quadrotor control, focusing on zero-shot transfer to new stabilizing control tasks. Our experiments demonstrate that it can predict a controller for new quadrotors through zero-shot learning that has equivalent accuracy to PG, which had to train on the system. As with the benchmarks, it is effective for warm start learning with PG and outperforms the state of the art (Sinapov *et al.* 2015). We also demonstrated that our approach leads to speedups up to two-three order of magnitudes.

Collective Learning between Diverse Agents over a Lifetime

My thesis will address the problem of collective multi-agent learning in MTL and lifelong learning settings. Beyond providing the substantial benefit of collective learning to our distributed learning problem, this effort will also address fundamental problems in distributed multi-task and transfer learning of interest to the machine learning community. The main challenge is to model and define a network topology that can impose suitable constraints on task parameters. Consensus among the neighboring agents is the simplest constraint but I plan to consider less restrictive constraints that can be learned. I will focus on the following questions:

- How can heterogenous agents learn tasks consecutively and ensure optimal performance over all tasks?
- How can distributed agents learn in an asynchronous regime that allows for inter-agent collective learning?
- Under what conditions can we guarantee beneficial transfer between heterogeneous agents?
- What theoretical guarantees can we provide on the sample complexity of multi-agent lifelong learning?
- How can we ensure long-term scalability to numerous tasks and changes in task distributions?
- How can we learn the constraints that can facilitate inter-agent knowledge transfer?

To study the problem of collective lifelong learning with multiple agents, we can extend the objective of Ruvolo & Eaton (Ruvolo & Eaton 2013) because it can support the notion of different agents having different (potentially time-varying) distributions over the tasks by imposing local regularizations and customizing the network topology. From each individual agent's (local) perspective, when learning the current task, its goals are to: 1.) rapidly train the best model for each task, given limited training data for the task, its local knowledge, and knowledge shared by other agents, and 2.) maximize the performance of models for all known tasks by sharing knowledge between its tasks. From the multi-agent system's (global) perspective, agents share useful knowledge to realize network constraints. We can formulate the collective learning problem as heterogeneous agents simultaneously solving an objective online, which optimizes the loss of each task model while encouraging shared structure. All agents benefit from sharing knowledge between their own tasks and between agents, both in terms of learning speed and performance.

References

Ruvolo, .P and Eaton, E. 2013. ELLA: an efficient lifelong learning algorithm. ICML, 28:507–515.

Sinapov, J., and Narvekar, S., and Leonetti, M., and Stone, P. 2015. Learning inter-task transferability in the absence of target task samples. ICAAS.

Isele, D., and Rostami, M., and Eaton, E. 2016. Using task features for zero-shot knowledge transfer in lifelong learning. IJCAI.

Wang, J., and Kolar, M., and Srebro, N. 2016. Distributed multitask learning. AISTATS.