

The Importance of Selective Knowledge Transfer for Lifelong Learning

Eric Eaton

Bryn Mawr College
Computer Science Department
Bryn Mawr, PA USA
eeaton@cs.brynmawr.edu

Terran Lane

University of New Mexico
Department of Computer Science
Albuquerque, NM USA
terran@cs.unm.edu

Abstract

As knowledge transfer research progresses from single transfer to lifelong learning scenarios, it becomes increasingly important to properly select the source knowledge that would best transfer to the target task. In this position paper, we describe our previous work on selective knowledge transfer and relate it to problems in lifelong learning. We also briefly discuss our ongoing work to develop lifelong learning methods capable of continual transfer between tasks and the incorporation of guidance from an expert human user.

Introduction

Versatile agents situated in rich, dynamic environments must be capable of continually learning and refining their knowledge through experience. These agents will face a variety of learning tasks, and can transfer knowledge between tasks to improve performance and accelerate learning. In this context, a learning task can be as simple as discovering the effects of an operator on the environment, or as complex as accomplishing a specific goal — anything that can be learned can be considered a task. As the agent experiences and learns a model for each task, it gains access to new data and knowledge. Since the agent will persist in its environment for some time, constantly learning, it may encounter new tasks that are similar to those that it has already learned.

Knowledge transfer can facilitate the learning of these new tasks by reusing information and knowledge from similar previously learned tasks. In the individual transfer learning scenarios commonly explored in the transfer learning literature, the previously learned tasks are known as the *source tasks*, and the new task is known as the *target task*. In lifelong learning scenarios, the agent will face an opportunity for transfer each time it encounters a new task, with the set of source tasks continually expanding.

As research in transfer learning moves toward lifelong learning, it becomes increasingly important to select the proper source knowledge for transfer to the new task. Realistically, only *some* of the source tasks will transfer well to the target task. Identifying which source tasks will transfer well is non-trivial, and can be the difference between successful

and unsuccessful transfer. This problem is of particular concern, since the transfer of inappropriate source knowledge can interfere with learning through the phenomenon of *negative transfer*, decreasing performance on a task.

In this position paper, we describe our previous work on selective knowledge transfer and relate it to problems in lifelong learning. Then, we briefly discuss our ongoing work toward incorporating these ideas into a lifelong learning system that can learn continually over a variety of tasks and incorporate user guidance to correct failures encountered during the transfer process.

Selective Knowledge Transfer

It is not necessarily possible to select the source knowledge to transfer to a new target task by examining only the surface similarities between the tasks. The selection must support the process of knowledge transfer by choosing source knowledge based on whether it will transfer well to the target task. In our previous work, we developed methods that identify the source knowledge to transfer based on this concept of *transferability* to the target task. Intuitively, transferability is the amount that the transferred information is expected to improve learning performance. It may or may not be possible to determine transferability based on surface similarities between the tasks, depending on the domain.

Selective transfer of instances between tasks

Instance transfer methods reuse data from the source tasks to augment the training data for the target task. We developed a novel task-based boosting method (Eaton and desJardins 2011) that identifies the appropriate source instances to transfer to a target task. Our approach iteratively constructs an ensemble of classifiers, reweighting each source and target instance as in AdaBoost by increasing the weight of mispredicted instances. Additionally, for each source task, the weights of all instances for that source task are updated based on its transferability to the target task. Source tasks that show positive transferability to the target task have their weights increased, while source tasks that show negative transferability have their weights decreased. Therefore, each source instance is reweighted both individually and based on its source task's transferability. This process identifies the source instances that best transfer to the target task while

providing a mechanism to continually adjust all instance weights throughout the boosting process.

The main challenge in applying instance transfer methods to lifelong learning scenarios is ensuring scalability, since these approaches typically store all training instances (and associated parameters) for all tasks. Sampling approaches could help alleviate these problems by storing only a subset of the data for each source task. While instance transfer methods are useful for small numbers of tasks with limited size training sets, scalability problems currently limit their applicability for the longer task sequences encountered during lifelong learning. For these scenarios, model parameter transfer methods ensure better scalability, as described next.

Selective model parameter transfer between tasks

Model parameter transfer methods identify which parameters of the previously learned source models can be reused when learning a target task. A variety of methods for model parameter transfer have been proposed, including using biased regularization (Kienzle and Chellapilla 2006; Marx et al. 2005) and hyper-parameter transfer in hierarchical Bayesian models (Niculescu-Mizil and Caruana 2007; Roy and Kaelbling 2007). In our method for selective transfer of model parameters (Eaton et al. 2008), we first embed the learning tasks on a manifold that captures the transferability between tasks. We estimate the transferability between pairs of tasks using parallel cross-validation and form an undirected graph over the learning tasks with edges between tasks weighted by their transferability. This graph represents a discrete approximation of a manifold that captures the transferability between tasks.

Each location on the manifold can be represented by the parameter vector that characterizes the learned model for the corresponding task at that location. We then learn a function over the manifold using these known parameter vectors as samples of the function. To determine the parameters to transfer to the new task, we first estimate its location on the manifold to identify tasks that would transfer well to it, and then evaluate the function at that location to predict the appropriate parameters that should be transferred in learning the new task's model.

This approach improves scalability by requiring only the storage of the model parameters and limited sets of training data, and by localizing the computation to specific regions of the space. Therefore, it is much more amenable to lifelong learning with an ever-increasing set of learning tasks. However, this method requires some recomputation of the transferability manifold with each new task, involving expensive (but parallel) cross-validation to estimate the transferability.

Ongoing Lifelong Learning Research

We are currently developing methods for transfer in lifelong learning that use a combination of automation and interaction with an expert user to achieve robust, continual transfer between learning tasks. This research builds directly upon our previous work in selective knowledge transfer and our work on interactive learning (Eaton et al. 2010), focusing on three key objectives:

1. **Learning a common set of sparsely shared model components**, each of which captures a chunk of knowledge. The model components are sparsely shared between tasks to allow for scalability of the source knowledge repository. Knowledge is transferred between tasks by sharing common components in the models.
2. **Enabling the transfer of knowledge in reverse to the source tasks** through these shared model components, thereby allowing learning on a new task to improve performance on related previously learned tasks—an essential capability for scalable lifelong learning. This idea of reverse transfer is relatively unexplored in transfer learning and is a novel contribution of our approach.
3. **Assimilating guidance and instruction from expert users** through interaction with the system to provide missing source knowledge, correct mistakes in transfer, and guide the transfer process toward the desired goals.

We are currently developing interactive transfer learning methods based on these objectives, and integrating them into a lifelong learning software framework. To evaluate this approach, we are applying these methods to visual scene annotation in still images and simulated 3D environments, and to multi-modal data fusion for computational neuroscience.

One of the key issues we face in applying selective transfer methods to lifelong learning is bounding the cost of determining the transferability between tasks. Parallel cross-validation provides accurate results, but incurs significant computational expense, especially when there are many source tasks. Therefore, we are also exploring the use of task metadata, model distribution and topological similarity, and user guidance to identify the appropriate source knowledge for a target task. User instruction can also provide alternative training in the case where there are gaps in the source knowledge or mistakes in automated source selection, enabling the repair of transfer errors and improved performance.

Acknowledgments

This paper was supported by the Office of Naval Research under Award #N00014-11-1-0139.

References

- Eaton, E.; and desJardins, M. 2011. Selective transfer between learning tasks using task-based boosting. To appear in *AAAI '11*.
- Eaton, E.; desJardins, M.; and Lane, T. 2008. Modeling transfer relationships between learning tasks for improved inductive transfer. In *ECML '08*, 317–332.
- Eaton, E.; Holness, G.; and McFarlane, D. 2010. Interactive learning using manifold geometry. In *AAAI '10*, 437–443.
- Kienzle, W., and Chellapilla, K. 2006. Personalized handwriting recognition via biased regularization. In *ICML '06*, 457–464.
- Marx, Z.; Rosenstein, M. T.; Kaelbling, L. P.; and Dietterich, T. G. 2005. Transfer learning with an ensemble of background tasks. In *NIPS '05 Workshop on Inductive Transfer*.
- Niculescu-Mizil, A., and Caruana, R. 2007. Inductive transfer for Bayesian network structure learning. In *AISTATS '07*.
- Roy, D. M., and Kaelbling, L. P. 2007. Efficient Bayesian task-level transfer learning. In *IJCAI '07*, 2599–2604.