Transfer Learning and Representation Discovery in Intelligent Tutoring Systems

Kimberly FERGUSON [a], Beverly Park WOOLF [b] and Sridhar MAHADEVAN [a]

[a] *Autonomous Learning Laboratory, {kferguso, mahadeva}@cs.umass.edu*

[b] *Center for Knowledge Communication, bev@cs.umass.edu*

University of Massachusetts Amherst, Computer Science Department

140 Governor's Drive, Amherst, MA 01002

Abstract. We describe a novel framework developed for transfer learning within reinforcement learning (RL) problems. Then we exhibit how this framework can be extended to intelligent tutoring systems (ITS). We compose an algorithm that automatically constructs a graphical representation based on the transfer framework. We evaluate this on a real-world ITS example and show that the model constructed by our approach performs better than previously published results. We propose that transfer learning is a useful and related area to explore for furthering ITS.

Keywords. intelligent tutoring systems, graphical models, transfer learning, student modeling

Introduction

While learning, humans and animals perform a natural abstraction of information that has been difficult for the machine learning community to emulate. Humans and animals do not start each new task or situation with a blank slate; they instinctually reuse and modify knowledge gathered from previous experiences. This use of previously learned knowledge speeds up the learning of new tasks. Within the education community, this ability is often referred to as *generalization*, in AI it has been called *transfer learning* (TL) and is an important area of research.

The aim of TL is to reuse information learned about one situation to improve performance in a new, and related situation. This research to develop Tl methods is immediately applicable to and can be enriched by how humans teach and learn. Additionally, it is a useful tool for enhancing intelligent tutoring systems (ITS). Applying our recent work in TL within reinforcement learning (RL) to ITS further demonstrates that our novel framework is a useful approach, truly related to how humans learn. The real-world example of student modeling shows the advantage of thinking of learning in terms of a proposed taxonomy of: value, knowledge, behavior, and representation transfer.

This paper describes a novel framework identifying different types of information transfer and different situations in which transfer might be used first for RL problems and then how this taxonomy can be extended to ITS.

1. Transfer Learning Framework

In RL, various terminologies have been used to describe transfer learning–which describes different situations where any type of information is reused. The terms knowledge transfer and behavior transfer have been used interchangeably, and the lack of consistency has been detrimental to successfully applying transfer learning concepts in related areas of research. *Transfer situations* are dependent on the difference between the source and target of transfer and will not be described in this paper. We define four different *types of information transfer* dependent upon the type of information being transferred:

- *Value Transfer (v-transfer)* refers to when specific scalar values are learned in one situation and then reused in a novel situation. Value transfer, which pertains to specific numerical knowledge being reused, is a higher level transfer than knowledge transfer, which encompasses broader conceptual knowledge.
- *Knowledge Transfer (k-transfer)* refers to when an agent collects a knowledge base of information and then reuses this information in a novel situation. Examples of k-transfer include the transfer of rules, expert advice, features, and relational information.
- *Behavior Transfer (b-transfer)* refers to when an agent collects a toolkit of one or more behaviors and then reuses these behaviors in a novel situation.
- *Representation Transfer (r-transfer)* [3] refers to when the underlying representation is learned and then reused in a novel situation. For an example from research in human learning, consider the studies that have shown that there is a specific stage when children transition to being able to correctly represent and generalize shapes. An experimental example is when young children try to fit blocks of different shapes into the correctly shaped slots. The purpose is to know that *Slot 1* represents the same shape as *Block 2* and then using that representation information, generalize that *Block 2* fits into *Slot 1*.

2. Skill Model Experimental Results

Our experiments use the Wayang Outpost intelligent tutoring system and involve learning the hidden characteristic of a student's mastery of geometry skills [4,1]. We cannot observe a student's mastery of a skill directly; we have to infer this knowledge from their answers to problems involving these skills [6]. In particular, we built a Bayesian network (BN) based on 159 student pretests of 28 problems using different combinations of 12 skills, use Expectation Maximization (EM) [2] to estimate the missing parameters, then ran inference to predict a student's individual mastery of each skill. Exactly how the skills should be linked within the network is unknown. We have experimented with different structured BN models and use the Bayesian Information Criterion (BIC) [8] to evaluate them.

The previously highest scoring model was a Hierarchical Skill Difficulty Model, hand-constructed based on the idea that students must first learn easier mathematical skills to build a foundation for the more difficult skills. This can be viewed as a transfer problem, where students transfer the information learned from the easier problems to speed up the learning of more difficult skills and solve more difficult problems. We show improved results with a Hierarchical Transfer Learning Model which is automatically constructed based on our transfer framework. The experiment for the Flat Skill Model, in which all skills are assumed to be independent, is the standard baseline to which we compare all other models. The results are displayed in Table 1.

Model	Max BIC	Avg BIC	Std Dev	Var
Flat Skill Model	-5250.7	-5306.6	10.5	111.2
Hierarchical Skill Difficulty Model	-4901.9	-4989.6	45.4	2062.3
Hierarchical Transfer Learning Model	-4630.8	-4726.2	48.4	2342.6

Table 1. Bayesian Information Criterion (BIC) Comparison: Note that maximum BICs and average BICs are based on 30 random runs and the best scoring model is highlighted. Simply put, given a set of data and probability distribution generating the likelihood, the larger the likelihood, the better the model fits.

3. Discussion

We presented a transfer learning framework for RL and extended it to intelligent tutoring systems. We hand-coded the features and algorithm used to automatically construct a student model based on this framework. This Hierarchical Transfer Learning Model outscores the previously best hand-constructed Hierarchical Skill Difficulty Model by virtually the same amount that the published model outscores the baseline Flat Skill Model. This provides evidence that the transfer framework is based on a foundation in education and human psychology. It also exhibits that the transfer framework can be advantageous for large real world problems in addition to its original application—RL domains. The automatically constructed graphical model can then be used to infer an individual student's mastery of each skill. A more precise BN allows for better student modeling. With improved student modeling, we can not only make predictions of how a student will do on a particular problem, but also measure how much a student's skills improve overall. Future work includes generalizing this approach to benefit a variety of ITS domains (other than just geometry) by formalizing the feature extraction and representation discovery algorithm. To conclude, transfer learning research has made advances in AI based on how humans learn and teach and should be studied in conjunction with ITS.

References

[1] C. Beal, I. Arroyo, J. Royer and B.P. Woolf, Wayang Outpost: An intelligent multimedia tutor for high stakes math achievement tests, American Educational Association annual meeting, 2003.

[2] A. Demster, N. Laird and D. Rubin, Maximization-likelhood from incomplete data via the EM algorithm, Journal of Royal Statistical Society, Series B, 1977.

[3] K. Ferguson, Proto-Transfer Learning in Markov Decision Processes Using Spectral Methods. University of Massachusetts Amherst, Technical Report (TR-08-23), 2008.

[4] K. Ferguson, I. Arroyo, S. Mahadevan, B.P. Woolf and A.G. Barto, Improving Intelligent Tutoring Systems: Using Expectation Maximization To Learn Student Skill Levels, Proceedings of the 8^{th} International Conference on Intelligent Tutoring Systems, June, 2006.

[5] E. Hannan and B. Quinn, The determination of the order of an autoregression, Journal of Royal Statistical Society, Series B, 41, 190–195, 1979.

[6] A. Jonsson, J. Johns, H. Mehranian, I. Arroyo, B.P. Woolf, A.G. Barto, D. Fisher and S. Mahadevan, Evaluating the feasibility of learning student models from data, AAAI Workshop on Educational Data Mining, 2005.

[7] K. Murphy, The Bayes Net Toolbox for Matlab, Computing Science and Statistics, 33, 2001.

[8] G. Schwarz, Estimating the Dimension of a Model, Annals of Statistics, 6, 461–464, 1978.