The Role of Transfer in Learning (extended abstract)

Sebastian Thrun
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
http://www.cs.cmu.edu/∼thrun/

Introduction

Virtually all of today's approaches to artificial neural network learning generalize considerably well if sufficiently many training examples are available. However, they often work poorly when training data is scarce. Various psychological studies have illustrated that humans are able to generalize accurately even when training data is extremely scarce. Often, we generalize correctly from just a single training instance. In order to do so, we appear to massively re-use knowledge acquired in our previous lifetime.

Lifelong learning is a framework that addresses the issue of knowledge re-use and inductive transfer in learning. In lifelong learning, it is assumed that the learner faces an entire family of learning tasks, not just a single one. When facing a new learning task, the learner may *transfer* knowledge acquired in previous learning tasks to boost generalization. Three questions are of fundamental importance for any approach to lifelong learning: The *what* that is being transfered, the *how* it is being transferred, and the *when* it is that it is being transferred.

Transfer

To successfully transfer knowledge across multiple learning tasks, a learner must identify aspects that its past (and future) learning tasks have in common. Recent research has produced a variety of approaches that are capable of transferring knowledge across multiple inductive learning tasks (see the survey and references in (Thrun, 1996)). Different approaches differ

- in the way they generalize when facing the first learning task, and
- in the way their generalization is affected when previously learned knowledge is transferred.

Using object recognition from color camera images as an example, a recent study compared a variety of lifelong learning with each other, and with the corresponding conventional learning methods (Thrun, 1996). In particular, we examined the generalization accuracy that was obtained after presenting only a single view of the target object (along with a counterexample). The approaches that were capable of transferring knowledge were also provided with views of five additional objects. The idea was that those approaches could learn some of the *invariances* in object recognition, and change the way they generalize to incorporate these invariances.

The results are remarkable. Those approaches capable of transferring knowledge

	error
Back-Propagation with pre-learned invariances	25.2%
nearest neighbor with pre-learned distance metric	24.8%
neighbor neighbor with pre-learned data representation	25.6%

consistently outperformed those that were not:

	error
conventional Back-Propagation	41.3%
conventional nearest neighbor	39.6%
Shepard's interpolation	39.6%

Moreover, the results seem to suggest that the generalization error merely depends on the particular learning method (*e.g.*, neural network vs. nearest neighbor). Instead, the fact that knowledge is transferred from previous object recognition tasks has the strongest impact on the result.

Selective Transfer

Obviously, in real life not every learning task is equally related to every other one. In the study above, we knew that *all* learning tasks were related in the same way (they all were object recognition tasks), so that all approaches could just blindly transfer knowledge among all of them.

In a second study involving a variety of mobile robot perception tasks (involving the recognition of people, landmarks, locations, obstacles), we investigated the robustness of lifelong learning approaches with respect to un-related tasks (Thrun & O'Sullivan, 1996). The results were not surprising: In cases where all tasks were well-related, transferring knowledge improved the generalization accuracy significantly, especially when training data was scarce:

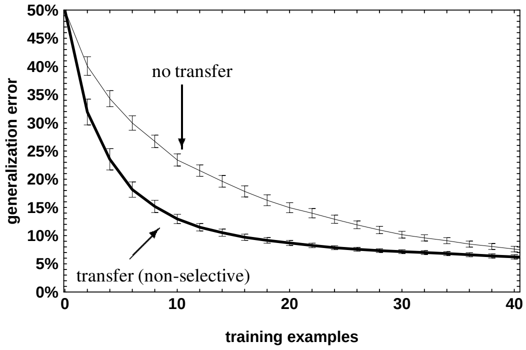

When, however, many tasks were unrelated, transfer did even hurt the overall performance:

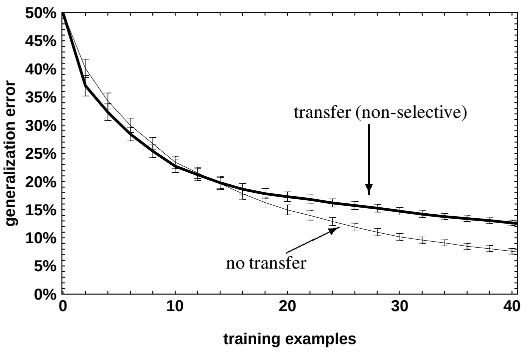

These findings illustrate that blindly transferring knowledge may be problematic in practice.

The TC algorithm transfers knowledge *selectively* (Thrun & O'Sullivan, 1996). It does this by arranging learning tasks into a hierarchy, based on their "relatedness," Relatedness is determined using statistical tests that empirically measure the effectiveness of transfer. The following hierarchy

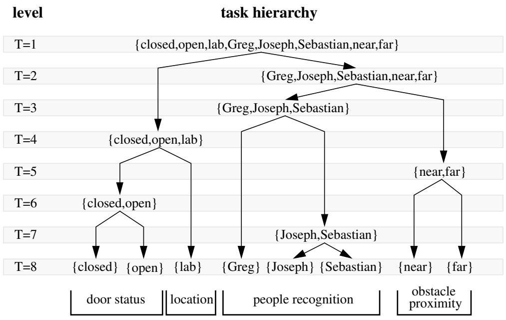

has been obtained in the mobile robot perceptual domain. The most notable result here is that different *types* of learning tasks (namely: tasks involving people, door status, location, obstacles) were grouped into different branches of the hierarchy. In other words, the computer discovered the different *types* of learning tasks.

The task hierarchy enables a learner to transfer knowledge *selectively*, from the most appropriate class of previous learning tasks. The results,

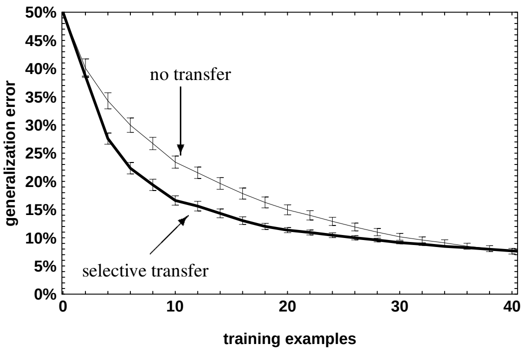

which are superior to those obtained with non-selective transfer, illustrate the role of proper task selection in the transfer of knowledge.

Learning To Act

The ideas presented here are also applicable to reinforcement learning (Sutton, 1991). Reinforcement learning addresses the problem of learning to act from delayed reward. The SKILLS algorithm (Thrun & Schwartz, 1996), a version of reinforcement learning which selectively transfers knowledge across different learning tasks, discovers partial action policies in multiple reinforcement learning tasks based upon a minimum description length argument. These partial policies can be re-used as building blocks in other reinforcement learning tasks.

Initial results, obtained for a simple grid-world scenario, are encouraging:

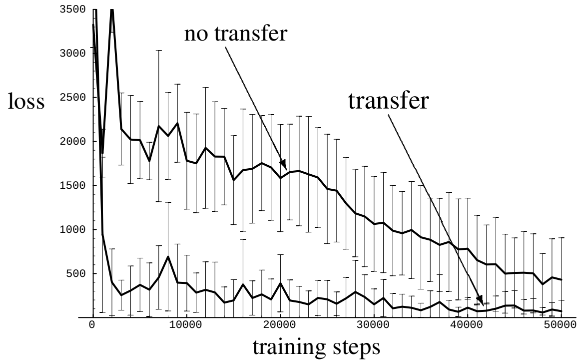

These curves illustrate that reinforcement learning converges faster, if knowledge is transferred from previous learning tasks (in this example: four tasks, two of which are actually related).

These findings are well in tune with results obtained with different learning methods. For example, when training a mobile robot to learn to navigate to a designated target object in an in-door office environment, we also found that reinforcement learning converges significantly faster when knowledge (in this case: neural network action models) acquired in previous learning tasks is being re-used (Thrun, 1996).

Conclusion

We draw three primary conclusions from this research: First, transfer, if applied correctly, is very likely to improve the results of learning, given that more than just a single learning tasks is available. Second, the lifelong learning problem—learning from many related tasks—is easier than the problem of learning from a single task, despite the fact that lifelong learning algorithms tend to be more complex. Third, since we firmly believe that transfer plays an important role in human learning, approaches that transfer knowledge among different learning tasks appear to be cognitively more plausible than approaches that do not.

References

Thrun, S. (1996). Explanation-Based Neural Network Learning: A Lifelong Learning Approach. Kluwer Academic Publishers.

Thrun, S. & O'Sullivan, J. (1996). Discovering structure in multiple learning tasks: The TC algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning*. San Mateo, CA: Morgan Kaufmann Publishers.

Sutton, R. S. (1991). Integrated Modeling and Control Based on Reinforcement Learning and Dynamic Programming, In *Advances in Neural Information Processing Systems 3*. San Mateo, CA: Morgan Kaufmann Publishers.

Thrun, S. & Schwartz, A. (1996). Finding Structure in Reinforcement Learning. In *Advances in Neural Information Processing Systems 7*. Cambridge, MA: MIT Press.