Machine Learning, 28, 5 (1997)

Guest Editors' Introduction

Many inductive learning systems focus on a single task. Research in the area of *inductive transfer* addresses the question of how learning can be enhanced by using knowledge from multiple learning tasks. Transfer can reduce the number of training examples or the number of computer cycles required for achieving a particular level of performance. Recently, the field of transfer in learning has received considerable attention within the machine learning community (see [Pratt and Jennings, 1996]).

As an example, consider how humans recognize people in photographs. We are remarkably good at recognition from just a single image, even when the pose differs, when different clothes are worn, or when the person has aged. For us, learning to recognize a new person is not an isolated learning task. Instead, we face many similar learning tasks over our entire lifetime, providing the opportunity to transfer knowledge between them. For example, once we have learned that the color of a person's clothes is irrelevant for the identity of a person, generalizing across different images is greatly simplified. In general, the presence of more than one learning task provides a synergistic opportunity to learn domain-specific knowledge such as invariance, which can be exploited when learning other tasks. Research on transfer in inductive learning seeks to find computational mechanisms for learning and transferring knowledge across different learning tasks.

This issue contains papers written by some of the leading researchers in the area of inductive transfer. Each paper explores transfer from a slightly different angle, using different assumptions about the learning task(s). Baxter's paper provides a Bayesian analysis of the complexity of supervised learning when training examples for more than one learning task are available, assuming that these tasks are appropriately related. This paper is complemented by Caruana's, which provides experimental results obtained with a particular learning algorithm for learning multiple tasks, along with a discussion as to when this algorithm might be applicable. Ring then investigates transfer in reinforcement learning. His paper presents an algorithm which can exploit knowledge over multiple reinforcement learning tasks whose complexity increases over time. The final paper by Schmidhuber, Zhao, and Wiering also investigates transfer in the context of reinforcement learning. Just like Genetic Programming, their approach learns directly in policy space. It is able to identify policies that modify themselves, thus it is potentially able to transfer learned knowledge within and across multiple learning tasks.

This issue would not have been possible without the transfer workshop held at the 1995 Conference "Neural Information Processing Systems," which was jointly organized by R. Caruana, D. Silver, J. Baxter, T. Mitchell, and by ourselves.

Lorien Pratt, Evolving Systems, Inc.
Sebastian Thrun, Carnegie Mellon University

[Pratt and Jennings, 1996] Lorien Pratt and Barbara Jennings, "A Survey of Transfer Between Connectionist Networks", Connection Science, Vol. 8, No. 2, 1996, 163-184.