

Foreword

This special issue contains five papers selected from those presented at the Fourth Annual Workshop on Computational Learning Theory (COLT'91), held at Santa Cruz, California, on Aug 5-7, 1991. These papers went through the normal review process of *Machine Learning*.

The annual COLT conferences provide a forum for theoretical studies of machine learning. This research area has developed rapidly in the past few years, which have also seen the emergence of such sister conferences as the one on Algorithmic Learning Theory held in Japan and Euro-COLT. The papers presented in this issue are perhaps representative of this emerging discipline. They contain a variety of positive results on the algorithmic possibilities in machine learning, as well as some negative results indicating the limits. Dana Angluin and Donna Slonim prove that it is possible to learn monotone disjunctive normal form formulae in polynomial time with a membership oracle that provides answers only to a random subset of the possible queries. David Helmbold and Philip Long show how to track concepts that are slowly changing over time under the pac learning paradigm. Robert Schapire presents a sophisticated general result which implies, for example, a polynomial time algorithm for inferring read-once Boolean formulae over the uniform distribution. David Haussler, Michael Kearns and Robert Schapire provide a study of learning curves that unifies the statistical physics and Vapnik-Chervonenkis dimension approaches. Finally the paper by Andrew Barron studies the classification accuracy that can be achieved by artificial neural nets through learning. We thank the authors for submitting their contributions in a timely fashion and the referees for their careful reviews.

Ming Li
Leslie Valiant