Preface

This volume contains papers from the final in a series of workshops called "Computational Learning Theory and 'Natural' Learning Systems," or "CLNL" (pronounced "kernel"). The first, held September 1990 in Princeton, New Jersey, was sponsored by Siemens Corporate Research (SCR), MIT and Princeton University. Its purpose was to explore the emerging intersection of theoretical learning research and natural learning systems (COLT, AI and Neural Nets). Over 100 participants attended to hear 16 oral and 25 poster presentations The second workshop, held in September 1991 in Berkeley, California, was sponsored by SCR, MIT and U.C. Berkeley. The third workshop, held in September 1992 in Madison, Wisconin, was sponsored by SCR, MIT and U. of Wisconsin. Finally this, the fourth workshop, was held in early September 1993 in Provincetown, Massachusetts (those of you who took the boat ride recall the temperature!), and was sponsored solely by SCR. Over 50 papers were submitted from 8 countries (including papers from 14 US states); 16 of these were selected for oral presentation and 17 for posters. There were also 4 invited presentations:

- Thomas Dietterich (Oregon State and Arris Pharmaceutical): "Scaling Up Machine Learning: Practical and Theoretical Issues"
- Leo Breiman (UC Berkeley): "Current Tree Research"
- Yann le Cun (AT&T Bell Labs): "Large Heterogeneous Learning Machines for Perceptual Tasks"
- Ron Rivest (MIT): "Reflections on Machine Learning"

About 100 people attended (including the 30 who arrived on our rented boat).

These unusual workshops attracted researchers interested in the intersection of three historically distinct styles of learning research: *Computational Learning Theory* has undergone a renaissance, *Neural Network* research has seen exponential growth and *Symbolic Machine Learning* has become a dominant influence in the AI field. In recent years, perhaps partly due to the "conciousness raising" that occurred due to these workshops (note that many of the program chairs of the prevailing learning conferences have attended a CLNL workshop at least once), there have been attempts to co-locate Machine Learning and COLT, while the premier neural nework conference

(NIPS) has become more and more representative of best work at the interesection of theory and experiment, symbols and signals.

Over the years, workshop participants have been encouraged to examine general issues in learning systems, which could provide constraints for theory, while at the same time theoretical results were interpreted in the context of experiments with actual learning systems. The organizers followed one simple heuristic: actively pursue topics in the intersection of these fields. To maximize the likelihood of lively discussion and substantive exchange, speakers with theoretical topics were juxtaposed with speakers working in experimental domains. We also sought a diverse set of invited speakers, in addition to the ones listed above:

First workshop (1990): Eric B. Baum (NEC Research Labs); Thomas Mitchell (CMU); David Rumelhart (Stanford); Leslie Valiant (Harvard).

Second workshop (1991): Bernardo A. Huberman (XEROX PARC); Nils J. Nilsson (Stanford); Geoffrey Hinton (U. Toronto); David Haussler (UCSC).

Third Workshop (1992): Michael Kearns (AT&T Bell Labs); Richard Sutton (GTE[1]); Alan Lapedes (Los Alamos); Pat Langley (NASA AMES[2]); Grace Wahba (U. Wisconsin at Madison).

These workshops were intended to provide a forum for the development and creation of new areas of research in Learning and Adaptive systems. We hope they have fostered tolerance in theorists for the importance and diversity of experimental approaches to the study of succcessful learning systems and biased empiricists at least to watch the evolving theory of learning systems.

[1]Now at University of Massachusetts.
[2]Now at ISLE and Stanford University.