To retrieve a copy of the paper

"The Relationshop Between PAC, the Statistical Physics Framework, the Bayesian Framework, and the VC Framework",

anonymous ftp to ftp.santafe.edu. The paper is in pub/dhw_ftp, as relating.frameworks.ps.Z.

***

Alternatively, the paper appears in "The mathematics of generalization", D. Wolpert (Ed). Here is a description of that book and ordering information:

TITLE: The Mathematics of Generalization: Proceedings of the SFI/CNLS Workshop on Formal Approaches to Supervised Learning

Edited by D. Wolpert

Other (first) authors:
- L. Breiman
- P. Cheeseman
- J. Denker
- T. Dietterich
- D. Haussler
- G. Hinton
- S. Nowlan
- N. Tishby
- G. Wahba

==========

Table of Contents:

Reflections After Referring Papers for NIPS
Leo Breiman

The Probably Approximately Correct (PAC) and Other Learning Models
David Haussler and Manfred Warmuth

Decision Theoretic Generalizations of the PAC Model for Neural Net and Other Learning Applications
David Haussler

The Relationshop Between PAC, the Statistical Physics Framework, the Bayesian Framework, and the VC Framework {a heavily revised version of a paper that was posted to connectionist.net about six months ago}
David H. Wolpert

Statistical Physics Models of Supervised Learning
Naftali Tishby

On Exhaustive Learning
David H. Wolpert and Alan S. Lapedes

A Study of Maximal-Coverage Learning Algorithms
Hussein Almuallim and Tom Dietterich

On Bayesian Model Selection
Peter Cheeseman

Soft Classification, a.k.a. Risk Estimation, via Penalized Log Likelihood and Smoothing Spline Analysis of Variance
Grace Wahba, Chong Gu, Yuedong Wang, and Richard Chappell

Current Research
Leo Brieman

Preface to Simplifying Neural Networks by Soft Weight Sharing
Geoffrey E. Hinton and Steven Nowlan

Simplifying Neural Networks by Soft Weight Sharing
Geoffrey E. Hinton and Steven Nowlan

Error-Correcting Output Codes: A General Method for Improving Multiclass Inductive Learning Programs
Thomas G. Dietterich and Ghulum Bakiri

Image Segmentation and Recognition
John S. Denker and Christopher C. J. Burges

==========

This book grew out of a workshop held under the auspices of the Center for Non-linear Studies at Los Alamos and the Santa Fe Institute. The idea for the workshop arose from a perception that there were many different fields that address supervised learning, but by and large these fields were not communicating with one another. (Examples of such fields are neural nets, conventional Bayesian statistics, conventional sampling theory statistics, computational learning theory, AI, and machine learning.) In particular, there were many different mathematical frameworks for addressing supervised learning. All had their own jargon, their own concerns, and their own results. And for the most part they weren’t interacting.

This was clearly a less than optimal state of affairs; we all have much to learn from one another, not only in terms of raw mathematical results, but also (perhaps more importantly) in perceptions of what the crucial issues are and how they should be addressed. Unfortunately, although it seems that this problem is abating, the rate of improvement is quite small. It seems possible that a general lack of communication amongst its practitioners will characterize supervised learning theory for some time to come.

The purpose of the workshop was try to (begin to) rectify this situation. A small group of researchers from several of the different supervised learning fields was brought together and, in effect, forced to mingle. The format of the workshop was an intensive two-day session of talks and discussion.

This volume is an attempt to try to replicate the success of the workshop in a broader context. Its purpose is to do for the reader what the workshop did for its participants: help a practitioner in one of the fields that make up supervised learning become acquainted with the relevant work by his or her colleagues in other fields.

Obviously (and unfortunately) it isn’t possible to duplicate in a reader of a book the experience of "an intensive two-day session ... (of being) forced to mingle ... (with) researchers from different fields". Given the different format, slightly different means are needed to achieve the same ends. Accordingly, it was decided that the papers in this volume should not so much be a formal compendium of the

talks presented at the workshop as an overview of the work being performed by the researchers who attended the workshop. Some of the work represented in these papers hadn't even been completed at the time of the workshop. Some of the other papers are reprints of work published shortly before or soon after the workshop. However all of the papers were chosen by their authors with the same goal in mind: to help those from other supervised learning fields get acquainted with the lay of those authors' lands. Moreover, the instructions to the authors were that they should not try to provide tutorials on their individual fields. (There are many other sources for such tutorials.) Rather they should present current cutting-edge perspectives and work that provide an intuitive understanding of what their field "is all about".

==========

The order numbers are 40985 for the hardcover and 40983 for the paperback. The prices are:

Paperback 0-201-40983-6 $31.25
Hardcover 0-201-40985-2 $59.25

It is recommended that people order through their home institutions (book stores or libraries) which may have a contract or working relationship with the publisher, Addison-Wesley. Otherwise they can call (800) 447-2226 to order by credit card.

Alternatively, they can pay by check by writing to

Advanced Book Marketing
Addison-Wesley Publishing
One Jacob Way
Reading, MA 01867, USA.