

## Introduction to the Special Issue on Computational Learning Theory

**Philip M. Long**

GISLONGP@NUS.EDU.SG

*Genome Institute of Singapore*

*1 Science Park Road*

*The Capricorn, #05-01*

*Singapore 117528, Republic of Singapore*

This special issue contains papers on Computational Learning Theory. The authors were invited to submit full papers corresponding to their preliminary versions presented at the 2001 Conference on Computational Learning Theory and the 2000 IEEE Symposium on the Foundations of Computer Science. The full papers were subjected to the usual rigorous JMLR refereeing process.

“Tracking a Small Set of Experts by Mixing Past Posteriors,” by Bousquet and Warmuth, concerns an online learning model, in which objects are encountered one at a time. In the case of two-class classification, the algorithm must predict the class of each object, and then it finds out the correct classification. The goal is to make few incorrect predictions (mistakes). Bousquet and Warmuth investigate the case in which, each time the algorithm encounters an object, it receives the predictions of a panel of experts. Previous work had shown that an algorithm can make few mistakes in online prediction if there is a way, with benefit of hindsight, to make few incorrect predictions by listening to one expert at a time, and making few changes to which expert you listen to. Bousquet and Warmuth describe online learning algorithms that are able to make few mistakes when you can (again, with benefit of hindsight), make few mistakes by switching back and forth a small number of times among an even smaller collection of good experts. Their guaranteed limit on the number of mistakes by their algorithm is close to the best possible in terms of the number of experts, the number of good experts, and the number of switches between good experts. It was strong enough to solve an open problem posed by Freund, earning the authors a \$500 reward. The techniques developed in this work have since been profitably applied in caching.

Auer’s “Using Confidence Bounds for Exploitation-Exploration Tradeoffs” also concerns online learning, except that the algorithm encounters situations one at a time, and must decide how to act in each situation in order to accumulate large reward over time. In many problems of this type, the algorithm must strike a balance between taking actions that maximize its immediate reward and taking actions that obtain information it can use to increase future rewards. In past theoretical work on the problems addressed in Auer’s paper, algorithms determined the probability of choosing different actions based on estimates of the immediate reward that would be obtained by choosing them. Auer shows that using confidence intervals on these estimates can lead to stronger theoretical performance guarantees for these problems.

Kalai and Vempala, in “Efficient Algorithms for Universal Portfolios,” provide a polynomial-time algorithm for apportioning funds among stocks that implements Cover’s UNIVERSAL

strategy. Previously known algorithms took exponential time. Cover’s strategy applies to a discrete time model in which the algorithm iteratively chooses how much to invest in each of a collection of stocks, and then finds out how their prices change. It has been shown that the UNIVERSAL strategy will accumulate a lot of wealth over a given period if there is a ‘constant rebalanced portfolio’ (determined with benefit of hindsight), that does well over that period. In a constant rebalanced portfolio, the percentage of funds invested in any given stock does not change over time.

Much of the practical success of Support Vector Machines (SVMs) is due to the application of nonlinear kernels. These kernels enable computationally efficient use of large collections of features that are functions of the raw attributes of the objects being classified. In “Limitations on Learning via Embedding in Euclidian Half Spaces,” Ben-David, Eiron, and Simon describe a variety of results that can be interpreted as suggesting that using SVMs with nonlinear kernels might not take reasonable advantage of many inductive biases strong enough to be exploited in other ways.

“Rademacher and Gaussian Complexities: Risk Bounds and Structural Results,” by Bartlett and Mendelson, concerns issues related to model selection. A recent line of research has identified alternatives to the VC-dimension that are more refined for measuring the effective complexity of a class of models that make use of the training data; these new measures can sometimes provide improved generalization estimates that can be used for model selection. After describing a number of ways in which some of the new measures are related to each other, Bartlett and Mendelson begin by establishing a guarantee on generalization in terms of the Rademacher complexity, for an abstract, decision-theoretic framework. Applying these results to model selection requires computing an appropriate complexity measure; in many applications, it is not obvious how to do this. Bartlett and Mendelson provide tools that enable one to efficiently construct estimates for many applications.

Boosting algorithms, which have been applied profitably in many domains, work by applying a subalgorithm repeatedly to a dataset in which the examples have been reweighted. Boosting increases the weight of examples which are predicted incorrectly by the subalgorithm, but it has been shown that it is useful to avoid concentrating too much weight on any given example (for example, to limit the effect of noise in the data). “On Boosting with Polynomially Bounded Distributions,” by Bshouty and Gavinsky, describes a general way in which boosting algorithms can be transformed to avoid concentrating too much weight on any example, while preserving their ability to boost.

I wish to thank the referees for their excellent work, Editor-in-Chief Leslie Kaelbling and Managing Editor David Cohn of JMLR for their guidance and help, and most of all, the authors for contributing their papers.