

## Introduction to the special issue on COLT 2006

Avrim Blum · Gabor Lugosi · Hans Ulrich Simon

Published online: 28 September 2007  
Springer Science+Business Media, LLC 2007

This special issue of *Machine Learning* is dedicated to the Nineteenth Annual Conference on Learning Theory (COLT 2006) held at Pittsburgh, PA, USA, June 22–25, 2006. The papers in this issue were selected from those presented at that conference, and the authors invited to submit completed versions of their work. Once received, these papers underwent the usual refereeing process of *Machine Learning*.

The field of Learning Theory provides a mathematical foundation for the study of machine learning. The analysis proceeds in a formal model so as to provide measures for the performance of a learning algorithm or for the inherent hardness of a given learning problem. The variety of applications for algorithms that learn is reflected in the variety of formal learning models. For instance, we can distinguish between a passive model of “learning from examples” and active models of learning where the algorithm has more control over the information that is gathered. As for learning from examples, a further decision is whether or not we impose statistical assumptions on the sequence of examples. Furthermore, we find different success criteria in different models (like “approximate learning” versus “exact learning”).

The papers in this special issue offer a broad view on the current research in the field including studies on several learning models (such as Teaching, Statistical Query Learning,

---

The conference proceedings, including preliminary versions of these papers, appeared as Lecture Notes in Artificial Intelligence, vol. 4005, Springer, 2006.

A. Blum (✉)  
Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3891, USA  
e-mail: avrim@cs.cmu.edu

G. Lugosi  
ICREA and Department of Economics, Pompeu Fabra University, Ramon Trias Fargas 25–27, 08005  
Barcelona, Spain  
e-mail: gabor.lugosi@gmail.com

H.U. Simon  
Fakultät für Mathematik, Ruhr-Universität Bochum, NA 1/73, 44780 Bochum, Germany  
e-mail: hans.simon@rub.de

Online Prediction, Reinforcement Learning, and Multiple Output Identification). Below we briefly introduce each of the papers and on the way provide some background information about their respective underlying learning models.

The model of *Teaching* studies the minimum number of well-chosen examples needed to ensure that any learning algorithm that produces a consistent function from the given class will in fact output the correct target function. The paper “DNF are Teachable in the Average Case” by Lee, Servedio, and Wan considers this problem for the class of DNF formulas having at most  $s$  terms. Even though it is known that there exist *some* functions in this class that require an exponential number of examples to uniquely identify in this way, they show that *most* functions in the class require a teaching set of size only  $O(ns)$ . They also study this average-case notion of teaching dimension for several other classes.

The Statistical Query (SQ) model is an elegant abstraction of Valiant’s PAC learning model. In this model, instead of having direct access to random examples (as in the PAC model), the learner obtains information about random examples via an oracle that provides estimates of various statistics about the unknown concept. Any learning algorithm that is successful in the SQ model can be converted, without much loss of efficiency, into a learning algorithm that is successful in the PAC model even in the presence of noise uniformly applied to the class labels of the examples. Furthermore, most efficient PAC learning algorithms can be easily redesigned so as to work in the SQ model. The paper “Unconditional Lower Bounds for Learning Intersections of Halfspaces” by Klivans and Sherstov is concerned with limitations on SQ learning. For example, they show that the intersection of halfspaces requires exponentially many queries. The key to their analysis is the observation that there exist (so-called) Bent-functions which can be cast as an intersection of (a small number of) halfspaces. This paper has some interesting links to questions in cryptography and complexity theory.

In most models concerned with online prediction of individual data sequences, no assumptions (stochastic or whatsoever) are made about the generation of data, and prediction algorithms are evaluated by means of competitive analysis. In this context, regret bounds are used to guarantee that the algorithm does not perform much worse than the best predictor in a given comparison class. The paper “A Primal-Dual Perspective of Online Learning Algorithms” by Shalev-Shwartz and Singer presents a novel framework for the design and analysis of online learning algorithms. They consider the dual of the loss-minimization problem and use the amount by which the dual increases as a measure of progress. This leads to new algorithms whose design is inspired by the duality approach and sheds also some new light on existing algorithms. In the paper “Tracking the Best Hyperplane with a Simple Budget Perceptron”, Cavallanti, Cesa-Bianchi, and Gentile evaluate the classical Perceptron algorithm with respect to two combined objectives. First, they would like to make the algorithm robust in situations where the optimal classifier changes over time. Second, they impose a bound on the number of support vectors that the Perceptron is able to keep in memory simultaneously. The authors find an elegant and successful way of combining these two objectives. The paper “Logarithmic Regret Algorithms for Online Convex Optimization” by Hazan, Agarwal, and Kale considers the online convex optimization problem, a fairly general online decision-making problem in which at each time step  $t$ , the algorithm must choose a point  $x_t$  inside a given convex set and then is presented with a convex cost function  $f_t$ , paying cost  $f_t(x_t)$ . The goal is to perform nearly as well as the best fixed point  $x$  in the set in hindsight. This paper presents a number of algorithms that are able to achieve particularly strong regret bounds when the cost functions exhibit at least certain limited forms of curvature. Such cost functions include those appearing in the universal portfolio problem, among others. The paper “Competing with Wild Prediction Rules” by Vovk is

concerned with comparison classes that contain highly irregular prediction rules such as in benchmark classes that cannot be cast as reproducing kernel Hilbert spaces. Vovk develops Banach-space methods to construct a prediction algorithm with good regret bounds in this quite general setting.

In the paper “Active Sampling for Multiple Output Identification” by Fine and Mansour, the authors propose a simplified version of the active learning model, in which the goal is simply to find at least one example of each possible label value. This setting is motivated by problems such as hardware verification and software testing. The paper studies this model for a variety of function classes, and presents a separation dimension that approximately characterizes the number of samples needed to succeed in this setting for the case of binary output values.

In the *Reinforcement Learning* model, the action of a learner leads to a state-transition and to a (state-dependent) immediate reward. The state transition function is unknown to the learner who aims at maximizing the long-term (possibly discounted) cumulative reward. This model is quite general and learning policies are often evaluated empirically or asymptotically only. The paper “Learning Near-optimal Policies with Bellman-residual Minimization Based Fitted Policy Iteration and a Single Sample Path” by Antos, Szepesvari, and Munos presents a near-optimal policy and a finite-sample bound on its performance in the setting with a continuous state-space where the trajectory of some behavior policy is given.<sup>1</sup>

We would like to thank all the referees for their efficient work and thorough reports. Special thanks go to the members of the program committee of COLT 2006 who helped us to select a sample of excellent papers that nicely represent the richness of and the diversity within the field of learning theory. We are very grateful to all authors for submitting their papers and for their work in improving and polishing their articles. We would like furthermore to thank the editorial staff of Machine Learning for their help. Finally, we are particularly grateful to the Editor-in-Chief, Foster Provost, for the opportunity to compile this special issue.

Pittsburgh, Barcelona, and Bochum, September 12, 2007

Special Issue Editors

---

<sup>1</sup>Due to time limitations, this paper will appear in a subsequent issue.