Machine Learning 35, 191–192 (1999)

Introducing the Special Issue of Machine Learning Selected from Papers Presented at the 1997 Conference on Computational Learning Theory, COLT'97

JOHN SHAWE-TAYLOR
Royal Holloway, University of London

The aim of the COLT special issues is to provide a forum for COLT conference papers which though theoretical have been judged to have significant practical ramifications and hence to be of potential interest to a broader spectrum of Machine Learning readers.

The papers included in this special issue cover a broad range of topics, their practical significance being in some cases of the form 'how to do it', while in others it is of the form 'this could be hard'. The paper most clearly aiming at a particular practical application is by Blum and Kalai entitled, 'Universal portfolios with and without transaction costs'. The paper considers the problem of rebalancing an investment portfolio on-line in such a way that the result is competitive with the best constant rebalancing determined in hindsight. They extend Cover's universal algorithm to take into account transaction costs, and at the same time give a randomized implementation which is significantly faster in practice. Other applications of the approach are also described.

Dalmau's paper 'A dichotomy theorem for learning quantified Boolean formulas' is in the 'this could be hard' category. It studies the question of learning quantified Boolean formulas constructed from a finite set of basic functions using conjunctions, constants and finally existential or universal quantification. He shows how the sets of basic functions can be identified as being in one of two categories. For sets in the first category learning is 'easy' and can be achieved by known algorithms and techniques, while for sets in the second category learning is not possible under standard cryptographic assumptions.

The paper 'Estimation of time-varying parameters in statistical models; an optimization approach' by Bertsimas, Gamarnik and Tsitiklis, studies an approach to nonparametric regression, reducing the problem to convex programming, an optimization problem that is efficiently solvable. They show strong convergence results for their approach assuming only that the underlying regression function is Lipzchitz continuous. Hence, the contribution of the paper is both theoretical and practical, while at the same time providing a link with a more statistical view of learning.

The final paper included in the special issue is entitled, 'Derandomizing stochastic prediction strategies', by Vovk. This paper describes a new approach to analysing the use of expert advice. The approach taken here is that, rather than having a pool of experts (as for example in Blum and Kalai's paper), there is one stochastic expert. Adopting this view makes it possible to rework previous analysis in the area as well as obtaining new

results on tracking the best expert. Applications of the technique can be used when merging overconfident experts and also in fitting polynomials.

While in many cases COLT papers do not concentrate on developing or demonstrating efficient algorithms on large data sets, I believe that the papers presented here do contain novel ideas and approaches which if translated into practical applications could give improved performance. I, therefore, commend the papers in this issue as original contributions to our understanding of the problems of Machine Learning.