Guest Editor's Introduction

LISA HELLERSTEIN

Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208-3118

This special issue of Machine Learning contains four selected papers presented at the Fifth Annual ACM Workshop on Computational Learning Theory (COLT), held in Pittsburgh, Pennsylvania, July 27-29, 1992.

Research in computational learning theory covers a diverse range of topics, and the papers in this special issue reflect that diversity. There is something in these papers for (practically) everyone — foundational models, algorithms, applications, logics, and memory systems. In what follows, I will discuss the contents of these papers and their relation to broader research issues in computational learning theory, and provide references for some recent related research. Readers unfamiliar with the field of computational learning theory may also be interested in some background reading; possibilities include three textbooks (Kearns & Vazirani, 1994; Anthony & Biggs, 1992; Natarajan, 1991) and two survey articles written by Angluin (1992, 1993). There are also other special issues of Machine Learning devoted to previous COLT conferences (Li & Valiant, 1994; Blumer & Case, 1992; Pitt, 1990).

One of the central problems in computational learning theory is the development and investigation of formal learning models. The goal of a formal learning model is to capture essential properties of actual learning methods, and to enable formal analysis and the development of learning algorithms. The PAC (probably approximately correct) model of learning, introduced by Valiant (1984), was designed to achieve that goal. The extensive research generated by the model attests to its degree of success. However, many researchers have pointed to deficiencies in the PAC model, and some have introduced variant models designed to correct these deficiencies.

The first paper in this special issue, by M. J. Kearns, R. E. Schapire, and L. M. Sellie, introduces a generalization of the PAC model called the *agnostic learning model*. The PAC model makes strong assumptions about the form of the function being learned (the *target function*). However, these assumptions are not always justifiable in practice. In much of empirical machine learning research, the learning algorithm makes virtually no assumptions about the target function. The algorithm is, however, constrained to output a hypothesis from a restricted class of functions; its object is to find the “best” approximation to the target function within the hypothesis class. The agnostic learning model captures these features of empirical machine learning research, while at the same time including the PAC model as a special case.

Kearns, Schapire, and Sellie prove a number of negative results showing the difficulty of agnostic learning, including an equivalence between agnostic learning and PAC learning with malicious errors (a particularly difficult noise model). They also show that agnostic learning is difficult even when the hypothesis class is the simple class of boolean conjunctions. In contrast to these negative results, they show that a number of non-trivial

problems are tractable in the agnostic learning model. For example, they give a dynamic programming method that can be used to learn agnostically when the hypothesis class consists of piecewise polynomials. The research in this paper is an initial exploration of the problems of agnostic learning; many open problems remain. Recently, W. Maass has proved results on agnostic learning with analog neural nets, and with simple hypotheses such as rectangles (Maass, 1994a, 1994b).

The basic ideas of the PAC model, and the fundamental research in the model, continue to have a broad impact in the formal study of machine learning. The second paper in this issue, "A Theory for Memory-Based Learning," by J-H Lin and J. S. Vitter, uses the framework of the PAC model to unify and analyze methods of learning smooth functions such as nearest-neighbor search, space decomposition, and clustering. These methods are termed "memory-based" because they store the values of the function at representative points, and use the stored values to estimate the value of the function at other points. Such methods have been used in applications such as the real time control of robots. However, there have been few theoretical results regarding the use of such methods for this purpose. Using techniques from the theory of PAC learning, Lin and Vitter give precise answers to questions such as the effect of memory size in learning smooth functions in memory-based systems.

The third paper in this issue bridges the fields of computational learning theory and knowledge representation. In "Learnability of Description Logics with Equality Constraints," W. W. Cohen and H. Hirsh study the problem of PAC learning a restricted class of first order logics known as description logics. More specifically, they explore the learnability of a description logic called CORECLASSIC. This logic is a subclass of CLASSIC, which is used in practical knowledge representation systems. Cohen and Hirsh prove that (under standard complexity assumptions), CORECLASSIC cannot be PAC-learned. They then give a PAC learning algorithm for learning a non-trivial restriction of the CORECLASSIC language. Building on the work in this paper, Cohen and Hirsh have since proved additional related results (Cohen & Hirsh, 1994), and Frazier and Pitt have shown the learnability of all of CLASSIC in a more powerful learning model (Frazier & Pitt, 1994). This theoretical work is a welcome addition to the extensive experimental work on learning first order logics. It also complements theoretical work on PAC learning subsets of Prolog, which have a different syntax than decision logics (Haussler, 1989; Frisch & Page, 1991; Džeroski, et al. 1992).

The PAC learning model and its variants (including agnostic learning) model learning from random data. In query models of learning, the learning algorithm is allowed to ask questions concerning the target function. These questions model learning strategies such as experimentation, hypothesis testing, and asking questions of a teacher or expert.

The final paper in this issue, "On-line learning of rectangles and unions of rectangles," by Z. Chen and W. Maass, gives an algorithm for learning d -dimensional rectangles $\prod_{i=1}^d \{a_i, a_i + 1, \dots, b_i\}$ over a discrete domain $\{1, \dots, n\}^d$. That is, their algorithm assumes that the target function f has the output value 1 on points inside such a rectangle, and the output 0 on points outside the rectangle. The algorithm is designed for the model of "on-line learning with equivalence queries". In this model, the learning algorithm proposes a sequence of hypotheses. If a proposed hypotheses h is not equal to the target

function f , the learning algorithm is given a counterexample x — a point in the domain such that $f(x) \neq h(x)$. The algorithm can use this information to construct its next hypothesis. The goal of the algorithm is to construct a hypothesis that is equivalent to the target function. This model is also known as the *equivalence query* model (Angluin, 1987), and is closely related to the *mistake-bound* model (Littlestone, 1988).

The algorithm of Chen and Maass proposes at most $O(d^2 \log n)$ incorrect hypotheses. It solves a basic problem while at the same time introducing novel algorithmic techniques. In particular, it incorporates an unusual solution to the credit assignment problem, and a new algorithm for performing binary search in the presence of errors. The paper also includes an algorithm for learning the union of two rectangles in the plane. Since the paper first appeared, a number of other papers have proved related results (Auer, 1993; Chen, 1993; Chen & Homer, 1993a, 1993b; Cohen & Hirsh, 1994; Goldberg, Goldman, & Mathias, 1994; Bshouty, 1994; Maass & Warmuth, 1994).

The papers in this special issue represent only a small sample of the many excellent papers presented at the 1992 COLT conference. I therefore encourage readers to obtain a copy of the conference proceedings (ACM Press, 1992). In the meantime, I hope that the papers included here will stimulate interest and lead to future research. Read, and enjoy!

References

- ACM Press (1992). *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*. Pittsburgh, PA: ACM Press.
- Angluin, D. (1987). Queries and concept learning. *Machine Learning*, 2, 319–342.
- Angluin, D. (1992). Computational learning theory: Survey and selected bibliography. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on the Theory of Computing* (pp. 351–369). Victoria, BC: ACM Press.
- Angluin, D. (1993). Learning with Queries. In E. B. Baum, (Ed.), *Computational Learning & Cognition: Proceedings of the 3rd NEC Research Symposium*. Proceedings in Applied Mathematics, vol. 64 (pp. 1–28). Philadelphia: Society for Industrial and Applied Mathematics.
- Anthony, M. & Biggs, N. (1992). *Computational Learning Theory: An Introduction*. Cambridge, England: Cambridge University Press.
- Auer, P. (1993). On-line learning of rectangles in noisy environments. In *Proceedings of the Sixth Annual ACM conference on Computational Learning Theory* (pp. 253–261). Santa Cruz, CA: ACM Press.
- Blumer, A. & Case, J. (eds). (1992). *Machine Learning, Special Issue on Computational Learning Theory* 9(2/3).
- Bshouty, N. H. (1994). On Learning Discretized Geometric Concepts. Unpublished manuscript.
- Chen, Z. (1993). Learning unions of two rectangles in the plane with equivalence queries. In *Proceedings of the Sixth Annual ACM conference on Computational Learning Theory* (pp. 243–252). Santa Cruz, CA: ACM Press.
- Chen, Z. & Homer, S. (1993a). Fast learning unions of rectangles on the plane with queries. Computer Science Department Technical Report, Boston University, BUCS-TR-10-93.
- Chen, Z. & Homer, S. (1993b). The bounded injury priority method and the learnability of unions of rectangles. Unpublished manuscript.
- Cohen, W. W. & Hirsh, H. (1994). Learning the CLASSIC Description Logic: Theoretical and Experimental Results. In J. Doyle, E. Sandewall, & P. Torasso (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR94)* (pp. 121–133). Bonn, Germany: Morgan Kaufmann.

- Džeroski, S., Muggleton, S., & Russell, S. (1992). Pac-learnability of determinate logic programs. In *Proceedings of the 1992 Workshop on Computational Learning Theory* (pp. 128-135). Pittsburgh, PA: ACM Press.
- Frazier, M. & Pitt, L. (1994). CLASSIC learning. In *Proceedings of the 1994 Workshop on Computational Learning Theory* (pp. 23-34). New Brunswick, NJ: ACM Press.
- Frisch, A. & Page, C. D. (1991). Learning constrained atoms. In *Proceedings of the Eighth International Workshop on Machine Learning* (pp. 427-431). Ithaca, New York: Morgan Kaufmann.
- Goldberg, P. W., Goldman, S. A. & Mathias H. D. (1994). Learning unions of boxes with membership and equivalence queries. In *Proceedings of the 1994 Workshop on Computational Learning Theory*. New Brunswick, NJ: ACM Press.
- Haussler, D. (1989). Learning conjunctive concepts in structural domains. *Machine Learning*, 4, 7-40.
- Kearns, M. J. & Vazirani, U. V. (1994). *An Introduction to Computational Learning Theory*. Cambridge, MA: The MIT Press. To be published.
- Li, M. & Valiant, L. G. (eds). (1994). *Machine Learning, Special Issue on Computational Learning Theory* 14(1).
- Littlestone, N. (1988). Learning quickly when irrelevant attributes abound: a new linear-threshold algorithm. *Machine Learning*, 2, 285-318.
- Maass, W. (1994a). Agnostic PAC-learning of functions on analog neural nets. In *Advances in Neural Information Processing Systems vol. 6*. San Mateo, CA: Morgan Kaufmann. *Neural Computation*, 6(5), 877-884.
- Maass, W. (1994b). Efficient agnostic PAC-learning with simple hypotheses. In *Proceedings of the 1994 Workshop on Computational Learning Theory* (pp. 67-75). New Brunswick, NJ: ACM Press.
- Maass, W. & Warmuth, M. K. (1994). Efficient learning with virtual threshold gates. Manuscript in preparation.
- Natarajan, B. K. (1991). *Machine Learning: A Theoretical Approach*. San Mateo, CA: Morgan Kaufmann.
- Pitt, L. (ed). (1990). *Machine Learning, Special Issue on Computational Learning Theory* 5(2).
- Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, 27, 1134-1142.