GUEST EDITORIAL

New Theoretical Directions in Machine Learning

Looking back over the last decade of work in machine learning, one thing is apparent: there is an art to designing efficient and effective learning algorithms. Too much of an art. Many researchers have argued that if the field is ever to fully mature, we must find a sound and adequate theoretical basis on which to build a science of machine learning.

A learning algorithm is an algorithm. Although it would be unwise to subsume machine learning entirely under the general theory of algorithms and computational complexity, the techniques and perspectives from this field are pertinent to developing a theoretical basis for machine learning. This special issue of *Machine Learning* presents some recent progress in applying these techniques and perspectives to the computational problems inherent in learning.

Early efforts along these lines were based primarily on the framework introduced by Gold (1967). This early work is reviewed in the excellent survey by Angluin and Smith (1983). More recently, Valiant (1984) has introduced a probabilistic framework for the study of learning algorithms. Kearns, Li, Pitt, and Valiant (1987) and Haussler (1987) have surveyed applications of this framework to AI concept-learning problems; Rivest (1987) has also carried out recent work in this paradigm.

The Valiant framework has been more successful than the Gold framework in addressing some of the requirements that are typically placed on a learning algorithm in practice. To allow greater emphasis on computational efficiency, it requires only that a good approximation to the target concept be found with high probability, rather than requiring exact identification of the target concept. Recent results have demonstrated that there are efficient learning algorithms that achieve this type of “probably approximately correct” identification (see Angluin, this issue) for many types of target concepts (Kearns et al., 1987; Haussler, 1987).

Although this framework has proven quite powerful, in its present state it still fails to address a number of important issues, including:

- (1) the problem of designing efficient incremental learning algorithms;
- (2) the relationship between probabilistic and worst-case performance;
- (3) the role of learner-generated queries in concept identification;¹
- (4) the problem of noisy training examples; and
- (5) the problem of unsupervised learning.

Each of the research papers in the following pages addresses one or more of these problems.

Littlestone's article deals with the first issue above. He introduces a framework for incremental learning in which the primary performance criterion is based on the total number of mistakes made during learning, computed in the worst case over all possible ordered presentations of the instances. He uses results by Angluin (this issue) to relate this analysis to Valiant's criterion. Littlestone also presents a new variant of the classical perceptron learning algorithm, showing how it can be adapted to learn a variety of basic concept types with a nearly minimal number of mistakes. The algorithms he gives represent dramatic improvements over previous incremental algorithms for simple (i.e., syntactically short) target concepts on instance spaces with many attributes. Like the classical perceptron algorithm, Littlestone's method is extremely efficient and well suited for implementation in a connectionist architecture.

Angluin's paper deals with issues (2) and (3) in the above list. She considers learning algorithms that obtain information about the target concept by making queries. Among the types of queries she considers are questions of the form "is x an instance of the target concept?", where x is any instance in the instance space,² and "is h equivalent to the target concept?", where h is any hypothesis in the hypothesis space. She calls the former a *membership* query and the latter an *equivalence* query.

Angluin's performance criterion is based on the number of queries needed to exactly identify the target concept. In the case of equivalence queries, this criterion is essentially the same as the one used by Littlestone, and Angluin also relates it to Valiant's probabilistic criterion. Along with her analytical results, she gives query algorithms for identifying a wide variety of different types of target concepts.

¹This is an aspect of Valiant's (1984) original framework that, until now, has received little attention.

²Sammur and Banerji (1986) and Subramanian and Feigenbaum (1986) have also studied this type of query.

The paper by Angluin and Laird deals with the fourth issue: noisy training examples. The authors extend the Valiant framework by assuming that each independent random example of the target concept is misclassified with probability η , where $0 \leq \eta < 1/2$. First they take a general look at the strategy of minimizing the number of disagreements between the hypothesis and the training sample. Based on this analysis, they then develop a polynomial-time algorithm for learning k -CNF target concepts (Valiant, 1984) from noisy random examples that can tolerate a noise level η approaching $1/2$. They compare these results with those in Valiant (1985) and Kearns and Li (1987), obtained under the assumption that training examples are maliciously modified by an adversary of the learning algorithm.

Finally, Pitt and Reinke consider the issue of unsupervised learning. They develop a formalism for representing and analyzing conceptual clustering algorithms and demonstrate it with a number of examples. Their criterion for success is the ability to efficiently produce a clustering of the instances that maximizes a given objective function, which is based on individual cluster tightness and overall distance between clusters. In their main theorem, they show that a simple variant of the standard agglomerative algorithm for hierarchical clustering will succeed for a wide variety of objective functions.

In dealing with the issues identified above, the present authors have all moved outside the original Valiant framework in different ways. Hence, despite these and other recent advances, it seems clear that the science of machine learning is still far from the comprehensive theoretical foundation that it needs. We hope the papers in this issue of *Machine Learning* will inspire others to join in the further development of this foundation.

David Haussler
University of California, Santa Cruz
HAUSSLER@SATURN.UCSC.EDU

References

- Angluin, D., & Smith C. (1983). Inductive inference: Theory and methods. *Computing Surveys*, 15, 237–269.
- Gold, E. M. (1967). Language identification in the limit. *Information and Control*, 10, 447–474.
- Haussler, D. (1987). Bias, version spaces and Valiant's learning framework. *Proceedings of the Fourth International Workshop on Machine Learning* (pp. 324–336). Irvine, CA: Morgan Kaufmann.

- Kearns, M., & Li, M. (1987). *Learning in the presence of malicious errors* (Technical Report TR-03-87). Cambridge, MA: Harvard University, Aiken Computation Laboratory.
- Kearns, M., Li, M., Pitt, L., & Valiant, L. G. (1987). Recent results on Boolean concept learning. *Proceedings of the Fourth International Workshop on Machine Learning* (pp. 337–352). Irvine, CA: Morgan Kaufmann.
- Rivest, R. L. (1987). Learning decision lists. *Machine Learning*, 2, 229–246.
- Sammut, C., & Banerji, R. (1986). Learning concepts by asking questions. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2). Los Altos, CA: Morgan Kaufmann.
- Subramanian, D., & Feigenbaum, J. (1986). Factorization in experiment generation. *Proceedings of the Fifth National Conference on Artificial Intelligence* (pp. 518–522). Philadelphia, PA: Morgan Kaufmann.
- Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, 27, 1134–1142.
- Valiant, L. G. (1985). Learning disjunctions of conjunctions. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 560–566). Los Angeles, CA: Morgan Kaufmann.