EDITORIAL

Exploratory Research in Machine Learning

Exploratory research contributes to the continued vitality of every discipline. The aim of exploratory research is to identify new tasks—tasks that cannot be solved by existing methods. Once a new task has been found, exploratory research seeks to develop a precise definition of the task and to understand the factors that make the task different from previously-solved tasks.

Until recently, most research in machine learning was primarily exploratory. However, during the past decade, some areas of the field—particularly inductive learning—have matured to the point that careful, quantitative experiments are now possible and proved theoretical results have been obtained. Although these trends are extremely healthy and long overdue, there is a danger that the increased attention to these products of mature research may discourage researchers from undertaking and publishing research of a more exploratory nature. The goal of this editorial is to emphasize the importance of exploratory research and to encourage the publication of high quality exploratory results in *Machine Learning*.

A Model of the Research Process

To appreciate the role of exploratory research, it is helpful to examine the various phases of the research process. We begin by defining each of the phases and indicating the kinds of results obtained in each phase.

Research begins with a phase of exploration, usually driven by specific problems in specific domains. For example, the early research in inductive concept learning was motivated by neural modeling (Rosenblatt, 1957), psychological modeling (Hunt, Marin & Stone, 1966), and minimization of switching circuits (Michalski, 1969) to name only a few of the pioneering projects. Exploratory research seeks to understand a new problem and develop a precise task definition. In the area of inductive concept learning, a clear definition of the task eventually emerged in the work of Mitchell (1978). Subsequent exploratory research has produced less constrained and more realistic definitions (e.g., Valiant, 1984).

Exploration is usually followed by (or associated with) a phase in which algorithms and methods are developed. In addition to the methods developed from the exploratory research mentioned above (perceptrons, CLS, and A^q), the past decade has seen an explosion in the number of learning methods (such as ID3 (Quinlan, 1983), PLS (Rendell, 1983), Back-propagation (Rumelhart, Hinton, & Williams, 1986), Stagger (Schlimmer, 1987), etc.). The primary research result in this phase is simply the development of a new method.

The goal of the third phase of research is to perform empirical evaluation of the methods that have been developed. The first empirical evaluations of machine learning algorithms were performed using the training set/test set methodology by Hunt, Marin, & Stone (1966). Subsequent evaluations focused on additional criteria such as the number of examples required to obtain good performance (Quinlan, 1983), noise tolerance (Quinlan, 1986), concept drift (Schlimmer & Granger, 1986), and incremental learning (Schlimmer & Fisher, 1986). In this third phase, methods are typically evaluated in isolation, although researchers are encouraged to proceed to comparative studies (phase four) as quickly as possible.

Once more than one method is available, the fourth phase of research can begin. In this phase, algorithms are compared against one another under controlled situations (e.g., same data, same representation, etc). The goal of the research is to determine which algorithms are better (and under what conditions). Inductive concept learning research is actively proceeding through this phase, as demonstrated in the papers by Fisher (1987), Schlimmer & Fisher (1986), Utgoff (1988), Quinlan (1988), Mingers (1989), Mooney, Shavlik, Towell, & Gove (1989), Weiss & Kapouleas (1989), Fisher & McKusick (1989).

A fifth phase of research, which may proceed independently of phases three and four, is the theoretical analysis of the task and the methods. For inductive concept learning, there has been an explosion of activity in the past five years, as represented by Valiant (1984), Blumer, Ehrenfeucht, Haussler, & Warmuth (1987), Ehrenfeucht, Haussler, Kearns, & Valiant (1988), Judd (1987), Pitt & Warmuth (1988), Kearns & Valiant (1988), and many others. The goal of this research is to develop upper and lower bounds on the performance of *any* algorithm. Another goal is to analyze specific algorithms to determine their relationship to these upper and lower bounds. Theoretical analysis must also determine the sensitivity of its results to slight changes in the task definition.

A sixth phase of research attempts to tie together all that is known about a problem and provide a theoretical foundation for the field. One form that this can take is a generative theory of methods that tells how to design a good method for any particular problem. In the inductive concept learning area, the minimum description length principle (Rissanen, 1978) and its Bayesian justification shows some promise of providing such a foundation, although much research remains to be done in this area.

The Role of *Machine Learning*

Machine Learning has the responsibility to encourage and publish high quality research in all of these phases. Recently, editorials in *Machine Learning* have attempted to articulate good techniques for the empirical evaluation of machine learning methods (i.e., phases three and four) and the development of good theoretical work (i.e., phase five). Nearly all of the manuscripts received by the journal fall into these three phases. This is appropriate, and I expect it to continue in the future. However, I would like to give some attention to the other phases, especially exploratory research.

An exploratory paper should include the following: (a) a precise statement of a new learning problem or learning situation, (b) a justification for why this is a new problem rather than being only superficially different, (c) a discussion of the feasibility of solving the problem, (d) a description of the issues that are believed to be important in attacking this new problem (e.g., tradeoffs, important variables that should be controlled, etc.), and (e) a suggested agenda for future research in the area. Each of these points should be justified by appealing to properties of some particular domain or domains. To demonstrate the novelty of the task, for example, one good approach is to show that existing methods fail (or are inapplicable). To show that it is feasible to solve the problem, a convincing strategy is to propose and implement a method that solves the problem. Without such justification and illustration, the paper is unlikely to pursuade the readers that this is an important kind of new problem or that it is a problem where research progress can be made.

An exploratory research paper should enable (and encourage) readers to pursue the research items on the agenda of part (e). To achieve this, the description of the research problem must be so clear that any reader could apply the description to decide whether another problem in a new domain was an instance of the research problem in question. It should be possible for the reader to replicate and explore, in a new domain, the problems and issues described.

Three good examples of exploratory research papers are (a) Amarel's (1968) paper on reformulation in the missionaries-and-cannibals problem, (b) the first paper describing LEX (Mitchell, Utgoff, & Banerji, 1983), and (c) the first LEAP paper (Mitchell, Mahadevan, & Steinberg, 1985). Each of these papers described either hand examples or partial implementations of systems. None of the papers provided empirical verification of the kind required in later research phases. But each clearly described a new kind of learning situation and analyzed the important issues and open research questions in the area. Each paper was tightly coupled to a domain, which provided the motivating examples for the claims in the paper.

One critical feature that each of the three papers share is exceptionally clear and convincing writing. High quality writing is even more important in exploratory papers than it is in papers describing other phases of research. This is because the material in exploratory papers necessarily involves new domains and new ways of looking at problems. Consequently, it is less familiar to readers and tends to fall outside established paradigms.

There are many areas of the machine learning field that are ripe for exploration and that would benefit from the publication of good exploratory research papers. Consider, for example, the area of knowledge integration, which studies the problem of how to integrate new knowledge into a large existing knowledge base quickly and easily. There is a great need for a precise definition of this task and a methodology for evaluating proposed knowledge integration methods.

Another area in which exploratory research is proceeding is the area of learning robots. What are the learning tasks facing a robot? How can we measure our progress as we attack these tasks? A paper answering these questions would be a valuable contribution to the field.

Concluding Remarks

According to the research model described above, a research field cannot remain vital unless it devotes some portion of its energies to exploratory research. It is this research that identifies new, important problems to attack. Without exploratory research, it is easy for a field to degenerate into "algorithm polishing" activities that yield only modest improvements in performance and stick within the safe boundaries of established paradigms.

This editorial has described the makeup of a high-quality exploratory research paper. I encourage authors to consider these criteria and to submit exploratory research papers to *Machine Learning.*

Thomas G. Dietterich
Editor
Department of Computer Science
Oregon State University

References

Amarel, S. (1968). On the representation of problems of reasoning about actions. In D. Michie (Ed.), *Machine Intelligence* (Vol. 3). Edinburgh: U. of Edinburgh Press.

Blumer, A., Ehrenfeucht, A., Haussler, D., & Warmuth, M. K. (1987). Occam's razor. *Information Processing Letters, 24,* 377–380.

Ehrenfeucht, A., Haussler, D., Kearns, M., & Valiant, L. (1988). A general lower bound on the number of examples needed for learning. *COLT 88: Proceedings of the Conference on Learning Theory* (pp. 110–120). Cambridge, MA: Morgan Kaufmann.

Fisher, D. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning, 2,* 139–172.

Fisher, D. H., & McKusick, K. B. (1988). An empirical comparison of ID3 and back-propagation. *IJCAI-89: Eleventh International Joint Conference on Artificial Intelligence* (pp. 788–793). Detroit, MI: Morgan Kaufmann.

Hunt, E. B., Marin, J., & Stone, P. J. (1966). *Experiments in induction.* New York: Academic Press.

Judd, J. S. (1987). Learning in networks is hard. *Proceedings of the First International Conference on Neural Networks* (pp. 685–592). San Diego, CA: IEEE.

Kearns, M., & Valiant, L. G. (1988). *Learning Boolean formulae or finite automata is as hard as factoring* (Technical Report No. 14-88). Cambridge, MA: Harvard University, Aiken Computation Laboratory.

Michalski, R. S. (1969). On the quasi-minimal solution of the general covering problem. *Proceedings of the Fifth International Federation on Automatic Control, 27,* 109–129.

Mingers, J. (1989). An empirical comparison of selection measures for decision-tree induction. *Machine Learning, 3,* 319–42.

Mitchell, T. M. (1978). *Version Spaces: An approach to concept learning* (Technical Report No. STAN-CS-78-711). Stanford, CA: Stanford University, Department of Computer Science.

Mitchell, T.M., Mahadevan, S., & Steinberg, L. I. (1985). LEAP: A learning apprentice for VLSI design, *IJCAI-85: Ninth International Joint Conference on Artificial Intelligence* (pp. 573–580). Los Angeles, CA: Morgan Kaufmann.

Mitchell, T. M., Utgoff, P. E., & Banerji, R. B. (1983). Learning by experimentation: Acquiring and refining problem-solving heuristics. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 1). San Mateo, CA: Morgan Kaufmann.

Mooney, R., Shavlik, J., Towell, G., & Gove, A. (1989). An experimental comparison of symbolic and connectionist learning algorithms. *IJCAI-89: Eleventh International Joint Conference on Artificial Intelligence* (pp. 775–780). Detroit, MI: Morgan Kaufmann.

Pitt, L., & Warmuth, M. K. (1989). The minimum DFA consistency problem cannot be approximated within any polynomial. *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing* (pp. 421–432). ACM.

Quinlan, J. R. (1983). Learning efficient classification procedures and their application to chess endgames. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 1). San Mateo, CA: Morgan Kaufmann.

Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning, 1,* 81–106.

Quinlan, J. R. (1988). An empirical comparison of genetic and decision-tree classifiers. *Proceedings of the Fifth International Conference on Machine Learning* (pp. 135–141). Ann Arbor, MI: Morgan Kaufmann.

Rendell, L. (1983). A new basis for state-space learning systems and a successful implementation. *Artificial Intelligence, 20,* 369–92.

Rissanen, J. (1978). Modeling by shortest data description. *Automatica, 14,* 465–471.

Rosenblatt, F. (1957). *The perceptron: A perceiving and recognizing automaton* (Technical Report No. 85-460-1). Ithaca, NY: Project PARA, Cornell Aeronautical Laboratory.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, & McClelland, J. L. (Eds.), *Parallel Distributed Processing* (Vol. 1). Cambridge, MA: MIT Press.

Schlimmer, J. C., & Fisher, D. (1986). A case study in incremental concept formation. *Proceedings of the National Conference on Artificial Intelligence, AAAI-86* (pp. 496–501). Philadelphia, PA: Morgan Kaufmann.

Schlimmer, J. C., & Granger, R. H. Jr. (1986). Beyond incremental processing: Tracking concept drift. *Proceedings of the National Conference on Artificial Intelligence, AAAI-86* (pp. 502–507). Philadelphia, PA: Morgan Kaufmann.

Schlimmer, J. C., & Granger, R. H. Jr. (1986). Incremental learning from noisy data. *Machine Learning, 1,* 317–354.

Utgoff, P. E. (1988). ID5: An incremental ID3. *Proceedings of the Fifth International Conference on Machine Learning* (pp. 107–120). Ann Arbor, MI: Morgan Kaufmann.

Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM, 27,* 1134–1142.

Weiss, S., & Kapouless, I. (1989). An empirical comparison of pattern recognition, neural nets, and machine learning classification methods. *IJCAI-89: Eleventh International Joint Conference on Artificial Intelligence* (pp. 781–787). Detroit, MI: Morgan Kaufmann.