

Active Learning for Adaptive Scheduling: A Statistical Approach

Jonathan Gratch
Beckman Institute,
University of Illinois
gratch@cs.uiuc.edu

Steve Chien
Jet Propulsion Laboratory
California Institute of Technology
chien@aig.jpl.nasa.gov

Background

We view learning as a process of hypothesis selection - given some set of alternative hypotheses, a performance metric, and a fixed distribution of examples, a learning algorithm should select (with high probability) a hypothesis that is (close to) the best in terms of its performance over the example distribution. Canonical examples are selecting the concept description with lowest classification error over a distribution of exemplars [9], or selecting the planning heuristic that most improves average planning performance [8]. Numerous hypothesis selection techniques have been proposed both in the machine learning and statistical communities. Learning proceeds by estimating the merit of the alternative hypotheses over randomly selected training examples. Techniques differ in how they attempt to minimize the number of training examples necessary to ensure the quality of the selected hypothesis.

Our research has focused on active learning strategies for reducing the cost of selecting a hypothesis. A hypothesis selection algorithm can be described in terms of its *allocation strategy*: this is a policy that determines how training observations are allocated to the alternative hypotheses. In general, there may be many different allocation strategies that perform equivalently in terms of the quality of the selected hypothesis, but differ in terms of their efficiency. Given a selection problem, S , and a set of allocation strategies, A , one could, *in theory*, rank these allocation strategies by the expected cost of selecting a hypothesis (where cost is defined

in terms of number of observations, cost of observations, etc.). An allocation strategy is *rational* for the problem S if it has the minimum expected cost over the set of alternatives A .

To improve the efficiency of hypothesis selection, we have studied the following active learning method: (1) provide a selection algorithm with a space of possible allocation strategies; (2) each time a selection problem must be solved, the selection algorithm actively determines (an approximation to) the rational policy for that problem, and allocates observations according to this rational policy. We have shown that, by considering certain restricted "allocation strategies spaces," the active learning can significantly increase the efficiency of hypothesis selection. These results are not solely theoretical - the approach has been applied to the problem of identifying good search control heuristics for a real-world scheduling problem at NASA.

The principal results of our work relevant to active learning are as follows:

1. We have shown on both synthetic and real-world data that active learning can significantly reduce the cost of selecting hypotheses [1,4,5].
2. We have demonstrated the applicability of hypothesis selection methods to a real-world scheduling problem [2,5,7] and used active learning techniques in this application - thus demonstrating the applicability of active learning to a real-world problem.

Relevance to Suggested Symposia Topics

Our work is relevant to the suggested symposia topics in several ways,

Theory: We have derived theoretical bounds on the performance improvement of active learning over non-active learning under certain assumptions [1].

Algorithms: we have developed an approximately rational active learning algorithm for a general class of hypothesis selection problems [1,4].

Evaluation: we have demonstrated our techniques on both synthetic and natural data sets [4].

Relevant Publications

[1] J. Gratch, S. Chien, and G. DeJong, "Improving Learning Performance through Rational Resource Allocation," *Proceedings of the Twelfth National Conference on Artificial Intelligence*, Seattle, WA, July 1994, pp. 576-581.

[2] J. Gratch, S. Chien, and G. DeJong, "Learning Search Control Knowledge for Deep Space Network Scheduling," *Proceedings of the Tenth International Conference on Machine Learning*, Amherst, MA, June 1993, pp. 135-142.

[3] J. Gratch, G. DeJong, and Y. Yang, "Rational Learning: Finding a Balance Between Utility and Efficiency," in *Selecting Models from Data: Artificial Intelligence and Statistics IV*,

P. Cheeseman and R. Oldford (eds.), Springer-Verlag, 1994, pp. 11-20.

[4] S. Chien, J. Gratch, and M. Burl, "On the Efficient Allocation of Resources for Hypothesis Evaluation: A Statistical Approach," *IEEE Transactions on Pattern Analysis and Machine Intelligence* (in press).

[5] S. Chien, J. Gratch, M. Burl, "A Statistical Approach to Adaptive Problem-solving for Large-Scale Scheduling and Resource Allocation Problems," *Proceedings of the 1994 AAAI Spring Symposium on Decision-theoretic Planning*, Palo Alto, CA, March 1994, pp. 27-33.

[6] J. Gratch, G. DeJong, S. Chien, "Deciding When and How to Learn," *Proceedings of the 1994 AAAI Spring Symposium on Goal-driven Learning*, Palo Alto, CA, March 1994, pp. 36-45.

[7] J. Gratch, S. Chien, G. DeJong, "Learning Search Control Knowledge to Improve Schedule Quality," *Proceedings of the 1993 Workshop on Knowledge-based Production Planning, Scheduling, and Control*, Chambery France, August 1993, pp. 159-168.

[8] J. Gratch and G. DeJong, "COMPOSER: A Probabilistic Solution to the Utility Problem in Speed-up Learning," *Proceedings of the Tenth National Conference on Artificial Intelligence*, San Jose, CA, 1992, pp. 235-240.

[9] A. Moore and M. McC, "Efficient Algorithms for Minimizing Cross Validation Error," *Proceedings of the Tenth International Conference on Machine Learning*, New Brunswick, MA, July 1994.