Improving Learning Performance Through Rational Resource Allocation

Jonathan Gratch[*], Steve Chien[+], and Gerald DeJong[*]

*Beckman Institute	+Jet Propulsion Laboratory
University of Illinois	California Institute of Technology
405 N. Mathews Av., Urbana, IL 61801	4800 Oak Grove Drive, Pasadena, CA 91109-8099
{gratch, dejong}@cs.uiuc.edu	chien@aig.jpl.nasa.gov

Abstract

This article shows how rational analysis can be used to minimize learning cost for a general class of statistical learning problems. We discuss the factors that influence learning cost and show that the problem of efficient learning can be cast as a resource optimization problem. Solutions found in this way can be significantly more efficient than the best solutions that do not account for these factors. We introduce a heuristic learning algorithm that approximately solves this optimization problem and document its performance improvements on synthetic and real-world problems.

1. Introduction

Machine learning techniques are valuable tools both in acquiring important scientific concepts and in support of decision making under uncertainty. Unfortunately, learning can involve a significant investment of resources. There may be monetary cost in obtaining data and computational cost in processing it. Usually such factors are addressed by informal or intuitive judgements rather than a rational analysis of the costs and benefits of alternative learning operations.

There is a significant learning cost in many diverse application areas. In speed-up learning there is substantial cost associated with processing each training example [Tadepalli92]). In some classification problems it is extremely expensive to obtain data (e.g. protein folding problems) and it is essential to make the most effective use of what data is available. Somewhat paradoxically, cost is also an issue when there is an overabundance of data. In this case it is expensive to use all of the data and one needs some criteria to decide how much data is enough to achieve a given level of performance [Musick93]. Finally, learning may involve ethical issues, as when experiments require giving potentially harmful treatments to human subjects. Under these circumstances it is a moral imperative to utilize as few subjects as possible and to quickly recognize and discard those treatments that worsen the patients condition.

This article discusses factors that influence cost and considers how to use rational analysis (i.e., [Doyle90, Rus-

sell91]) of these factors to minimize learning cost. We discuss this in the context of parametric hypothesis selection problems, an abstract class of statistical learning problems where a system must select one of a finite set of hypothesized courses of action, where the quality of each hypothesis is described as a function of some unknown parameters (e.g. [Gratch92, Greiner92, Kaelbling93, Moore94, Musick93]). A learning system determines and refines estimates of these parameters by "paying for" training examples.

We show how such problems can be cast as resource optimization problems, and that solutions found in this way can be significantly more efficient than solutions that do not account for the cost of gathering information (more than an order of magnitude). Surprisingly, standard hypothesis selection algorithms do not reason about information cost, and are thus less efficient then they might be. We introduce a rational hypothesis selection algorithm that approximately solves the resource allocation problem and empirically document the analytically predicted improvements in efficiency. This algorithm is quite general and can handle situations where the cost of processing data is initially unknown.

2. Hypothesis Selection Problems

Hypothesis selection problems are an abstract class of learning problems where one hypothesis must be chosen from a predefined set based on performance over an unknown distribution of problems or tasks. Performance is characterized by a hypothesis' *expected utility* over the distribution, which must be estimated from training data. Hypothesis selections are at the core of many machine learning approaches. For example, the *utility problem* in speed-up learning is a selection problem in which a problem solving heuristic is chosen from a set of proposed candidates, where expected utility is defined as the average time to solve a problem [Gratch92, Greiner92, Minton88]. The *attribute selection problem* in classification learning is a problem of selecting one of a set of attributes on which to split, where utility is equated with information gain [Musick93]. In reinforcement learning a system must select an action, where utility is equated with expected reward [Kaelbling93].

Several factors affect the cost of identifying a good selection. For example, there may be some cost in obtaining each training example. Furthermore, there can be additional cost for each hypothesis that is evaluated over a given training

example.[1] The challenge is to choose examples and evaluations in such a way as to maximize the likelihood of a good selection with a minimum of learning cost.

Choosing the best hypothesis is problematic as the underlying probability distributions are typically unavailable. Rather than requiring a hypothesis selection algorithm to always select the best hypothesis, algorithms typically obey some probabilistic requirement on the properties of the hypothesis that they select. Several alternative requirements have been proposed. In this paper we adopt the *probably approximately correct* (PAC) requirement favored by computational learning theory [Valiant84]. Under this requirement a hypothesis selection algorithm selects a hypothesis that with high probability is close to the best.

The expected utility associated with a hypothesis can be estimated by observing its performance over a finite set of training examples. However, to satisfy the PAC requirement an algorithm must reason about the discrepancy between the estimated and true utility of each hypothesis. Formally, let there be k hypotheses. Let H_{sel} denote the expected utility of the selected hypothesis and (without loss of generality) let H_i, $i=1..k-1$, be the expected utility for the remaining hypotheses. Let $\hat{H}_i$ be an estimate of the expected utility of the hypothesis. The PAC requirement is that hypothesis *estimated* to be best must be within some user-specified constant ε of the best hypothesis with probability $1-\delta$. It suffices to bound the probability that a hypothesis is estimated to be worse than the selected hypothesis given that it is in fact better, for each of the pair-wise comparisons:

$$\sum_{i=1}^{k-1} \Pr\left[\hat{H}_i < \hat{H}_{sel} - \epsilon | H_i > H_{sel} + \epsilon\right] \quad \leq \quad \delta \qquad (1)$$

Thus the problem of bounding the probability of error reduces to bounding the probability of error of each of the $k-1$ comparisons of H_{sel} to H_i.

To assess these probabilities we must adopt certain statistical assumptions. In this article we adopt the normal parametric model for reasoning about statistical error. This assumes that the difference between the expected utility and estimated utility of a hypothesis can be accurately approximated by a normal distribution (see [Hogg78] for an explanation of the robustness of this common assumption which is grounded in the Central Limit Theorem). The expected cost associated with processing data is also assumed to be normally distributed. Choosing a different parametric model would change the subsequent analysis but analogous results should hold for the conventional models.

With the normality assumption the probabilities in Equation 1 are a reduced to a function of the estimates, the number of examples, n, used for each estimate, the closeness parameter ε, and an unknown variance term, σ^2. Variance measures how much each observation can differ from its expected value, which can be estimated from the data.[2] To simplify the presentation we ignore the ε parameter in the discussion that follows ([Chien94] offers more details). For a given pair-wise comparison, δ_i, the (simplified) probability of incorrect selection is:

$$\delta_i = \Phi\left(-(H_{sel} - H_i)\frac{\sqrt{n}}{\sqrt{\sigma_{sel,i}^2}}\right) \qquad (2)$$

where the function Φ is the quantile function of the standard normal distribution. Intuitively, Equation 2 shows that the probability of a mistake diminishes as the difference in expected utility between the hypotheses increases, as the number of training examples increases, and as the variance of each hypothesis decreases. This relationship can be used to determine how many training examples to allocate to each comparison. If we wish to achieve a given bound of δ_i, then by simple algebra the number of examples needed for a given pair-wise comparison is:

$$n_{sel,i} = \frac{\sigma_{sel,i}^2}{(H_{sel} - H_i)^2}[\Phi^{-1}(\delta)]^2 \qquad (3)$$

where Φ^{-1} is the inverse of the quantile function of the standard normal distribution.

While the variance and true expected utilities are unknown, a class of statistical approaches called *sequential* approaches has been designed for such problems [Govindarajulu81]. These techniques develop estimates of the unknown parameters from a small initial sample size and then incrementally refine these estimates after each subsequent training example. For example, after some number of examples a sequential technique would estimate that the hypothesis it will eventually select is the one with the current highest estimated utility. Such techniques terminate sampling based on an estimate of the sufficient number of training examples. Section 4 introduces a sequential hypothesis selection algorithm that uses a sequential approximation to Equation 3 to decide when to stop sampling.

3. The Value of Rational Learning

The PAC requirement constrains but does not completely determine the behavior of a hypothesis selection algorithm. We would like an algorithm to satisfy the requirement with the minimum cost possible. Several of the factors that contribute to the cost are unknown before learning begins. For this reason standard (non-rational) hypothesis selection al-

---

1. For example, in classification learning a potentially large set of examples must be partitioned for each hypothesized split. In speed-up learning the learning system may have to re-solve the example problem for each candidate heuristic.

2. We "block" examples as in [Moore94] to further reduce sampling complexity. Blocking forms estimates by averaging the difference in utility between hypotheses on each observed example, which can substantially reduce the variance in the data when hypotheses are related (e.g. when each hypothesis is a variant on a basic search control strategy). It is trivial to modify the algorithm to work for the case where it is not possible to block data.
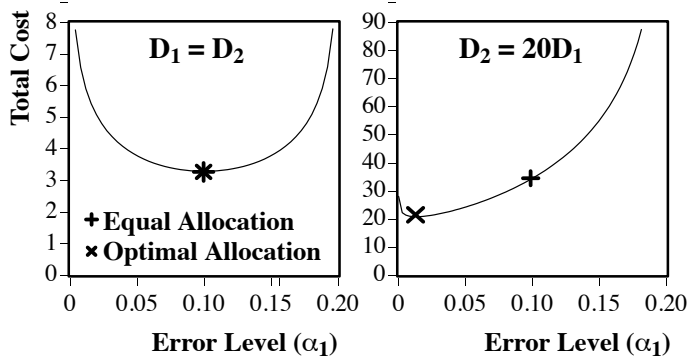
Figure 1. An illustration of the difference between equal and optimal allocation with equal and unequal disparity indices.
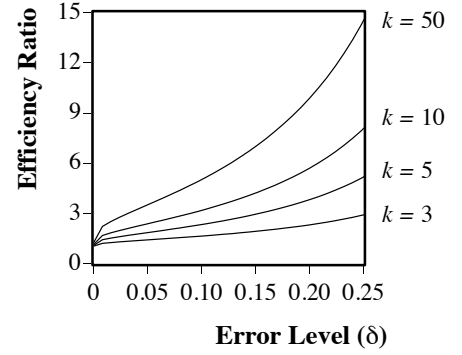

Figure 2: The potential benefits of rational analysis. Shows the ratio of equal allocation cost to optimal cost for several error levels and number of hypotheses.

gorithms ignore these factors when making their selection. This section discusses the relevant factors and shows that they can be folded into a single value, the *disparity index*. We show that in theory an algorithm can achieve large performance improvements by exploiting this information, if only it were available. Comparable performance improvements can be achieved in practice using sequential techniques, as we show in the next section.

Equation 3 illuminates the factors that affect selection cost. In order to satisfy the PAC requirement we must, for each non-selected hypothesis, bound the probability that it is better than the selected hypothesis. The total cost is is the sum of the cost of processing each training example. Equation 3 shows that the number of examples allocated to the two hypotheses increases as the variance increases, as the difference in utility between the hypotheses decreases, or as the acceptable probability of making a mistake decreases.

The first two factors are determined by the environment, but the last, the probability threshold associated with each comparison, can conceivably vary and thus be placed under the control of the hypothesis selection algorithm. The algorithm must only ensure that the *sum* of these probabilities remain less than δ (Equation 1). If one comparison requires a great many examples and another very few, it seems possible that allowing greater error for the first and less for the second might reduce the total cost. In fact, allowing the algorithm to judiciously allocate error to each comparison can result in a substantial reduction in overall cost.

Reducing the cost of selection can be cast as an optimization problem. Total cost is the sum of the number of examples allocated to each comparison (from Equation 3) times the average cost to process an example. Let $c_{sel,i}$ denote the average cost per example to compare the selected hypothesis with hypothesis i. Let α_i be the error level allocated to the comparison. The optimal allocation of error can be determined by solving the following optimization problem:[3]

---
3. This assumes that the cost of processing examples for one comparison is independent of the other comparisons. A more complex analysis is needed to faithfully represent cases where there is significant sharing of cost between comparisons.

Resource Optimization Problem

Choose α_i to minimize $\quad \displaystyle\sum_{i=1}^{k-1} c_{sel,i} \frac{\sigma^2_{sel,i}}{(H_{sel}-H_i)^2}[\Phi^{-1}(\alpha_i)]^2$

Subject to the constraint that $\quad \displaystyle\sum_{i=1}^{k-1} \alpha_i \le \delta$

Of course in an actual hypothesis selection problem the expected utility of the hypotheses, and perhaps the variance and cost will be unknown before learning begins. Without considering such information the only reasonable policy is to assign an equal error level to each comparison (i.e. $\alpha_i=\delta/[k-1]$). However, comparing this *equal allocation* policy with the optimal solution shows that equal allocation can be highly sub-optimal. To see this, consider the case with three hypotheses, $k=3$, which results in two comparisons with error α_1 and $\delta-\alpha_1$. The selection cost is:

$$D_1[\Phi^{-1}(\alpha_1)]^2 + D_2[\Phi^{-1}(\delta-\alpha_1)]^2 \qquad (4)$$

$$\text{where } D_i = \frac{c_{sel,2}\sigma^2_{sel,i}}{(H_{sel}-H_i)^2}$$

The value D_i is called the *disparity index* for comparison i.

To be optimal, α_1 must be chosen so as to minimize the total cost. The equal allocation policy assigns α_1 equal to $\delta/2$. Equation 4 indicates that the equal allocation solution is optimal *only* in the case where the two disparity indices are equal. This is illustrated in Figure 1, which shows the cost equation as a function of α_1, first in the case where the disparity indices are equal, and then when there is a difference between their values. The minimum point under this curve is the optimal cost and the value of α_1 at this point indicates the optimal error allocation. In contrast, the equal allocation policy yields a cost that may differ significantly from this minimum.

In practice it is unlikely that the disparity indices will be equal all for comparisons. Even if the example cost is similar for every hypothesis the variance and expected utilities

of hypotheses will almost certainly differ. The inefficiency of equal allocation increases as the differences between disparity indices increases. The inefficiency also increases as with the number of hypotheses. By taking the difference in disparity indices to the limit it can be shown that for k hypotheses the ratio of equal allocation cost to the optimal cost can be up to $[\Phi^{-1}(\delta/[k-1])]^2 / [\Phi^{-1}(\delta)]^2$. The ratio can be quite large as illustrated in Figure 2. Thus, ignoring disparity information can result in costs up to an order of magnitude greater with as few as ten hypotheses under consideration. This result also shows that the ratio cannot grow without bound and that equal allocation is near optimal for cases with few hypotheses and a small error level.

4. Rational Example Allocation

If the disparity indices were known advance, an algorithm could optimize the cost of selecting a hypothesis. Although this information is unavailable before learning begins, with a sequential approach the algorithm can develop increasingly accurate approximations to this information in the course of learning. These approximations can be almost as effective as the true information in guiding learning behavior. In this section we introduce a rational hypothesis selection algorithm that exploits these approximations to minimize selection cost. This is compared with an efficient non-rational approach similar to Moore and Lee's BRACE algorithm [Moore94]. The superiority of the rational approach is documented on artificial and real-world data sets.

4.1 Interval-Based Selection Algorithm

We first introduce the basic hypothesis selection approach. Rational and non-rational algorithms derived from this approach differ in how they choose hypotheses to further evaluate. The algorithm initially evaluates all hypotheses over a small initial set of n_0 training examples. This is to develop initial parameter estimates and to enhance the robustness of the normality assumption. The algorithm then incrementally processes training examples, deciding to evaluate one or more hypotheses on that example. Learning proceeds incrementally until, to the required level of confidence, one hypothesis ϵ-dominates. The basic approach is as follows:[4]

With hypotheses $H_1..H_k$
Evaluate all hypotheses over n_0 training examples
While no selection
    Let H_{sel} be hypothesis with highest estimated utility

$$\text{IF} \quad \sum_{i=1}^{k-1} \Pr\left[\hat{H}_i < \hat{H}_{sel} - \epsilon \mid H_i > H_{sel} + \epsilon\right] \leq \delta$$

        THEN select H_{sel}
    ELSE Obtain next example
        Evaluate those hypotheses chosen according to
        rational or non-rational policy as outlined below

---

4. See [Chien94] for complete discussion of such rational and non-rational algorithms. The probability is computed with equations analogous to Equation 2.

Equal Allocation Policy. The non-rational algorithm follows the equal allocation policy. Each cycle through the loop allocates an additional example to a pair-wise comparison if its probability of error remains above the fixed level of $\delta/[k-1]$. Eventually every comparison will drop below this error threshold and the procedure will terminate.

Marginal Rate of Return Policy. Using estimates of the disparity indices, the rational algorithm calculates the increase in confidence and the cost of allocating an additional example to each comparison. At each cycle through the main loop the algorithm allocates an example to the comparison with the highest *marginal rate of return*. This is the ratio of increased confidence to increased cost.

This rational policy tries to maximize the decrease in statistical error per unit cost, although we cannot guarantee the strategy achieves the optimal error allocation. Complications include the fact that estimates of the disparity factors differ from their true value and the initial sample size parameter restricts the algorithm's degrees of freedom. Nonetheless, this policy has performed well empirically. The marginal rate of return is estimated using an equation analogous to Equation 2, substituting in estimated for actual utility values. After processing the n_0 initial training examples the algorithm estimates the expected utility, variance, and cost of the various comparisons. The change in error can be estimated by considering how the error would change assuming the current parameter estimates are correct:

$$\Phi\left(-(\hat{H}_{sel} - \hat{H}_i)\frac{\sqrt{n}}{\sqrt{\hat{\sigma}^2_{sel,i}}}\right) - \Phi\left(-(\hat{H}_{sel} - \hat{H}_i)\frac{\sqrt{n+1}}{\sqrt{\hat{\sigma}^2_{sel,i}}}\right) \quad (5)$$

The estimated marginal rate of return for a comparison is computed by dividing this estimate of the reduction of error by the estimated cost of processing an additional training example.

4.2 Empirical Evaluation

We illustrate the performance of the algorithms on simulated and real-world data. The first evaluation uses simulated data with high disparity to illustrate that the rational algorithm achieves performance improvements comparable to what is predicted by the theoretical analysis. The second evaluation uses data from a NASA scheduling application to illustrate the robustness of the approach on a real-world hypothesis selection problem.

4.2.1 Simulated Data. A rational algorithm should significantly outperform a non-rational approach when there is a large difference between the costs, variances, or expected utilities of the various hypotheses. We test this hypothesis for several number of hypotheses and error levels. For all experiments ϵ is set at 1.0 and δ varies from 0.05 to 0.25, in 0.05 increments. We perform tests with three, five, and ten hypotheses. The training examples are randomly generated: All utility values and example costs are normal-

ly distributed according to some expected value and variance, denoted N(value,variance). For all experiments, H_1 is distributed $N(74,50)$ with cost $N(20,1)$, H_2 is distributed $N(72,50)$ with cost $N(50000,1)$. All remaining hypotheses are distributed $N(5,50)$ with cost $N(20,1)$. For each configuration the algorithms are run 5000 times and the reported results are the average over these trials.

Figure 3 summarizes the predicted and observed *efficiency ratio*. This is the cost to select using equal allocation over the cost to select using rational allocation. The performance improvement due to rational allocation is surprisingly close to the limit. This suggests that for this data set the rational algorithm has identified a near optimal error allocation. Note that for large error the observed efficiency drops below the predicted level. This is a consequence of the initial sample size parameter n_0. The rational algorithm is forced to take at least this many examples on every comparison, while in this problem configuration less would suffice to achieve the probability bound. The implication is that when the hypothesis evaluation problem is easy (requires perhaps fewer than n_0 examples to make a selection) the efficiency will be effected more by the choice of the initial sample size than the allocation policy. An interesting issue we have not sufficiently explored is possible strategies for setting the initial parameter size.

4.2.2 NASA Scheduling Data. The test of real-world applicability is based on data drawn from a NASA scheduling application detailed in [Gratch93]. This data provides a test of the applicability of the techniques. Both algorithms assume estimated utility varies normally from the expected utility. In fact, this common assumption is violated by the data as most of the scheduling heuristics are bi-modally distributed. This characteristic provides a rather severe test of the robustness of both approaches.

The heuristic system was developed to schedule communications between earth-orbiting satellites and ground-based antennas. In the course of development extensive evaluations were performed with variant scheduling heuristics. The purpose of these evaluations was to choose a heuristic that generated satisfactory schedules quickly on average. This is easily seen as a hypothesis evaluation problem. Each of the heuristics corresponds to a hypothesis. The cost of evaluating a hypothesis over a training example is the CPU time required to solve the scheduling problem with the given heuristic. The utility of the training example is simply the negation of its cost. In that way, choosing a hypothesis with maximal expected utility corresponds to choosing a scheduling heuristic with minimal average cost.

The application involves several hypothesis selection problems, four of which we use in this evaluation (A, B, C, and D). Each selection problem consists of a set of scheduling heuristics, and data on the heuristics' performance over about one thousand scheduling problems. For the purpose of these experiments the data sets are assumed to correspond exactly to the underlying probability distributions. An experimental trial consists of executing a technique over examples drawn from one of these data sets. Each time a training example is to be processed, some problem is drawn randomly with replacement from the data set. The actual utility and cost values associated with this scheduling problem is used. As in the synthetic data, each experimental trial is repeated 5000 times and all reported results are the average of these trials. In this data the cost and expected utilities of hypotheses are relatively close to each other so the difference between the disparity indices is relatively small across comparisons.

Each trial used an error level of 0.05 or 0.25 and ε equal to 4.0. The results are summarized in Table 1. For each algorithm this shows the average number of examples required to select a hypothesis, the total cost of those examples, and the observed probability that the selection was correct for each of the four selection problems.
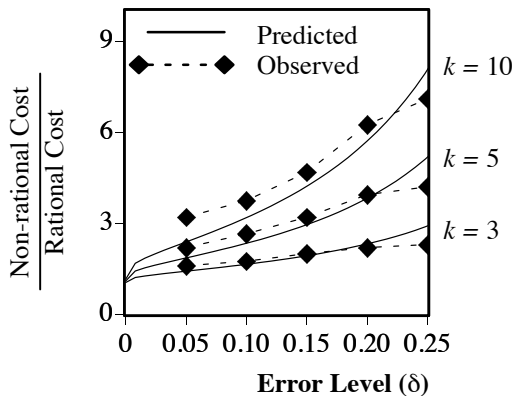
Figure 3. Shows predicted and observed efficiency of the rational allocation policy (the ratio in cost between the non-rational and rational policies). The rational policy shows a significant increase in efficiency.

	Parameters			Equal Allocation			Rational Allocation			Cost Ratio
	k	ε	δ	Ex.	Cost	Pr.	Ex.	Cost	Pr.	
A	3	4	0.25	180	277	1.00	77	120	1.00	2.3
			0.05	908	1,391	1.00	648	998	1.00	1.4
B	2	4	0.25	30	47	1.00	30	46	1.00	1.0
			0.05	74	115	1.00	76	117	1.00	1.0
C	7	4	0.25	1189	1758	0.88	779	1148	0.77	1.5
			0.05	2,371	3,493	0.94	2,153	3,184	0.94	1.1
D	7	4	0.25	3,274	4,993	0.93	2,241	3,429	0.88	1.5
			0.05	7,972	12,037	0.96	7,621	11,583	0.94	1.0

Table 1. Average number of observations, cost, and probability of correct selection for scheduling data.

Both algorithms performed robustly. In each selection problem the PAC requirement was achieved or nearly achieved. This result is particularly remarkable given the data's significant departure from normality. The rational algorithm provides a modest improvement over the equal al-

location algorithm on three out of the four selection problems. The improvement increased with higher error level in accordance with theoretical predictions. In both the scheduling and artificial data the rational algorithm tended to exhibit statistical error closer to the requested bound. The equal allocation strategies excessive conservatism is due to its inflexibility in allocating statistical error in cases where a hypothesis could be discarded with less than n_0 datapoints.

While the scheduling improvements may seem modest, there are three points that must be emphasized. First, the number of hypotheses was small and improvements should grow with the number of hypotheses. Second, in absolute terms the savings are significant. For example, the 350 examples saved in selection problem D translate into about fifteen hours of CPU effort. Finally, in no case did the rational algorithm perform worse. Thus there is little loss, and potential for substantial improvement with rational allocation.

5. Related Work and Conclusions

This analysis can be extended in a number of ways. In many learning situations one may be reluctant to assume normality. For example, when selecting attributes in a decision tree a multinomial model may be more appropriate. We suspect comparable results will hold for a wide range of statistical models but further analysis is necessary. Selection problems could be formalized in a bayesian statistical framework as in [Moore94, Rivest88]. This would eliminate the need for an initial sample but require a rigorous encoding of prior knowledge. Related to this, Howard [Howard70] has extensively investigated a bayesian framework for assessing learning cost in the case of single hypothesis problems.

While this article has focused on minimizing cost in the context of hypothesis selection, the ability to assess both the benefits and costs of learning has been investigated in a variety of contexts both inside and outside of artificial intelligence. For example the tradeoff between goal-directed action and exploration behavior has been studied in reinforcement learning [Kaelbling93]. Another active area of investigation involves the selection of an inductive bias for classification learning tasks. A weaker bias allows higher potential accuracy but requires more data. The selection of an appropriate bias depends on the availability and cost of obtaining training examples as well as usefulness of better prediction (see [desJardins92]). The same issue arises in neural networks and in statistics when one must choose a network topology or statistical model that balances the tradeoff between the fit to the data and the number of examples required to reach a given level of predictive accuracy. Finally, these learning issues can be seen as part of the more general area of *limited rationality*. This is the problem of developing a theory of rational decision making when in the presence of limited reasoning resources [Russell91, Wellman92].

To summarize, we argue that learning algorithms must assess both the benefits and costs of learning. We provide a theoretical analysis of the factors that contribute to learning cost. By reasoning about a value called the disparity index a learning algorithm can achieve the same level of benefit at substantially reduced cost. We introduce a heuristic algorithm that empirically achieves the predicted performance improvements over a non-rational approach. While the improvements on any given hypothesis selection problem may lie well below the theoretical limit, the rational algorithm is unlikely to perform worse and may perform significantly better. Therefore there seems little reason not to adopt this or an analogous rational approach.

References

[Chien94] S. A. Chien, J. M. Gratch and M. C. Burl, "On the Efficient Allocation of Resources for Hypothesis Evaluation in Machine Learning: A Statistical Approach," Technical Report, University of Illinois (forthcoming).

[desJardins92] M. E. desJardins, "PAGODA: A Model for Autonomous Learning in Probabilistic Domains," Ph.D. Thesis, University of California, Berkeley, CA, April 1992.

[Doyle90] J. Doyle, "Rationality and its Roles in Reasoning (extended version)," *AAAI90*, Boston, MA, 1990.

[Govindarajulu81] Z. Govindarajulu, *The Sequential Statistical Analysis*, American Sciences Press, Columbus, OH, 1981.

[Gratch92] J. Gratch and G. DeJong, "COMPOSER: A Probabilistic Solution to the Utility Problem in Speed–up Learning," *AAAI92*, San Jose, CA, July 1992, pp. 235–240.

[Gratch93] J. Gratch, S. Chien and G. DeJong, "Learning Search Control Knowledge for Deep Space Network Scheduling," *ML93*, Amherst, MA, June 1993.

[Greiner92] R. Greiner and I. Jurisica, "A Statistical Approach to Solving the EBL Utility Problem," *AAAI92*, San Jose, CA, July 1992, pp. 241–248.

[Hogg78] R. V. Hogg and A. T. Craig, *Introduction to Mathematical Statistics*, Macmillan Inc., London, 1978.

[Howard70] R. A. Howard, "Decision Analysis: Perspectives on Inference, Decision, and Experimentation," *Proceedings of the IEEE 58*, 5 (1970), pp. 823–834.

[Kaelbling93] L. P. Kaelbling, *Learning in Embedded Systems*, MIT Press, Cambridge, MA, 1993.

[Minton88] S. Minton, in *Learning Search Control Knowledge: An Explanation–Based Approach*, Kluwer Academic Publishers, Norwell, MA, 1988.

[Moore94] A. W. Moore and M. S. Lee, "Efficient Algorithms for Minimizing Cross Validation Error," *ML94*, New Brunswick, MA, July 1994.

[Musick93] R. Musick, J. Catlett and S. Russell, "Decision Theoretic Subsampling for Induction on Large Databases," *ML93*, MA, June 1993, pp. 212–219.

[Rivest88] R. L. Rivest and R. Sloan, A New Model for Inductive Inference," *Second Conference on Theoretical Aspects of Reasoning about Knowledge*, 1988.

[Russell91] S. Russell and E. Wefald, *Do the Right Thing: Studies in Limited Rationality*, MIT Press, Cambridge, MA.

[Tadepalli92] P. Tadepalli, "A theory of unsupervised speedup learning," *AAAI92*, San Jose, CA, July 1992, pp. 229–234.

[Valiant84] L. G. Valiant, "A Theory of the Learnable," *Communications of the ACM 27*, (1984), pp. 1134–1142.

[Wellman92] M. P. Wellman and J. Doyle, "Modular Utility Representation for Decision–Theoretic Planning," *AIPS92*, College Park, Maryland, June 1992, pp. 236–242.