Budgeted Learning of Bounded Active Classifiers

Aloak Kapoor and Russell Greiner*

Department of Computing Science
University of Alberta
Edmonton T6J 2E8
{aloak|greiner}@cs.ualberta.ca

Abstract. Since resources for data acquisition are seldom infinite, the need exists for learners and classifiers that act intelligently under hard budgets. In this paper, we consider problems in which feature values are unknown to the learner and classifier, but can be acquired at a cost. The goal is a learner that spends its learning budget b_L acquiring training data so as to produce the most accurate active classifier that spends at most b_C per instance. From the learner's perspective, purchasing every feature of every instance is sure to approach the underlying distribution asymptotically, but will this yield the best distribution when only b_L dollars worth of data can be collected? In this work, we show empirically that the answer is *no* (especially for small b_L) and present alternate learning strategies that achieve superior performance on a variety of real-world datasets.

1 Introduction

While a doctor may have the option of using a wide variety of medical tests (including MRIs, blood work, etc.) to diagnose a patient, many medical plans involve capitation payments that restrict the per-patient cost of medical diagnosis and treatment. These physicians can only consider diagnostic strategies that spend at most a specified amount; they would clearly want to use the most accurate such strategy. In general, these strategies can operate sequentially: e.g. first performing test Blood_7 (at cost $C(\text{Blood}_7)$), then using this information to decide on the next action; perhaps performing Liver_3 if Blood_7 was positive, but performing Urine_2 if Blood_7 was negative, and so forth. Once the total cost of the tests performed reaches the capitation amount b_C (i.e. if $C(\text{Blood}_7) + C(\text{Urine}_2) + \cdots = b_C$), the strategy must stop collecting information and render a decision — e.g. "Cancer = true". We call such a strategy a "bounded active classifier" [1].

Earlier results [1] have shown that one can PAC-learn this optimal "bounded active classifier" $\text{BAC}^* = \arg\min_b\{\text{error}(b) | b \in \text{cost-}b_C\text{-active classifiers}\}$, assuming the *learner* has no *a priori* resource bound — i.e. it can purchase every feature of as many instances as necessary. Of course, if we are charging the

* Both authors wish to thank NSERC and iCORE for their generous support. Russell Greiner also thanks the Alberta Ingenuity Centre for Machine Learning.

classifier (read "physician") for each feature, it seems strange to provide this information for free to the learner (think "experimental designer"). This paper extends those earlier results by investigating the challenge of learning this BAC* when the *learner* has a fixed budget to spend to acquire the relevant training data — i.e., when the learner can spend only a total of b_L to produce the best classifier that can spend only b_C per instance. Thus, we investigate the problem of budgeted learning of a bounded active classifier.

Although the task is NP-hard in general [1], we demonstrate how to improve the running time of the optimal algorithm, propose a variety of tractable learning strategies, and run tests on datasets with strong, weak, and irrelevant features. Our learning strategies are able to beat the obvious round-robin approach (which spends equally on all features) by significant margins.

The rest of this paper is organized as follows. Section 2 formally introduces the framework for budgeted learning a bounded active classifier, highlights the simplifying assumptions we make, and derives complexity results. Section 3 discusses a variety of algorithms to find good approximate solutions to the problem, including the obvious round-robin purchasing as well as novel approaches. Section 4 gives empirical results that compares the proposed algorithms and discusses their effectiveness. Section 5 reviews related literature from machine learning and statistics, and Section 6 summarizes contributions and considers future work.

2 Formal Description

The "budgeted bounded-active-classifier learner", BBACL, is given the (non-negative) cost $C(X_i) \in \Re^+$ of acquiring each individual feature X_i of any single specified instance[1] and the loss matrix $L = [\ell_{i,j}]$ whose (i, j) element specifies the penalty for returning the class c_i when the true class is c_j; by convention we assume $\ell_{i,i} = 0$ and $\ell_{i,j} > 0$ for $i \neq j$. BBACL also knows how much the learner can spend $b_L \in \Re^+$, and how much the resulting active classifier can spend per instance $b_C \in \Re^+$.

Throughout, the BBACL can see the current "tableau", whose rows each correspond to an instance $i \in \{1, \ldots, m\}$ and whose first r columns each correspond to a feature, and whose $r + 1$st column is the class label. Initially, only the class label is specified; the other $m \times r$ entries are all unknown. In general, we will let $x_i^{(j)}$ refer to the value of the ith feature of the jth instance. At any point, BBACL can perform the (i, j) "probe", to determine the value of $x_i^{(j)}$, at cost $C(X_i)$. This also reduces BBACL's remaining budget from b_L to $b_L - C(X_i)$. Once this budget reaches zero, BBACL stops collecting information and returns a bounded active classifier BAC, which corresponds to a decision tree of bounded depth [2].

[1] If any tests costs $C(X_i) = 0$, we can simply gather that information for each instance. We assume that these costs are independent of each other, for one or more instances; see Section 6.

The score of any BAC B is its expected misclassification error:

$$Q(B) = \sum_{\mathbf{x},y} P(\mathbf{x}, y) L(B(\mathbf{x}), y) \tag{1}$$

Let $All(b_C)$ be the set of all such active classifiers that spend at most b_C per instance, then our goal is the BAC:

$$\text{BAC}^* = \underset{B \in All(b_C)}{\arg\min}\ Q(B) \tag{2}$$

that minimizes this error.

2.1 Simplifying Assumptions

For our work we will assume a constant misclassification cost $\ell_{ij} = 1$ for $i \neq j$ and $\ell_{ii} = 0$. Our algorithms will need to estimate the probabilities over the values of the features of an instance $P(x_i^{(j)})$ to decide which probe to perform. We will take a Bayesian stance by assuming there is a prior distribution over labeled instances, before seeing any data.[2] As a simplification, we will make the Naïve Bayes assumption, which means the distribution of $x_i^{(j)}$ is independent of $x_k^{(j)}$ (for $k \neq i$) as we know the value of the class y_j.[3] Hence, if instance j is labeled with class $+$, we will model the distribution of its ith feature $x_i^{(j)} \sim \text{Dir}(\, \alpha_{1,+}^{(i)}, \ldots, \alpha_{w,+}^{(i)}\,)$ as a Dirichlet distribution with parameters $\alpha_{j,+}^{(i)} > 0$, assuming X_i has w values [3]. These parameters are unrelated to the ones for negatively labeled instances $\alpha_{j,-}^{(i)}$ and also unrelated to the parameters values for other features $X_{h \neq i}$. Initially, we will assume that each such distribution is uniform $\text{Dir}(\, 1, \ldots, 1\,)$. If we later see 29 $Y = +$ instances with $X_i = +$ and 14 $Y = +$ instances with $X_i = -$, the posterior distribution for $x_i^{(j)}$ for a new $Y = +$ instance would be $\text{Dir}(\, 1+29,\ 1+14\,)$. The mean probability for $X_i = +$ here would be $30/(30+15) = 2/3$.

In general, if a variable X's prior distribution is $X \sim \text{Dir}(\, \alpha_1, \ldots, \alpha_w\,)$, then

$$P(X = i) \quad = \quad \frac{\alpha_i}{\sum_k \alpha_k} \tag{3}$$

If we then observe a sample S that includes a_i instances of $X = i$, then X's posterior distribution remains a Dirichlet, with new parameters

$$X|S \quad \sim \quad \text{Dir}(\, \alpha_1 + a_1, \ldots, \alpha_w + a_w\,) \tag{4}$$

[2] The sparsity of the data means the obvious frequentist approach of using simple frequencies is problematic.

[3] Note that Naïve Bayes models often produce good classifiers even for datasets that violate this assumption.

2.2 Complexity Results

[4] proves the following much simpler task is NP-hard: Given a set of coins with known prior distributions and a fixed total number of flips, decide when to flip which coin to decide which coin has the highest head probability. Our framework inherits that negative result. (Identify each coin f_i with a binary feature, whose head probability corresponds to the probability the class is true, given f_i is true: $P(c = +|f_i = +)$, and let $P(c = +|f_i = -) = 0$ for all features.) In addition, [1] shows that computing the best active classifier is NP-hard in general, even if we know the entire distribution. Again, our framework inherits that negative result. That paper also provided a tractable algorithm for (PAC-)learning a bounded active classifier given the underlying distribution (i.e. after observing the feature values); this observation motivates our interest in finding such optimal bounded active classifiers.

3 Algorithms

This section summarizes a number of "budgeted bounded-active-classifier learners". We focus on only the data collection part of the algorithms; after collecting $\$b_L$ worth of feature-values, each of the algorithms then passes its learned Dirichlet distributions to a dynamic program that produces the BAC* that minimizes Equation 2.

3.1 Optimal Policy

As our problem is a finite Markov Decision Process, there exists a deterministic optimal policy for spending the learning budget such that the expected total (expected) misclassification error[4] of the final bounded active classifier is minimized. Mathematically, the optimal learning policy is the one that minimizes Equation 5:

$$\sum_{i \in Outcomes} Prob(i) E(error_{BAC^*}|i) \tag{5}$$

where each "outcome" corresponds to a state in which our learning budget has been fully exhausted and has resulted in posterior Dirichlet distributions over the feature values.

Such a policy can be computed via a bottom-up dynamic program. To see this, note that we can compute the value of all possible outcomes where the learning budget has been exhausted, and then use these to compute the value of all possible outcomes where there is only \$1 left in the learning budget, and so on. Unfortunately, the number of outcomes (and hence the computational complexity) has a prohibitive lower bound:

[4] the first expectation is over the set of possible Dirichlet distributions the learner's purchases can result in, and the second expectation is over the possible instances (x,y) that can occur *given* the resulting Dirichlets

Proposition 1. *Let $|X_i|$ denote the domain size of feature X_i, $|S|$ denote the number of classes, $t = |S| \sum_i |X_i| - 1$, and each feature have unit cost. Then the bottom-up dynamic program must compute the value of* $\Omega\left(\left(\frac{b_L+t}{b_L}\right)^{b_L}\left(\frac{b_L+t}{t}\right)^{t}(t)^{\frac{-1}{2}}\right)$ *outcomes.*

Due to space constraints, we defer the proof to [5]. Given the exponential dependence on the learning budget and the domain sizes of all the features, the straightforward calculation of the optimal policy via dynamic programming can tractably solve only small problems.

We have considered methods of improving upon this naïve dynamic program, namely by reducing the number of subproblems that must be solved. Below we show an interesting way to achieve this reduction by exploiting the equivalence of two "permuted" states under the conditional independence assumption.

Definition 1. *A proper permutation for a feature X_i with t domain values is a bijective function $f : [1, t] \to [1, t]$ that reorders the t parameters for every Dirichlet distribution on X_i.*

Example 1. Let

$$(X_i|Y = 0) \sim \mathrm{Dir}(4, 2, 7), \quad (X_i|Y = 1) \sim \mathrm{Dir}(3, 8, 5)$$

Then a proper permutation for feature X_i is:

$$(X_i|Y = 0) \sim \mathrm{Dir}(7, 2, 4), \quad (X_i|Y = 1) \sim \mathrm{Dir}(5, 8, 3).$$

Proposition 2. *Let us identify a "state" of our problem by the value of b_L and the set of Dirichlets over the feature-class pairs. Consider any two states A and B, which have equal values of b_L and are such that the Dirichlets of A can be made equal to the Dirichlets of B by specifying a set of r proper permutations, one for each feature X_i. Under these conditions, the expected value of state A is equal to the expected value of state B when following an optimal policy, and the optimal action to take from state A is the optimal action to take from state B.*

We defer the proof to [5]. This proposition allows us to improve the naïve dynamic program by reusing the computed value of a state A for properly permuted versions of A. The real-time improvement using Proposition 2 is shown in Table 3.1 below. In the last case, the naïve dynamic program ran out of memory after more than two hours, while our improved version finished properly in under an hour. Unfortunately such improvements are not sufficient to remove the exponential complexity of the dynamic program, leading to the more tractable, suboptimal approaches we consider next.

3.2 Round Robin (RR)

This obvious algorithm simply purchases *complete* instances until its budget b_L is exhausted. It draws examples randomly, and so expects to have collected

Table 1. Proposition 2's reduction in computation time

b_L	b_C	Features	Classes	Domain Size	Naïve	Improved
2	4	6	2	4	161 sec	65 sec
3	2	4	3	3	888 sec	432 sec
4	3	4	3	3	8280 sec	3360 sec

data about members of each class y in proportion to $P(Y = y)$. If there are r unit-cost features, we expect to know everything about roughly b_L/r instances. Notice this approach implicitly assumes that all features are equally valuable in learning the target concept.

3.3 Biased Robin (BR)

A more selective approach than Round Robin is to purchase a single feature and test whether or not its observed value has increased some measure of quality. The Biased Robin algorithm follows such an approach, continually purchasing feature X_i as long as it improves quality, and otherwise moving to feature X_{i+1} (and of course looping back to X_1 after X_r). There are several choices for how to measure quality or loss and these are discussed in Section 3.6. As further motivation for this algorithm, [6] found it to be one of the best approaches for budgeted learning of a passive Naïve Bayes classifier, albeit with a a different loss function. This method also corresponds to the "Play the Winner" approach discussed in [7].

3.4 Single Feature Lookahead (SFL)

With a limited learning budget, one would like to avoid wasting purchases on poor features. This motivates a prediction-based approach, which uses a loss function to approximately predict the expected loss incurred after making each possible purchase. If only the next purchase is considered, then this reduces to the (1-step) greedy algorithm.

SFL uses this prediction based approach, but controls the level of myopia or "greediness" involved by providing an additional parameter, d, the lookahead depth. With a lookahead depth of d, SFL calculates the expected loss of spending its next $\$d$ continuously purchasing feature i of class j. That is, if S denotes our current set of Dirichlets and S' denotes the Dirichlets after spending $\min(\$d, \$b_L)$ purchasing feature X_i of a $Y = j$ instance, then the expected loss for (i, j) is:

$$SFL(i,j) = \sum_{S'} P(S'|S) Loss(S') \tag{6}$$

The feature-class pair (i, j) with lowest expected loss is purchased, the Dirichlets are updated based on the observed outcome of that purchase, and Equation

6 is computed again for all feature-class pairs. This process repeats until the learning budget is exhausted. SFL has shown strong performance on two previously investigated variants of the budgeted learning problem [6, 8], and is thus well-motivated for our task.

3.5 Randomized Single Feature Lookahead (RSFL)

Our experiments show that the SFL algorithm is biased toward the feature that currently looks best. That is, SFL often takes very long runs of purchasing a single discriminative feature-class pair and neglects to explore other potentially good features. This property can be problematic, particularly when a dataset contains several discriminative features that can jointly yield a more accurate BAC than any single feature by itself. The Randomized Single Feature Lookahead algorithm (RSFL), alleviates this problem by increasing exploration among the best looking feature-class pairs. The RSFL algorithm is very similar to SFL, as it too calculates the expected loss in Equation 6 for each feature-class pair. However, rather than deterministically purchasing the feature-class pair with the best SFL score, RSFL considers the best X feature-class pairs and for each feature-class pair (i,j) in this set, it chooses to purchase feature i of class j with probability:

$$\frac{\exp \frac{-SFL(i,j)}{\tau}}{\sum_{i,j} \exp \frac{-SFL(i,j)}{\tau}} \tag{7}$$

Here τ is a temperature controlling exploration versus exploitation. We set $\tau = 1$ throughout this paper.

3.6 Loss Functions

As mentioned earlier, several of our algorithms rely on a loss function

$$g : \{\text{Dirichlet distributions over features}\} \rightarrow \mathbb{R} \tag{8}$$

that measures the quality of a given probability distribution. After experimenting with several different choices of loss functions, we found Conditional Entropy Loss and Depth 1 BAC Loss to be effective.[5]

Conditional Entropy measures the uncertainty of the class label Y given the value of a feature X_i:

$$-\sum_x P(X_i = x) \sum_y P(Y = y|X_i = x) \log_2 P(Y = y|X_i = x) \tag{9}$$

The Biased Robin algorithm uses loss Equation 9 before and after the purchase of feature X_i to determine whether the purchase improved the prediction ability of X_i.

[5] The obvious loss function is just to use Equation 2 to compute the expected error of the optimal BAC. However, since loss functions can be called several times on a single purchase, the computational expense of computing Equation 2 is prohibitive.

On the other hand, the Depth 1 BAC loss function

$$\min_i \sum_x P(X_i = x) \min_y (1 - P(Y = y | X_i = x)) \tag{10}$$

is used by SFL, RSFL, and the greedy algorithm. This function calculates the expected misclassification error of the best Depth 1 BAC.

4 Experimental Results

To compare the algorithms, 50 repetitions of five-fold cross validation were performed on several datasets from the UCI Machine Learning Repository [9]. Datasets with continuous values were discretized using supervised entropy discretization [10]. Each dataset was randomly partitioned into five folds. The algorithms were run five times, and on each run a single fold was set aside for testing, while the remaining four were available for purchasing. We used the average of values from these five runs as the misclassification error of each algorithm on the whole dataset. This process was repeated a total of 50 times to reduce the variance and get a measure of the average misclassification error. Thus, each point in the graphs that follow represents 50 repetitions of five-fold CV.

In the first set of experiments all features have unit cost and the datasets contain some irrelevant features. We set the classifier's budget to $b_c = 3$, as this is large enough to allow several features to be relevant, but small enough to keep computations tractable. Initially, all Dirichlets parameters are set uniformly to 1. For reference, each graph also includes a gold standard "All Data" algorithm which is allowed to see the *entire* dataset, and thus represents the best that one can do using the Naïve Bayes assumption on the data.

Figure 1 shows the performance of the algorithms on the glass identification dataset. The glass dataset is a binary class problem with nine features whose domain sizes vary between one and three. The four features that have a domain size of one represent irrelevant information that any learning algorithm (especially one under a constraining budget) should avoid. Both RSFL and BR learn better than the obvious RR algorithm for all learning budgets considered. In fact, we found the optimal $b_C = \$3$ BAC produced by the "All Data" algorithm involves four different features, and these four features are precisely the ones that RSFL and BR purchase heavily during learning. This is in contrast to the RR purchasing behaviour that spends on all features equally, despite their unequal predictive power. Finally, SFL and Greedy spend their entire budget on only one or two features during learning, which accounts for their low accuracy BACs.

The Breast Cancer dataset contains ten features, only one of which is irrelevant to the concept. The Breast data is particularly interesting because nearly all its features are good predictors, but three features have markedly lower conditional entropy than the rest. To produce the lowest error BAC, the learning algorithms must discover the superiority of these three features. We find RSFL does exactly this, spending 20%, 21%, and 32% of its budget respectively on the

Fig. 1. Identical costs and some irrelevant features – RSFL and BR outperform RR.

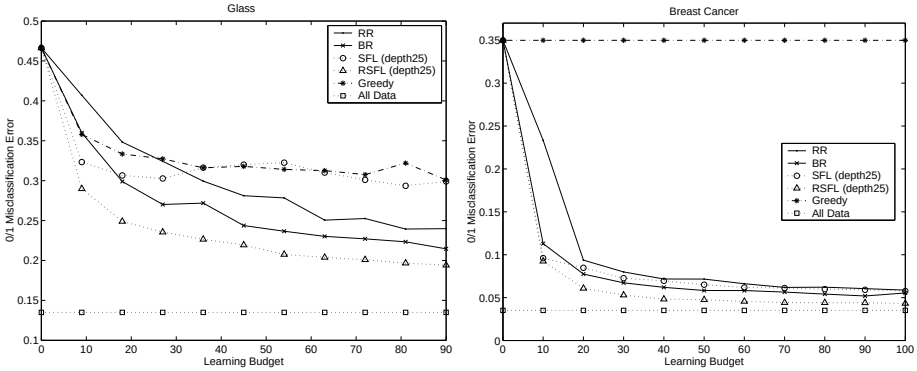

three strong features. In comparison, RR spends 10% of its budget on every feature which makes it much more difficult for it to separate the top features from the rest. BR also performs better than RR for all learning budgets considered.

The next set of experiments, shown in Figure 2, considers datasets without any irrelevant features. The Iris dataset has only four features and is a three class problem. Given that all four features are relevant, and that $b_C = 3$ in this experiment, the optimal BAC requests every feature at some point in its tree. With only four features to consider, RSFL is able to test them all effectively and produce better BACs than RR for all budgets considered. BR is also competitive with RR, except at some of the very low budgets where BR's exploration model prevents it from ever investigating some of the features.

Fig. 2. Identical costs, no irrelevant features - RR still suboptimal.

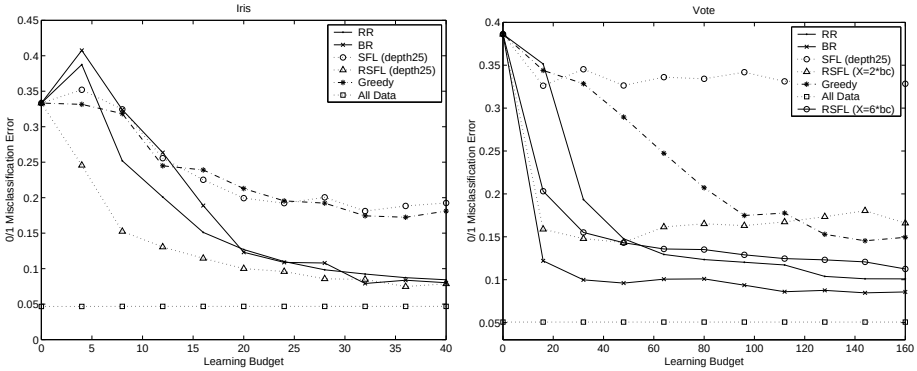

Figure 2 shows another binary class problem, the Vote dataset, that contains 16 relevant features. Many of these features have similar (high) predictive power. Once again we see that both RSFL and BR beat RR. RSFL asymptotes fairly

soon – it spends its budget finding a few strong features quickly and outputs a fairly low error BAC. As expected, at larger budgets RR collects enough information on every feature to find many more suitable candidates for its BAC than RSFL can. We show in the graph that one can improve the performance of RSFL by increasing the number of top feature-class pairs that RSFL considers on this dataset. We also observe that BR's exploration model is particularly well suited to this task because it is able to collect information on every feature at larger budgets, and this property is crucial on a dataset such as Vote with a large number of predictive features.

Our final set of experiments involved datasets where the features differed in cost. Both the Heart Disease dataset and the Pima Indians dataset have known cost data [9], and we use a scaled version of these costs in our tests. The scaled Heart Disease costs range from \$1 to \$7, and our tests are run with $b_c = \$7$. This dataset represents the worst case for RR, because the irrelevant features happen to be the most expensive ones. In fact, RSFL achieves the same error rate after \$100 that RR takes \$500 to reach. In the Pima dataset, feature costs are between \$1 and \$5, and we set $b_c = \$5$. The two irrelevant features have cost \$1, and the single best feature is \$4. Once again, BR and RSFL dominate RR for all budgets considered.

Fig. 3. Different feature costs - RSFL and BR dominate RR.

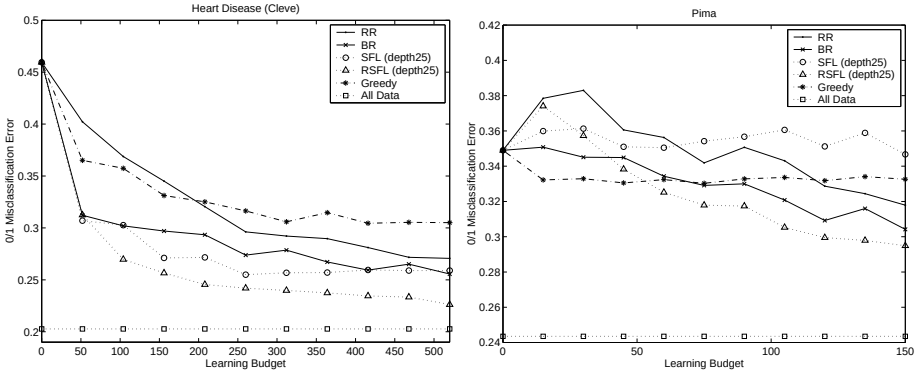

5 Related Work

There are a number of different senses of "costs" in the context of learning [11]. This research considers two costs: the costs paid by the learner to acquire the relevant information at training time to produce an effective classifier and also the costs paid by the classifier, at performance time, to acquire relevant information about the current instance. We impose hard constraints on the cost of tests that can be performed per instance, and on the expenses paid by the learner.

Many existing (sub)fields, such as active learning [12] and experimental design [13] (as well as earlier results such as [6]) focus on only the first of these costs – e.g. bounding how much the learner can spend to produce an accurate *passive* classifier. In addition, many of these systems request the *class label* for an otherwise *completely specified instance*. Thus they require only a single quantity per instance. Our problem is the complement of this: class labels are known but attribute information must be purchased. Unlike most of the other models, this means our work may need to consider the correlations amongst the many unknown properties of an instance. Other results seeking to reduce the sample complexity for learning include decision theoretic subsampling [14], on-line stopping rules [15], progressive sampling [16], and active feature value acquisition [17]. We note that these techniques differ from our approach because we place a firm budget on the learner's ability to acquire information, while these approaches typically allow the learner to purchase until some external stopping criteria (for instance, accuracy) is satisfied.

Recent work by [18] considers a subproblem that we encounter in our overall framework: how to represent the class distribution when only a firm budget of n training examples can be used. For example, if our budget allows for ten training examples, should we select five from class one and five from class two or skew our samples toward the majority class.

As for the costs paid by the classifier at performance time, both [19] and [1] attempt to produce a decision tree that minimizes expected total cost. However, neither work assumes an a priori resource bound on the learner, thereby allowing for unconstrained amounts of training data with which to build these classifiers. Again, our work makes the more realistic assumption that if data costs money at performance time, it very likely costs money at learning time as well.

Finally, we can view our model as a (fixed horizon, partially observable) Markov Decision Process (MDP) [20]: after performing a set of tests, the learner is in a "state" associated with those tests and their outcomes; it can then select some new test to run, which stochastically maps that state to a new state (depending on the outcome of that test). There is a vast literature on finding optimal policies for such MDPs [21]. We note that although the MDP formulation is theoretically clear, it has not yielded strong results in our experiments due to the dimensionality and lack of obvious features for function approximation. Recent research has provided new ideas for handling large MDPs [22], but such methods are still unsuitable to our task because of the exponential dependence on the size of the learning budget (i.e. the horizon). A simpler version of our problem also exists in the MDP framework [8], and the results of that work motivate several of the policies that we adapt for the budgeted learning of bounded active classifiers.

6 Conclusions

Many standard learning algorithms implicitly assume the features are always available for free, to both the learner at "training time" and later the classifier, at

"performance time". This paper extends those systems by explicitly considering these costs, at both training and performance time. It introduces the formal framework for budgeted learning a bounded active classifier, and presents some complexity results. It also proposes an improvement for the optimal algorithm which it proves works effectively. Moreover, this paper motivates and defines a variety of tractable learning strategies and shows they work effectively on data with identical and different feature costs. In particular, we demonstrated that our proposed strategies can often do much better than the obvious algorithm – "Round Robin" – especially when training data is limited. We devote future work to developing a tractable approximation algorithm with guarantees on learning performance.

References

1. Greiner, R., Grove, A.J., Roth, D.: Learning cost sensitive active classifiers. Artificial Intelligence (2002)
2. Dobkin, D., Gunopoulos, D., Kasif, S.: Computing optimal shallow decision trees. In: International Workshop on Mathematics in Artificial Intelligence. (1996)
3. Heckerman, D.: A tutorial on learning in bayesian networks. In: Learning in Graphical Models. The MIT Press (1999)
4. Madani, O., Lizotte, D.J., Greiner, R.: Active model selection. Technical report, University of Alberta (2004)
5. Website: http://www.cs.ualberta.ca/∼aloak/bl/bl.html (2005)
6. Lizotte, D.J., Madani, O., Greiner, R.: Budgeted learning of naives-bayes classifiers. In: Proceedings of Uncertainty In Artificial Intelligence. (2003)
7. Robbins, H.: Some aspects of the sequential design of experiments. Bulletin of the American Mathematical Society (1952)
8. Madani, O., Lizotte, D.J., Greiner, R.: Active model selection. In: Proceedings of Uncertainty in Artificial Intelligence. (2004)
9. S. Hettich, C.B., Merz, C.: UCI repository of machine learning databases (1998)
10. Fayyad, U., Irani, K.: Multi-interval discretization of continuous-valued attributes for classification learning. In: International Joint Conference on Artificial Intelligence. (1993)
11. Turney, P.: Types of cost in inductive concept learning. In: Workshop on cost sensitive learning (ICML). (2000)
12. Cohn, D.A., Ghahramani, Z., Jordan, M.I.: Active learning with statistical models. In: Advances in Neural Information Processing Systems. (1995)
13. Chaloner, K., Verdinelli, I.: Bayesian experimental design: A review. Statistical Science (1995)
14. Musick, R., Catlett, J., Russell, S.: Decision theoretic subsampling for induction on large databases. In: International Conference on Machine Learning. (1993)
15. Schuurmans, D., Greiner, R.: Sequential pac learning. In: Conference on Learning Theory. (1995)
16. Provost, F., Jensen, D., Oates, T.: Efficient progressive sampling. In: International Knowledge Discovery and Data Mining Conference. (1999)
17. Melville, P., Saar-Tsechansky, M., Provost, F., Mooney, R.: Active feature-value acquisition for classifier induction. In: ICDM. (2004)
18. Weiss, G.M., Provost, F.: Learning when training data are costly: the effect of class distribution on tree induction. Journal of Artificial Intelligence Research (2003)

19. Turney, P.: Cost-sensitive classification: empirical evaluation of a hybrid genetic decision tree induction algorithm. Journal of Artificial Intelligence Research (1995)
20. Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach. Prentice Hall (2002)
21. Sutton, R.S., Barto, A.G.: Reinforcement Learning. The MIT Press (1998)
22. Kearns, M., Mansour, Y., Ng, A.Y.: A sparse sampling algorithm for near-optimal planning in large markov decision processes. Machine Learning (2002)