

Decision-centric Active Learning of Binary-Outcome Models

Maytal Saar-Tsechansky* • Foster Provost**

*Red McCombs School of Business, The University of Texas at Austin, Austin, Texas 78712, USA

**Leonard N. Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012, USA
maytal@mail.utexas.edu • fprovost@stern.nyu.edu

It can be expensive to acquire the data required for businesses to employ data-driven predictive modeling, for example to model consumer preferences to optimize targeting. Prior research has introduced “active learning” policies for identifying data that are particularly useful for model induction, with the goal of decreasing the statistical error for a given acquisition cost (*error-centric* approaches). However, predictive models are used as part of a decision-making process, and costly improvements in model accuracy do not always result in better decisions. This paper introduces a new approach for active data acquisition that targets decision-making specifically. The new *decision-centric* approach departs from traditional active learning by placing emphasis on acquisitions that are more likely to affect decision-making. We describe two different types of decision-centric techniques. Next, using direct-marketing data, we compare various data-acquisition techniques. We demonstrate that strategies for reducing statistical error can be wasteful in a decision-making context, and show that one decision-centric technique in particular can improve targeting decisions significantly. We also show that this method is robust in the face of decreasing quality of utility estimations, eventually converging to uniform random sampling, and that it can be extended to situations where different data acquisitions have different costs. The results suggest that businesses should consider modifying their strategies for acquiring information through normal business transactions. For example, a firm such as Amazon.com that models consumer preferences for customized marketing may accelerate learning by proactively offering recommendations—not merely to induce immediate sales, but for improving recommendations in the future.

1. Introduction

With advances in computing power, network reach, availability of data, and the maturity of induction algorithms, businesses are taking advantage of automated predictive modeling to influence repetitive decisions, often as tools for extracting customer, competitor and market intelligence (Berry and Linoff, 2004). Consider an example: telecommunications companies face severe customer retention problems, as customers switch back and forth between carriers (the problem of “churn”). For each customer, at each point in time, the company faces a decision between doing nothing and intervening in an attempt to retain the customer. This paper focuses on modeling uncertain outcomes for a certain type of decision problem:¹ a firm or individual must decide repeatedly either to take a specific action or not to take the action. If the action is taken, it results in one of two uncertain but well-defined outcomes. We assume that the utilities associated with each outcome and with not taking the action can be estimated or are known.

¹ We note below where our theoretical derivation and our proposed techniques apply more generally; however, for most of the treatment in the paper we assume this particular setting. In the penultimate section we discuss extensions to this setting and limitations of the techniques we introduce.

Such decisions can be based on predictive models built from data on known outcomes; however acquiring such data can be costly. Consider our customer-retention example.² Outcomes can be modeled based on data about customers and their responses to the firm's actions, and firms collect such data in various ways. They undertake direct solicitations, for example via experimental special offers, via customer surveys, and via interactions such as Amazon's on-line acquisition of product ratings. Firms also acquire data directly from third parties. For example, Acxiom³ sells detailed consumer demographic and lifestyle data to many firms in support of their marketing efforts; other firms such as Abacus⁴ maintain and sell specialized consumer purchase information. Firms also collect information indirectly, in the course of normal business interactions, for example by observing responses to offers or the results of everyday merchandizing decisions. All these acquisitions involve costs to the firms.

For this paper, we consider the acquisition of a particular kind of data. Following the terminology used by Hastie et al. (2001) we refer to the data used to induce the predictive model as the *training* data. Each training data point is a historical *example* of the phenomenon being modeled (e.g., the targeting of a particular customer), described by various variables including a target variable (e.g., indicating whether or not the customer responded to the offer). Of particular importance to this paper, in training data each historical example's *label*, the target variable's value, is known. We focus on the acquisition of these labels, which can be costly. For example, obtaining response data for individual consumers involves solicitation costs, incentives required for revealing preferences, negative reactions to solicitations, etc. Firms also incur opportunity costs when labels are acquired over time through normal business interactions. For example, making a particular offer to a sample of web-site visitors, for the purpose of acquiring training labels, may preclude making another offer already known to be profitable.

² Later we will expand on this example, noting that more than one model may be used for a complex decision. For this paper we ignore some issues pertinent to this example, such as calculations of lifetime value, but see (Rosset et al. 2003) for a treatment from the data-mining perspective.

³ <http://www.acxiom.com/>

⁴ <http://www.abacus-us.com/>

To reduce the cost of label acquisition, researchers have studied the selective acquisition of labels (see Section 2.2), reasoning that focused selection of cases for label acquisition should yield more-accurate models for a given acquisition budget, as compared to the standard approach of acquiring labels for cases sampled uniformly at random. There are various label-acquisition strategies for reducing statistical error (*error-centric* approaches). However, business applications employ predictive models to help make particular business decisions. Of course, a more accurate model may lead to better decisions, but concentrating on the particular decisions themselves (a *decision-centric* approach) has the potential to produce a more economical allocation of the acquisition budget. Prior work has not addressed label acquisition to facilitate decision-making directly.

This paper makes several contributions, largely in the design-science setting described by Hevner et al. (2004). (1) We introduce decision-centric label acquisition as a general approach with intuitive advantages over the traditional error-centric approach. We demonstrate that, in comparison, the error-centric approach can be wasteful in a decision-making context, because some costly improvements to statistical accuracy can have no effect—or, counterintuitively, can have a detrimental effect—on decision-making accuracy.

(2) We introduce and contrast two different general strategies for decision-centric label acquisition. One strategy builds a model to estimate outcome probability, later to be combined with utility information. The other strategy constructs a new modeling task incorporating the utility information. These strategies are decision-centric and they are generic, i.e., they can be applied to any predictive modeling technique, and thus they extend the existing literature on active learning for improving a model’s statistical accuracy.⁵

(3) Using data from a direct-marketing campaign, we conduct an experimental comparison of various label acquisition strategies. The results show that an implementation of one of the two decision-centric strategies, which we call Goal-Oriented Active Learning (GOAL), clearly is preferable over the

⁵ Note that the problem that has come to be known as “cost-sensitive learning” constitutes a special case of optimal decision-making. Decision-centric techniques apply to such problems as well; see the Appendix for a discussion.

alternatives (including the alternative decision-centric approach, error-centric acquisition, and uniform random acquisition), resulting in better decision-making accuracy and thereby in larger profit. On the other hand, decision-centric acquisition produces models with higher statistical error than error-centric acquisition, demonstrating that the technique selects appropriately those error reductions that will improve decision-making and avoids wasteful improvements in statistical accuracy. Furthermore, for cases where utilities are not known with certainty, for example when they too must be estimated from the data, we analyze GOAL’s robustness to error in utility estimations. The results show GOAL to be relatively robust to utility-estimation error, with decision-making degrading gradually with increasing estimation error—in the worst-case converging to the performance of uniform random sampling. We also demonstrate an extension to GOAL for scenarios where label-acquisition costs are not uniform.

2. Active Learning: Terminology, Framework and Prior Work

As described above, a decision resulting in an action will yield one of two well-defined outcomes. Predicting which outcome will occur therefore can be modeled as a classification problem, and the decision-maker faces the opportunity to induce a classification model from training data. Formally, a classification model is a mapping of an input vector $x \in X$ to a discrete class $y \in Y$. The model is induced from a training set of labeled “examples”— (x, y) pairs—which are generalized into a concise model $M: X \rightarrow Y$. Differently from such a *categorical* classification model, a *probabilistic* classification model estimates a probability distribution over Y for a given input vector x . A maximum a posteriori classification rule would then map x to the class y with the highest estimated probability. Alternatively, the estimated probability can be used in decision-making to estimate the expected utility of the corresponding action.

2.1 Active Learning

A modeling technique’s predictive performance for a particular domain of application is captured by its “learning curve,” the relationship between the model’s performance (e.g., its accuracy) and the number

of training examples used to induce the model. Typically, a model's performance improves with the number of training examples, but the marginal improvements diminish as the number of training examples increases (cf., Cortes, et al., 1994; Perlich, et al., 2003). Consider a typical setting where there are many potential training examples for which labels can be acquired at a cost, e.g., consumers to whom we can send an offer to determine whether they will respond. Let us refer to examples whose labels are not (yet) acquired as *unlabeled* examples, and to examples whose labels already have been acquired as *labeled* examples. The goal of active learning (Cohn et al., 1994) is to acquire the labels of unlabeled examples judiciously in order to produce a better model as economically as possible. Specifically, for a given number of acquisitions, we would like the model's generalization performance to be better than if we had used the alternative strategy of acquiring labels for a representative sample of examples (via uniform random sampling). Active learning chooses to label examples estimated to be particularly informative for reducing model error, so ideally it results in a steeper learning curve exhibiting higher accuracy for any given number of acquisitions, or reduced acquisition cost for any given level of accuracy.

The challenge of an active learning method is therefore to estimate the relative contribution of possible training examples prior to acquiring their labels. Most existing methods identify examples for acquisition based on some notion of the uncertainty of the currently held model. For example, *uncertainty sampling* (Lewis and Gale, 1994) is a generic active learning method designed for inducing binary classifiers. Uncertainty sampling defines the most informative examples (whose labels should be acquired) as those examples for which the current classification model estimates probabilities closest to 0.5. The rationale is that the classification model does not capture strong discriminative patterns for predicting the class membership of these examples, and so the estimation of the classification boundary can be improved most by acquiring their labels for training.

We propose instead a strategy that employs the (estimated) costs of actions and benefits from outcomes, in addition to the estimation of outcome probabilities, in order to evaluate the potential contribution to the decision-making task from each acquisition.

2.2 Prior Work

The role of information acquisition in decision-support has been studied in a variety of settings. An important line of relevant research involves the classic multi-armed bandit problem originally proposed by Robbins (1952). Given k slot machines with different rates of return, a gambler has to decide which machine to play in a sequence of trials in order to maximize the overall reward. The challenge stems from the tradeoff between (i) the acquisition through experimentation of information about the returns of different machines and (ii) profit generation by playing the machine estimated to have the highest winning odds. Differently from active learning, the gambler estimates the success probability of each machine, whereas an active learner must induce a multivariate predictive model.

Prior work on technology choice also has considered the acquisition of costly information for estimating the success of alternative courses of action (McCardle 1985, Kornish 2006). McCardle (1985) analyzes technology adoption decisions, where at each phase the decision maker decides whether to acquire information to improve the estimation of technology success or to act based on the current knowledge. McCardle's decision maker must estimate a single parameter (i.e., the probability of success) and can only acquire observations of a single random variable. Active learning must consider the induction of a predictive model and choose between potential training examples of varying informative value.

Research on predictive models for business intelligence has focused primarily on modeling techniques (e.g., West et al., 1997; Moe and Fader, 2004). As discussed above, such intelligence relies on expending significant time and/or money to obtain data. Therefore, it is important to understand the fundamental properties of the information that will be most effective for inducing accurate models, so as to direct the acquisition of such information. The fundamental notion of active learning for building predictive

models has a considerable history in the literature. Work on optimal experimental design (Kiefer, 1959; Fedorov, 1972), or *OED*, examines the choice of observations for inducing parametric statistical models. The objective is to devise a distribution over the independent variables reflecting the contribution of label acquisition for these examples. OED generally depends upon a closed-form objective function that cannot be derived for non-parametric models. Simon and Lea (1974) describe conceptually how induction involves the simultaneous search of two spaces: the results of searching the hypothesis space can affect how the example space will be sampled, and vice versa. Winston (1975) suggests that the best examples to select for induction are "near misses," instances that are very similar to correctly classified instances. Most existing active-learning methods address categorical classification problems and estimate the value of acquisitions based on some notion of the uncertainty of the currently held model. This idea was first introduced by the Query by Committee (QBC) algorithm (Seung et al., 1992), where given a stream of training examples an example is considered informative, and its label is acquired, if different models from the current version space (Mitchell, 1997) assign the example to different classes. A variety of alternative measures of uncertainty have been proposed (e.g., Cohn et al., 1994; Lewis and Gale, 1994; Saar-Tsechansky and Provost 2004).

3. Active Learning for Decision-Making

Previously introduced active learning strategies identify and acquire information that is estimated to produce the largest reductions in the model's statistical prediction error. These error-centric strategies make no reference at all to the decision-making context in which the model will be used. Learning a model from data should allow the decision maker to better identify the most cost-effective decision—the action with the highest expected utility.⁶ We argue that knowledge of the structure and utilities of the decision-making task also can aid in deciding what additional data to acquire.

⁶ For clarity, we will assume that the utility for the business is immediate profitability. In practice, business utility would take into account other factors, such as customer satisfaction, shareholder satisfaction, regulator satisfaction, and so on.

Two different sorts of learning processes may support the identification of the most cost-effective decision:

1. **Decision Learning:** The optimal decision itself (rather than the outcome of the action) is cast as a categorical classification task, incorporating any and all utility information. Specifically, each example of a potential decision is labeled to reflect the (estimated) optimal decision and a targeting model is induced, mapping each example directly to the predicted highest-expected-utility decision.
2. **Response Learning:** The decision maker induces a predictive model that estimates probabilities of action outcomes, which are used subsequently by the decision maker to estimate and compare the expected utilities from alternative actions (action versus no-action in our case). The targeting strategy applies Von Neumann-Morgenstern expected utility theory (Von Neumann and Morgenstern, 1944): given the estimates, for each decision-making opportunity the decision-maker identifies the action that will produce the highest utility in expectation.

Both approaches employ learned models when inferring the optimal action. The fundamental difference between these two methods is that Response Learning applies induction to model the outcome uncertainties, e.g., in consumers' responses, and applies existing theory about strategies for optimal decision-making to derive the predicted optimal action. Decision Learning, by contrast, employs expected utility theory to label examples with the optimal action and applies induction to learn a model of optimal action from the data. As we will see, these differences can result in different decisions.

3.1 A Decision-Learning Approach: predicting the best action

We explore decision learning, first to evaluate its efficacy for decision-making in general, and then to assess how active learning can be applied for cost-effective label acquisition. As just discussed, decision learning will build a model that maps decisions directly to actions. Thus, the training examples will be descriptions of decisions to be made. Consider the case where an example describes a consumer and the decision in question is whether or not to mail her a solicitation. Central to the formulation of the

learning problem is the determination of the labels for the training examples. A possible (naïve) approach is to label an example decision as “solicit” if the corresponding historical solicitation had been profitable. However, this naïve approach does not take into account the role of expected utility theory in the evaluation of alternatives. Specifically, consumers who did not respond to the campaign (and that would be considered unprofitable) may indeed be worthy of solicitation if *in expectation* the amount they would contribute exceeds the cost of solicitation.⁷ Because of the stochastic nature of consumer responses the optimal decision is the action with the greatest utility in expectation. Regardless of whether or not a consumer responds in a given campaign, the example decision should be labeled as “solicit” if targeting a consumer with the corresponding data description would be profitable in expectation. Once examples are labeled with their decision-optimal labels, an error-based categorical classification model can be induced to estimate the decision-optimal mapping. Note that the optimal decision is never observed directly, even for consumers whose responses are known.⁸

Because the decision-making problem now is formulated as a classification task⁹ a standard active learning method for classification can be applied. In our evaluation we employ Uncertainty Sampling (discussed above; Lewis and Gale, 1994) to acquire consumers’ responses. We apply one modification to the original approach which has been shown to improve its performance substantially (Saar-Tsechansky and Provost 2004): examples are sampled from a distribution; an example’s probability of being sampled increases the closer its probability of being profitable is to 0.5.

3.2 A Response-Learning Approach: improving probabilities judiciously

As discussed, the alternative to Decision Learning is to concentrate on improving the estimations of response probability, and to apply decision theory to identify the optimal choice. However, if decisions are not always improved with more accurate probability estimations, traditional active learning strategies

⁷ Empirical results support this conclusion. In the data used below for the experiments, because of the low response rate, when using the naïve method almost all consumers are assessed to be unprofitable, and the corresponding decision is labeled “don’t solicit.” The resultant, induced model targets very few consumers and results in negative profits.

⁸ Using the notation introduced next, a consumer x_i is labeled as profitable if $\hat{f}_i \cdot U_i^S + (1 - \hat{f}_i) \cdot U_i^F - C \geq \Psi$.

⁹ This reformulation procedure is a generalization of the procedure used by MetaCost (Domingos, 1999) for cost-sensitive learning. See the Appendix for a comparison of the cost-sensitive learning setting and the decision-making setting.

may not be best—error-reducing active learning approaches may waste resources on acquisitions that reduce model error but produce little or no improvement in decision-making. We now derive the fundamentals for a new decision-centric active learning approach.

3.2.1 Theoretical Foundation of a Decision-Centric Acquisition Policy

Recall that our decision maker decides whether or not to initiate a business action, such as mailing a solicitation to a consumer. The decision maker would like to estimate the expected utility from the action. Let x_i be an example (e.g., the description of a consumer) and let f_i denote the (unknown) probability that the action with respect to x_i will be successful (e.g., consumer x_i will respond to the marketing campaign, or will renew her contract). Let us assume for the moment that the decision maker is considering this action in isolation, an assumption we will relax later, for example as with a single-offer targeted marketing campaign where the action would be to make the offer to each of a set of consumers. Given that the action is taken, let the utility of success and the utility of failure with respect to instance x_i be U_i^S and U_i^F respectively. Let the corresponding utility of inaction be Ψ_i . Finally let C denote the cost of action. To maximize utility, the action should be initiated if $f_i \cdot U_i^S + (1 - f_i) \cdot U_i^F - C \geq \Psi_i$, or equivalently, if the probability of a successful outcome exceeds the threshold given by

$$f_i^{Th} = \frac{\Psi_i + C - U_i^F}{U_i^S - U_i^F} \quad (1).$$

Since the true probability of success is unknown, in order for a decision maker to act optimally it is necessary to estimate f_i . Because training information is costly, we would like to reduce the cost of inducing an estimation model that will render decisions of a given quality. We have argued that simply reducing model error is not always a cost-effective strategy, and now we can be more specific. Consider the case in which the actual probability of success exceeds the threshold (suggesting action is better than inaction). For the induced model to allow a decision maker to act optimally it is sufficient that the estimated probability of success \hat{f}_i exceed the threshold as well, even if it were to be highly inaccurate.

Improvement of the probability estimation via information acquisition when the current estimation already specifies the correct decision would not affect decision-making, and therefore the cost of the improvement would be wasted. This observation suggests a complete deviation from the error-centric paradigm. In fact, if the true probability is only slightly above the threshold and the estimate has a non-negligible variance, improving the probability estimation by reducing the variance may adversely affect decision-making (cf., Friedman 1997), as we illustrate below.

A model is induced from a sample, so the model’s probability estimation \hat{f}_i can be treated as a random variable. Let Γ_i be the correct decision and let $\hat{\Gamma}_i$ be the decision derived using the model’s probability estimation. Similarly to Friedman’s analysis of incorrect classifications under 0/1 loss (Friedman, 1997), the probability of making a “wrong decision”—i.e., a decision that is inconsistent with the decision derived using the true probability of success—is given by:

$$P(\hat{\Gamma}_i \neq \Gamma_i) = I(f_i < f_i^{Th}) \int_{f_i^{Th}}^{\infty} p(\hat{f}_i) d\hat{f}_i + I(f_i \geq f_i^{Th}) \int_{-\infty}^{f_i^{Th}} p(\hat{f}_i) d\hat{f}_i \quad (2)$$

where the indicator function $I(L)$ is 1 if L is true and 0 otherwise.

In order to reduce the cost of building models that support accurate decisions, it is important to understand the circumstances under which costly improvements in probability estimation accuracy should be avoided. If we approximate $p(\hat{f}_i)$ with a normal distribution, the probability of making an

inconsistent decision is given by
$$P(\hat{\Gamma}_i \neq \Gamma_i) = \Phi \left[\text{sign}(f_i - f_i^{Th}) \frac{E\hat{f}_i - f_i^{Th}}{\sqrt{\text{var} \hat{f}_i}} \right] \quad (3)$$

where Φ denotes the right-hand-side tail of the standard normal distribution, E denotes an expectation and var denotes the variance of a random variable. Assume for illustration that a learner is used to induce a model from a training sample for estimating the probability that consumers would respond to a certain offer, and for a given consumer x_i the model produces (on average over different samples) a probability estimation \hat{f}_i such that the expected profit from an offer solicitation to x_i is higher than the expected profit of not making the offer, i.e., $\hat{f}_i > f_i^{Th}$. Also assume that the true probability of response

suggests the same ($f_i > f_i^{Th}$). So we expect the model to lead to the correct decision: make the offer. Under such circumstances it may not be cost-effective to acquire additional labels (consumer feedback) to improve the estimation; improving the estimation may increase the chance of decision-making error! Specifically, from (3) we see that indeed the larger the average probability estimation produced by the learner and hence the more extreme—and in the case in question the more inaccurate—the estimates are, the more likely it is that the model would produce the correct decision. This is because the overly extreme estimates reduce the chance that, due to estimation variance, the estimated expected profit from action would (mistakenly) fail to exceed that of inaction. There is an incentive, however, to remove such inaccuracy when the estimated probability and the true probability of response lead to different decisions.¹⁰ For example, assume that the true expected benefits from inaction exceed those of action, but that on average the learner induces a model that suggests otherwise. Thus, for cost-effective acquisition of training labels to support decision-making, one is well advised to take a decision-centric strategy, acquiring labels when an improvement in the estimation is likely to lead to higher expected utility, and avoiding acquisitions otherwise even if they might produce a more-accurate model.

The analysis remains essentially the same if we no longer consider a single action in isolation, but instead replace inaction with an alternative action with uncertain outcomes and the expected utility of this action, Ψ_i , can be estimated (potentially using another model to estimate the probabilities of the corresponding uncertain outcomes). The formulation would then have to include the corresponding costs of the different possible outcomes in the obvious fashion. However, for this paper we assume costly examples are acquired to *induce* one model at a time; we discuss this limitation further below.

In summary, ideally and in contrast with traditional error-centric methods, an active learning mechanism should incorporate (1) the costs of alternative courses of action and (2) the utilities derived from possible outcomes.

¹⁰ There is an incentive whenever the estimated probability \hat{f}_i and the true probability f_i are on different sides of the threshold f^{Th} .

3.2.2 Goal-Oriented Active Learning (GOAL)

Our analysis suggests that the most informative responses can be obtained for cases where the current model would lead to erroneous decisions. Unfortunately, this condition cannot be determined with certainty before responses are acquired. Consider, then, the following two scenarios concerning a decision as to whether to target particular consumers with an offer. In scenario A the estimated class probability of response is considerably higher than the threshold probability. In scenario B the estimated probability of response is only marginally greater than the threshold probability. In scenario A the evidence in the training data is strongly in favor of action—target the consumer. While one cannot determine with certainty the consumers for whom the current model would lead to erroneous targeting decisions, it is possible to employ a heuristic to capture this notion. Our response-learning policy assumes that a confident response probability estimation is also more likely to yield a correct decision, rendering additional acquisitions to improve the decision wasteful. This follows the motivation of uncertainty sampling (Lewis and Gale 1994) and most other active learning heuristics. More specifically, if the estimated probability of response is significantly larger or smaller than the probability for which the decision maker would be indifferent regarding two or more alternative actions, we will infer that there is substantial evidence in the data in favor of the corresponding action, and thus that the decision is more likely to be correct (and vice versa). Furthermore, when the current decision is erroneous, a more substantial change in the estimated probabilities is necessary to affect the decision in Scenario A as compared to Scenario B, requiring more training examples to sway the estimation in favor of inaction (all else being equal). Thus, the approach we propose here acquires labeled examples pertaining to decisions that are more likely to be erroneous and also are less costly to affect, i.e., decisions for which a relatively small change in the estimation can change the preference order of choice (when \hat{f}_i is closer to f_i^{Th}). Of course, although the design is guided by the theoretical development above, this is a heuristic method. Note also that this selection criterion generalizes the selection criterion of uncertainty sampling

(Lewis and Gale 1994). Specifically, with a particular (perhaps unlikely) choice of utilities,¹¹ this criterion reduces to the selection criterion of uncertainty sampling. Thus, the theoretical development just presented provides an alternative explanation for uncertainty sampling’s effectiveness in maximizing classification accuracy.

The new active learning technique we propose operates iteratively. At each phase, $n \geq 1$ examples are selected from the set of unlabeled examples UL ; their labels are acquired and the examples are added to the set of labeled training examples L . The estimated effectiveness of each potential acquisition is calculated as follows. Each example $x_i \in UL$ is assigned a score that is inversely proportional to the minimum absolute change in the probability estimation that would result in a decision different from the decision implied by the current estimation, i.e., the score of example x_i is inversely proportional to $|\hat{f}_i - f_i^{Th}|$.

For selection, as described in Section 3.1, rather than selecting the examples with the highest scores (as is common in active learning), a sampling distribution is created. Specifically, the effectiveness scores are considered to be weights on the examples, and examples are drawn from a distribution where the probability of an example to be selected for labeling is proportional to its weight. Incorporating a stochastic element in the selection of examples alleviates the danger of acquiring multiple identical examples (due to identical weights), helps to avoid selecting outliers, and generally is likely to choose examples from denser areas of the example space (all else being equal). Formally, the sampling-distribution weight $W(x)$ assigned to example $x_i \in UL$ is given by

$W(x_i) = \left[\kappa \cdot \left(\beta + |\hat{f}_i - f_i^{Th}| \right) \right]^{-1}$, where β is some small real number to avoid division by zero (in the empirical evaluation that follows $\beta=0.001$), and κ is a normalizing factor given by

¹¹ Specifically, $f_i^{Th} = 0.5$ if $(U_i^S + U_i^F) = 2(\Psi_i + C)$, i.e., the sum of the utilities of the outcomes of the main action is equal to twice the sum of the utility of the alternative action (in our case, of inaction) and the cost of the main action.

$\kappa = \sum_{i=1}^{\text{size}(UL)} \left(\beta + \left| \hat{f}_i - f_i^{Th} \right| \right)^{-1}$, to make W a distribution. Figure 1 presents pseudocode of the method, which we call Goal-Oriented Active Learning (GOAL).

Figure 1: The Goal-Oriented Active Learning (GOAL) Algorithm

Input: Set of unlabeled examples UL , initial set of labeled examples L , modeling technique I , stopping criteria, batch size n , cost of acquisition C , utilities U_i^f, U_i^s, Ψ_i for each example $x_i \in UL$,

While (stopping criteria not met)

- 1 Apply I to L , resulting in model M
- 2 Apply model M to UL
- 3 For each example $x_i \in UL$ compute $w(x_i) = \left[\kappa \cdot \left(\beta + \left| \hat{f}_i - f_i^{th} \right| \right) \right]^{-1}$
- 4 Sample from the probability distribution w , a subset S of n examples from UL without replacement
- 5 Remove S from UL , acquire labels for examples in S , and add them to L

End While

Output: Model M generated from L

The problem setting we employ here includes two alternative actions. However, the framework applies directly to a decision among n alternative actions if the expected utility from alternative actions can be estimated (possibly by using other predictive models). Specifically, assume consumer responses to action A are being modeled. Following GOAL's principles, it would be most cost effective to acquire responses for which it is likely that a small change in the consumers' estimated response probabilities to action A may alter the decision. In our analysis above we derive the likelihood of a decision change by comparing the expected utility from A to that of the alternative action and by computing a *distance* between the current distribution of uncertain outcomes and one that would lead to a decision change. The treatment in a multi-action (>2) scenario differs only by the choice of the relevant action whose expected utility should be compared to that derived for action A . A change in a consumer response estimation to A will change the preference order of choice if: (i) A is currently the action with the highest expected utility, but a change in the estimation of A 's outcome probabilities would lead to the action with the second-highest expected utility to be the action of choice; or if (ii) a change in the estimation of A 's outcome probabilities would result in the expected utility from A to exceed that of the action with the (current) highest expected utility. In (i) the relevant comparison is between A and the second-best action, whereas in case (ii) the relevant comparison is between A and the action currently estimated to have the highest expected utility.

4. A Direct Marketing Campaign Case Study

Our empirical evaluation employs data from a direct-marketing campaign, which comprise real consumer interactions and both information-acquisition costs and utilities from successful targeting. The data pertain to a charity’s periodic solicitations to potential donors; they were the data for the KDDCUP competition in 1998 (Zadrozny and Elkan, 2001; Wang et al., 2003) and are now publicly available (Blake et al., 1998). For these data, each solicitation costs 68 cents and the average response amount is \$15. The response rate is approximately 5%. Because of the low response rate and the cost of solicitation, informed decisions that minimize wasteful solicitations are critical to the success of the campaigns. Building the predictive models used to generate targeting decisions requires costly acquisitions of consumer responses. For cost-effective utilization of the charity’s donor base, it is important to identify the most informative solicitations for training the model so as to improve future targeting decisions.

4.1 Acquisition Strategies for the Direct Marketing Problem

Given an acquisition budget, an acquisition strategy solicits potential donors in a training phase and acquires their responses (e.g., whether or not they responded to the solicitation, and if so, in what amount). These form the labeled training data. Once the labels are acquired, a targeting model is induced from the training data and subsequently is employed in the main campaign to target potential donors. In the main campaign, a successful solicitation is one that results in a contribution that exceeds the solicitation cost. Therefore, the training objective is to reduce the total acquisition cost of training examples necessary to achieve a particular level of campaign profit, or alternatively to increase the campaign profit for a particular acquisition investment.

For Decision Learning we evaluate two acquisition strategies:

- (1) As discussed above, Decision Learning casts the decision problem as a classification task and we employ Uncertainty Sampling to identify useful response acquisitions.

(2) As an alternative to active learning we acquire responses from a set of donors sampled uniformly at random. Uniform random sampling (URS) is the most widely applied practice for the acquisition of labels based on a set of unlabeled training examples. In spite of its simplicity of implementation, URS is remarkably effective because it attempts (implicitly) to obtain a representative sample.

For Response Learning we evaluate three label-acquisition strategies:

(3) Acquisition of responses from a representative set of donors, using URS.

(4) An error-centric acquisition (ECA) method that focuses on the model's average error reduction.

Because probability estimates must be used to evaluate the expected profitability of alternative solicitations, an acquisition strategy that improves these estimations may also be effective for improving targeting decisions. To our knowledge BOOTSTRAP-LV (Saar-Tsechansky and Provost, 2004) is the only generic method designed specifically to reduce class-probability-estimation error, rather than classification error. BOOTSTRAP-LV follows the traditional error-centric paradigm, acquiring examples that reduce the average probability estimation error. Specifically, the method estimates the variance in learned models' response probabilities, for each potential example, and assigns a higher preference to the acquisition of responses from examples with higher variance.¹² In the experiments described here the number of bootstrap-samples used in BOOTSTRAP-LV is $k=10$, and performance does not vary significantly for larger values of k .

(5) GOAL. Let the estimated probability that a potential donor x_i would respond to a mailing be \hat{f}_i , the estimated contribution amount (described below) be \hat{U}_i^s , and the mailing cost be C . The profit from inaction is zero; hence a solicitation is initiated if $\hat{f}_i \cdot \hat{U}_i^s - C \geq 0$ and the threshold probability

¹² An associated technical report (Maytal Saar-Tsechansky and Provost, 2005) provides a comparison to other error-centric policies with these data. The study demonstrates that Bootstrap-LV results in the highest profits on average and is therefore the error-centric policy we chose for this study.

is $f_i^{Th} = \frac{C}{\hat{U}_i^S}$. Therefore, the weight assigned to acquiring donor i 's response in GOAL is given by

$$\left[k \left(\beta + \left| \hat{f}_i - f_i^{Th} \right| \right) \right]^{-1} = \left[k \left(\beta + \left| \hat{f}_i - \frac{C}{\hat{U}_i^S} \right| \right) \right]^{-1}.$$

4.2 Experimental Setting

In order to evaluate the five acquisition strategies, we compare the profits generated from solicitation decisions derived from the models induced with each. We now describe the induction methods examined, the data partitioning, and the method for calculating generated profits.

For estimating the probability of response, we employ Probability Estimation Trees (PETs)—unpruned C4.5 classification trees (Quinlan, 1993) for which the Laplace correction is applied at the leaves (Provost and Domingos 2003; Perlich et al. 2003); as discussed below similar results are obtained for other induction techniques. For this application, revenues from successful solicitations are not known in advance and therefore also must be estimated. To do so, we use a linear regression model induced from historical data on past donations.¹³ For the Decision-Learning approach we use all the predictors employed above to estimate the probability of response and the contribution amounts. In addition, we include the targeting costs, the estimated response probability, and the estimated contribution amounts as independent variables for the Decision-Learning approach. The response probabilities and contribution amounts are estimated from the corresponding available training data for each point on the learning curve. To minimize experimental variance we also employ Probability Estimation Trees for Decision Learning.

¹³ We employ predictors that were identified in earlier studies to be predictive of donation amounts. The predictors are: the amount of the most recent gift, the number of donations per solicitation, average donation amount in response to the last 22 promotions, and the estimated probability of donation as estimated by the model. Following Zadrozny and Elkan (2001) the probability estimation is incorporated as a predictor in the linear regression model to remove a sample selection bias. Because large gifts are rare there exists a selection bias towards one group of frequent donors who donate small amounts resulting in the regression model underestimating gifts by donors who contribute large amounts infrequently. To alleviate such a bias, Heckman (1979) recommends incorporating the probability of belonging to either group (i.e., the probability of making a donation) as a predictor in the regression model.

On a separate (holdout) set of donors, we compare the profits generated by each method for an increasing number of acquired responses. More specifically, at each acquisition phase the responses of M additional donors are acquired by each method and added to its respective training set. Each point on each curve shown hereafter is an average over 10 independent experiments. For each experiment, the data are randomly partitioned into: an initial set L of labeled training examples selected at random (used to build the first model); an unlabeled pool UL of donors from which each strategy will acquire additional responses, which then are added to L (cumulatively as the active learning progresses); and an independent test set T of potential donors whose responses and donations are known to the experimenters, for evaluating the strategies. To reduce experimental variance, the same data partitions are used by all methods.

The profit is calculated via the simulated process depicted in Figure 2 (recall that responses are known to the experimenters for the entire test set). For the Response-Learning approach, for each potential donor in the test set, a solicitation is mailed if the estimated expected revenue exceeds the solicitation cost. For the Decision-Learning approach, a solicitation is mailed if the profitability model estimates that the consumer is likely to be profitable (with probability ≥ 0.5). In both approaches, the cost of mailing is subtracted from the total profit whenever a solicitation is made; if a donor responds to a solicitation the actual donated amount is added to the overall profit.




Figure 2: Decision-making profitability calculation from charity solicitations

4.3 Results

Since we are not aware of such an experiment being carried out previously, first we show that Decision Learning is inferior to Response Learning when acquiring labels uniformly at random. Figure 3 plots the profit generated with each approach for an increasing number of training data. The curves clearly show that the estimation of response probabilities consistently results in significantly better decisions. Thus, while the decision task can in principle be formulated as a classification task, targeting decisions are less accurate on average when they are induced from the data, as compared to when they are derived directly from expected utility theory.




Figure 3: Profits from targeted solicitations. Targeting decisions are derived with Decision Learning and Response Learning for an increasing number of responses used for training.

Why is it advantageous to induce consumers’ response probabilities and to apply a policy derived from expected utility theory, rather than to learn the final targeting policy completely from data? First, note that both approaches have access to the same estimations for the probability of response \hat{f}_i and the utilities from different outcomes. The optimal targeting rule $EU(\hat{f}_i, U_i^S, U_i^F)$ comes directly from Von Neumann-Morgenstern expected utility theory and determines whether a consumer with an estimated probability of response \hat{f}_i , a contribution U_i^S , and a targeting cost U_i^F is profitable in expectation.

Response Learning estimates \hat{f}_i from the data and the optimal targeting rule $EU(\hat{f}_i, U_i^S, U_i^F)$ is applied to infer each consumer’s profitability. By contrast, with Decision Learning $EU(\hat{f}_i, U_i^S, U_i^F)$ itself must be induced from the data. Thus, unless (i) Decision Learning implicitly learns a better \hat{f}_i than learning it directly (which seems unlikely), or (ii) induction of $EU(\hat{f}_i, U_i^S, U_i^F)$ is perfect, then targeting decisions based on Decision Learning’s *induced* mapping $\hat{EU}(\hat{f}_i, U_i^S, U_i^F)$ will necessarily be inferior to decisions derived from Response Learning. Indeed, as shown by the initially increasing profit curve of Decision Learning in Figure 3, the expected utility mapping induced from the data improves with more training examples. However, because $EU(\hat{f}_i, U_i^S, U_i^F)$ is known exactly and can be readily deployed, there is no advantage in imperfectly relearning this knowledge via induction.

Figure 4 shows the average acquisition selection times and the profits obtained from targeting decisions using Response Learning (RL) when labeled training examples are acquired with GOAL, Error-Centric Acquisition (ECA), and Uniform Random Sampling (URS), as well as using Decision Learning (DL) when labeled training examples are acquired with Uncertainty Sampling (US). Profitability is plotted for increasing costs of response acquisitions. Initially, each method received (the same) 2000 training examples. At each phase ten responses and in total 5000 responses were acquired actively by each method. Note that initially all acquisition methods have access to the same, small set of responses, and therefore, the same consumer response model is induced by all RL methods resulting in the same performance (DL learns a different model, of course). As additional donors’ responses are acquired by each method, the sets of responses available for training begin to differ in composition, resulting in different learned models.

When Uncertainty Sampling is applied to acquire informative responses for Decision Learning, the resultant profits do not surpass the performance obtained with GOAL, nor do they surpass profits of Response Learning when responses are acquired at random. This demonstrates that active acquisition

using Uncertainty Sampling does not overcome Decision Learning’s handicap, its need to induce the optimal decision-making rule (imperfectly) from the data.




Figure 4: Profits generated and selection times when responses are acquired via GOAL, Error-centric acquisition (ECA), uniform random sampling (URS) for Response Learning (RL), and when responses are acquired via Uncertainty Sampling (US) for Decision Learning (DL).

For Response Learning we proposed three alternative response-acquisition policies that aim to improve decision-making for a given budget: GOAL, error-centric acquisition, and acquisitions drawn uniformly at random. As shown in Figure 4, GOAL improves the profits (on average) more than the other methods. For a given cost, GOAL obtains higher profits from better targeting decisions and is significantly more efficient computationally when compared with ECA.¹⁴ GOAL also generates higher profits as compared to when responses are acquired uniformly at random. Similarly, URS acquisitions are clearly inferior to those obtained with ECA.

As Figure 4 shows, as more responses are acquired and the composition of the training sets diverges, the relative advantage of GOAL becomes more apparent. GOAL produces statistically significantly higher profits than URS according to the Wilcoxon signed rank test ($p < 0.05$) once 10

¹⁴ Computations were performed on a Dell Precision Workstation with two Pentium P4 2.0 GHz CPUs and 512MB main memory, running Linux.

responses are acquired and onward, and it obtains statistically superior profits than ECA once 180 responses are acquired and until 2500 responses are acquired.

The difficulty of improving targeting decisions and profitability is well demonstrated by the number of acquisitions required in order to obtain a given improvement in profits. For example, ECA must acquire more than 3500 responses in order to obtain a 1.5% increase in profits; GOAL must acquire only about 600 responses, less than one-fifth of the responses required by ECA, to exhibit the same improvement in performance. The largest improvements are obtained in the early acquisition phases, where after only 400 responses are acquired by each method GOAL results in 1.8% higher profits on average compared to ECA. Qualitatively and quantitatively similar results are obtained when assessing decision-making error rather than profit, and in particular when separating the targeting of profitable donors and unprofitable donors (see Saar-Tsechansky and Provost, 2005).

These results provide strong evidence that GOAL effectively identifies acquisitions that can improve decision-making. Two possible explanations are (1) that the advantages conferred by GOAL are attributed to its ability to avoid wasteful improvements in probability estimation, or (2) because it actually does a better job of improving probability estimation for a given budget. Figure 5 compares the error of the probability estimates produced by GOAL with those generated by ECA (as always, on out-of-sample test sets). Probability estimation accuracy is measured with *BMAE* (Best-estimate Mean Absolute Error), computed as $BMAE = \frac{\sum_{i=1}^N |\hat{p}_{best}(x_i) - \hat{p}(x_i)|}{N}$, where $\hat{p}(x_i)$ is the probability estimated by the model under evaluation (and that was induced from the selected subset of the available examples); N is the number of test examples for which the models are evaluated; \hat{p}_{best} is a surrogate to the best estimated probability and is estimated by a “best” model induced using the entire set of available (labeled) examples $L \cup UL$ (and using a more complicated modeling approach, a Bagged-PET which generally produces superior probability estimations (Provost and Domingos, 2003; Perlich et al. 2003)).




Figure 5: Comparison of probability estimation error

In contrast to the pattern shown in Figure 4, on average the probability estimations obtained with GOAL for a given acquisition cost are considerably worse than those obtained with ECA. ECA’s improvements are statistically significant according to the Wilcoxon test ($p < 0.05$) after both strategies have acquired 500 examples.

Figure 4 and Figure 5 demonstrate that GOAL leads to improved decision-making for a given cost, while the average probability estimations it produces are inferior. As the discussion in Section 3.2 suggests, some improvements in probability estimation accuracy may not improve decision-making. Such improvements are wasteful since they are costly to achieve. In particular, costly accuracy-improving acquisitions will be cost-effective only if they lead to different targeting decisions. GOAL and ECA differ in the manner by which they value response acquisitions: ECA prefers acquisitions estimated to improve response probability estimation the most, regardless of its impact (or the lack thereof) on decision-making. GOAL, by contrast, promotes improvements in response probability estimation based on their estimated potential impact on improving decisions, regardless of the subsequent improvement in the model’s average probability estimation. Thus, for example, GOAL may prefer a small improvement in the average probability estimation over a large improvement if the former is more likely to improve decisions and profitability. By contrast, ECA will always acquire responses estimated to improve the probability estimation model the most, even if the acquisitions are not likely to improve decisions. Furthermore, as our analysis suggests, ECA may acquire information regarding already correct decisions

that reduces the estimation error but that increases the likelihood of incorrect decisions due to estimation variance. Thus ECA may (1) forgo acquisitions that lead to better decisions, and (2) acquire data that leads to worse decisions. Apparently, GOAL often avoids these pitfalls.

Finally, GOAL is generic: it does not make any assumption about the form of the model or about the induction algorithm. Hence, it can be applied with any model for estimating class probabilities. Whether it will be effective for various models must be assessed empirically. It produces qualitatively and quantitatively similar results for logistic regression and for naïve Bayes (Saar-Tsechansky and Provost, 2005).

4.4 The Effect of Error in the Estimation of Utilities on GOAL's Performance

To identify useful acquisitions GOAL employs information about the decision task: knowledge of the possible actions, the corresponding uncertainties, and the utilities of different outcomes. In some cases, uncertain utility-producing events must be estimated, such as the contribution amounts in the direct marketing task. Error-centric acquisition and random sampling do not base their acquisition decisions on these utilities, and thus these alternative policies are not adversely affected by the inaccurate estimations. We now examine the impact of the quality of estimation on GOAL's relative advantage.

To simulate estimations of varying quality, for each consumer we assume estimations of response amounts are drawn from a Gaussian distribution with a mean equal to the true amount and a standard deviation that increases to produce increasingly inaccurate estimates. The shapes of the resultant distributions are depicted in Figure 6 (for a consumer whose true contribution is \$15). As the standard deviation increases the probability that the estimated contribution is within $\pm \$2$ of the true value drops from 0.97 to 0.19 (for standard deviations of 1 and 10, respectively). Note, that a standard deviation of 2, which leads to a probability of 0.77 that the estimation is within $\pm \$2$ of the contribution, is already worse than the simple linear estimation model we employ, which leads to a standard deviation of 1. As before, we measure the marketing campaign's profitability in order to evaluate a policy's efficacy at improving the campaign's targeting decisions. Note that the accuracy of contribution

estimation can affect campaign profitability via two separate processes. It can affect GOAL’s response acquisitions—the phenomenon we aim to characterize here. It also can affect the targeting decisions themselves: for a given model, different decisions can be derived for different contribution estimations. For experimental control, we isolate the effect of estimation accuracy on GOAL’s response acquisitions. In particular, we vary the accuracy of the estimations during the acquisition phase to gauge its effect on acquisition efficacy, and we use the exact contributions for targeting. Thus, observed changes in profitability as the estimation degrades can be attributed only to GOAL’s response-acquisition policy.

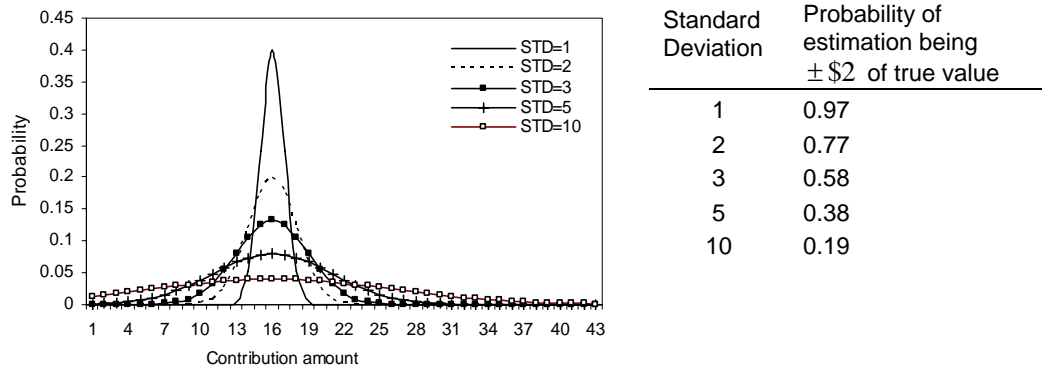



Figure 6: The distributions from which contribution estimations are drawn when the true contribution is \$15. As the standard deviation increases, the probability that the estimation is $\pm \$2$ from the true value drops from 0.97 to 0.19.



- a.** GOAL’s acquisitions converge to those of random sampling as the quality of the estimation degrades.
- b.** Error-Centric Acquisitions produce better results when the estimation is very poor.
- c.** GOAL’s advantage as compared to Decision Learning is most significant when the utility estimation is accurate.

Figure 7: Increase in profits using GOAL versus alternative policies, as a function of estimation error. Vertical axis shows percentage differences in profit of GOAL versus the alternative.

Figure 7a shows the improvement in profits generated by models induced with GOAL over those generated with URS (after 1000 acquisitions), as a function of the accuracy of contributions’ estimations. Differences in profits are shown as a percentage of the profit obtained with random acquisitions. The

relationship exhibits two desirable properties. First, GOAL is robust: it generates profits that are equivalent or better than those produced with URS for all levels of estimation error. Second, as the estimation error increases, the estimated contributions are drawn from a distribution that is less concentrated around the true mean, and more closely resembles a uniform distribution. As a result, GOAL’s performance converges to that of URS.

Figure 7b shows the improvement in profits generated with GOAL over those generated with Error-Centric Acquisition. GOAL’s relative advantage with respect to the error-centric policy shows a similar pattern: it is not very sensitive to estimation error and is advantageous even when the estimation is moderately poor (the probability that the contribution is within $\pm \$2$ of the true value being greater than or equal to 0.58). Only when the task description provided to GOAL is highly noisy and GOAL’s performance approaches that of a uniform acquisition policy, does the error-centric policy become preferable (recall that URS is inferior to ECA). Our prior analysis showed that GOAL’s ability to acquire information more economically than ECA stems from the fact that it selectively acquires responses only if they are likely to improve decisions. However, when the decision task is misrepresented by erroneous estimations, GOAL is less likely to identify truly useful acquisitions. By contrast, error-centric acquisitions consistently improve the average probability estimation.

Figure 7(c) shows how GOAL’s relative advantage with respect to Decision Learning varies with respect to the quality of the utility estimation. GOAL’s relative advantage is most significant when the utility estimation is most accurate, but it is also preferable when the estimation variance increases. As we note above, both approaches have access to the estimations for the probability of response \hat{f}_i and the utilities from different outcomes. These results show that although GOAL is indeed sensitive to errors in the utility estimation, GOAL’s advantage remains significant.

4.5 Non-uniform Response Acquisition Costs

GOAL aims to capture the potential impact of each response acquisition on future decisions. This formulation of the problem so far has assumed that the same cost is incurred to acquire each

response. If response acquisition costs were non-uniform, it would be useful to forgo an informative response if the same improvement could be obtained for a lower cost via a different acquisition. More generally, an acquisition would be preferable if it were expected to produce a larger improvement in profit per unit cost. To identify the optimal acquisition schedule it would be necessary to estimate the expected increase in profit from each potential acquisition and to acquire responses for which the expected improvement per unit cost is the highest. We will now provide a brief demonstration.

There are considerable challenges in obtaining a meaningful estimation of the expected increase in profitability. For example, it is necessary to estimate profitability on data representative of the target population. However, active learning acquires a sample that often is not representative of the population to which the model would be ultimately applied. While, as we show here, model induction benefits from such a sample, unless the sample is unbiased (Mookerjee, 2001) our ability to infer certain characteristics of the target population accurately may be limited. In particular, an estimation of expected performance on a biased sample can produce highly inaccurate estimations of performance on the target population (Baram et al., 2004). Thus, for this paper we use the heuristic measure employed by GOAL to capture the value of each potential acquisition as a proxy for the expected improvement in profit. To penalize the value of potential acquisitions for their cost, weights assigned by GOAL are normalized by the cost of acquisition to reflect the expected value per unit cost. We therefore prefer the acquisition of responses whose marginal impact as captured by GOAL, per unit cost, is the highest. We call this extension to GOAL, which incorporates acquisition costs, GOAL-AC.




Figure 8: Profits generated with GOAL-AC, GOAL, ECA, and URS where acquisitions costs are drawn from a uniform distribution with range [1, 5].

To evaluate the performance of GOAL-AC with non-uniform acquisition costs we assume acquisition costs are drawn from a uniform distribution with range [1, 5]. Figure 8 shows the profits generated with GOAL-AC, GOAL, URS and ECA for acquisitions of varying costs.¹⁵ As shown, responses obtained by GOAL-AC for a given cost are more informative on average and produce better decisions for a given acquisition budget. Higher acquisitions costs also lead to decreased profitability for the URS policy.¹⁶ As compared to GOAL, GOAL-AC effectively integrates information on the cost of each potential acquisition and on the expected impact of responses on future decisions in order to prioritize acquisitions.¹⁷

5. Limitations and Future Work

One main contribution of this paper is the presentation of a framework for decision-centric active learning. GOAL directs acquisitions to improve the learning of a single predictive model capturing the probability distribution of uncertain outcomes to a given action. These acquisitions could be in a decision context including other predictive models, possibly learned from data. For example, in our

¹⁵ Because the acquisitions of different responses incur different costs, the same acquisition costs are not incurred by all acquisition strategies. Therefore we plot based on the costs incurred by URS and show profits for other strategies for these costs by linearly interpolating profits from known (profit, cost) pairs obtained for higher and lower acquisition costs.

¹⁶ The URS policy increasingly targets more donors that are likely to donate relatively larger amounts. However, low response rates along with higher targeting costs lead to a decrease in revenue.

¹⁷ Note that because the simulated acquisition costs are higher on average than the true costs, lower profits are generated by all policies as compared to the respective profits shown in previous sections.

introductory example we discussed learning a model of response to an incentive for the purpose of controlling customer attrition. Such a decision process would likely also include a model for predicting the probability of customer attrition. However, we do not consider the development of policies for scheduling acquisitions for learning multiple models simultaneously. Decision-centric active learning involving not only actively selecting among data to build a model, but also (simultaneously) actively selecting from among data for building different models would require significant advances over what we have presented here, because we do not directly estimate the expected utility of the acquisition—so there is little basis for comparisons between acquisitions for different models. We hope that this limitation is not severe, because (at least in our experience) currently data mining campaigns for different models are undertaken separately. In our example, it is likely that a model predicting the likelihood of customer attrition already has been built and tested when the model for a new incentive is to be induced.

Our analysis suggests that the more informative acquisitions are those that are more likely to affect the ultimate decision-making. The manifestation of this principle in GOAL is the heuristic measure of the “mobility” of the decision. We address only binary uncertain outcomes, for which GOAL can calculate the minimum “distance” between the current estimated probability distribution and one that would lead to a decision change. When actions can lead to multiple (>2) uncertain outcomes, it is not clear what measure can be used to capture the “distance” between the current distribution and those which would alter the decision. In such setting, a comparison between the actual expected utility from action A and that of the relevant alternative may be a suitable measure of the mobility of the decision (Melville, et al., 2005). Indeed this also would address the aforementioned limitation regarding scheduling acquisitions for more than one model.

We show that GOAL’s effectiveness is undermined by highly inaccurate decision parameters. It is conceivable that in some domains information about the decision task, such as the utilities from various outcomes, can be obtained at a cost. This scenario suggests opportunities to develop acquisition policies that not only acquire information directly for model induction, but also (or instead) identify information

that can improve the performance of another acquisition strategy. In the context of GOAL, such a policy could suggest how to acquire information cost-effectively about consumers' likely response amounts so as to support GOAL's acquisitions of responses. For example, a policy could target acquisitions that would improve simultaneously the models for estimating response probability and response amount. The successful development of such policies may be dependent on the ability to estimate accurately the expected increase in profits per unit cost. However, the less directly related to profitability the impact of the information acquired is, the more challenging it will be to derive effective heuristic measures of the benefit of costly information to profitability.

Finally, the employment of active interactions with consumers also gives rise to a new challenge of balancing between the benefits of making offers or recommendations to consumers intended to increase sales at present versus informative interactions that can benefit learning. The tradeoff between activities initiated to support learning and those that exploit what is already known has been explored in the robotics literature; a similar framework may prove useful for business as well (Pednault et al., 2002).

6. Conclusion and Implications

Because the information required for effective predictive modeling often is costly to obtain, it is beneficial to devise mechanisms to direct the acquisition of data for cost-effective improvements to decision-making. The main contributions of this paper are the formulation of the decision-centric active learning problem, the development of decision-centric active learning techniques, and the demonstration via rigorous experimentation that a well-designed technique can improve substantially over alternative approaches. The framework can be adapted for developing strategies and techniques to address other decision-making settings.

GOAL's acquisition method is derived from theoretical observations regarding the conditions under which class probability estimation error is more likely to undermine decision-making. When applied to a direct-marketing problem, GOAL's decision-centric strategy identifies acquisitions that significantly improve models for targeting, as measured by campaign profitability. Examining the

relationship between error reduction and decision-making efficacy reveals that the economies exhibited by GOAL indeed are derived from acquiring labels that will affect solicitation decisions, sometimes at the expense of overall probability-estimation error reduction. An additional advantage of GOAL is that it does not require much computation.

This paper introduces the notion that decision-centric, active acquisition of modeling information can lead to more profitable model building and application. We have motivated the use of such techniques with examples of modeling consumer preferences and loyalty, which recently have been high-profile modeling applications. These are applications for which the acquisition of labels for training data carries clear costs. There are many other uses of predictive modeling in business, many of which have associated data acquisition costs. A further implication of this work is that decision-centric data acquisition strategies should be considered elsewhere as well.¹³

The notion of decision-centric *active* information acquisition also suggests that businesses should consider modifying their strategies for acquiring information through normal business transactions. A firm, such as Amazon.com, that models consumer preferences for customized marketing can accelerate learning by proactively offering recommendations—not merely to induce immediate sales, but for improving recommendations in the future. The decision-centric acquisition approach presented here suggests that such active acquisition of information may well result in better decisions in the future. For firms such as Amazon.com, such capacity potentially could be employed to accelerate induction of consumer preference models and to offer more accurate and effective recommendations, earlier.

Acknowledgments

We thank our editors Sumit Sarkar and Giri Kumar Tayi and our anonymous reviewers for their insightful, substantive comments. We also thank the organizers of The Second International Knowledge

¹³ From the research perspective there is a striking lack of publicly available data sets that include all the cost/benefit information required to evaluate methods such as this. This should be less of a limitation from a managerial point of view, where sources of both data and cost/benefit information should be available.

Discovery and Data Mining Tools Competition (KDD Cup 1998) and a national veterans' organization for the data set used in this paper.

This research was sponsored in part by the Air Force Research Laboratory, Air Force Materiel Command, USAF, under Agreement number F30602-01-2-0585. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL or the U.S. Government.

Appendix: Cost-Sensitive Learning vs. Decision-Centric Learning

The topic of study that has come to be known as “cost-sensitive learning” (CSL) addresses the task of inducing effective classification models given a set of known misclassification (error) costs. CSL is a special case of optimal decision-making, and it bears similarities to, but also some important distinctions from, the setting we study here. It is useful to outline this relationship in order to highlight the potential application of GOAL to CSL problems. CSL pertains to the prediction of an uncertain and discrete outcome/class: CSL aims to infer a model for predicting that instance x belongs to class i that minimizes the conditional expected loss: $L(i|x) = \sum_j p(j|x)C(i,j)$, $i \neq j$, where $p(j|x)$ is the probability that instance x belongs to class j , and $C(i,j)$ is the cost of classifying an instance to class i when it really belongs to class j .

In contrast, optimal decision-making pertains to general choice problems, where a (possibly costly) action is taken, after which the uncertainties (outcomes) are played out and the utilities collected. Different actions may incur different costs and an action does not necessarily correspond to the prediction that an instance belongs to a particular class. While traditional CSL defines a symmetric decision problem where the same set of outcomes may be played out after each action, in a general decision-making problem the alternative actions can lead to different sets of uncertainties. For example, when a consumer is targeted in a marketing campaign, a cost is incurred, and the subsequent uncertainties correspond to consumers' responses such as a purchase or a decline of the offer. If the consumer is not targeted, however, no cost is incurred, and there are no uncertainties that follow.

Optimal decision-making also must consider the benefits from accurate decisions, which are normally assumed in CSL to be zero.

Thus CSL constitutes a special case of optimal decision-making, with some unique properties that do not carry over to the more general case. To demonstrate the possible ramifications of the differences between the two tasks, consider the following result for CSL derived for binary class problems. When the misclassification costs in CSL are uniform, one would be indifferent between different actions (assigning an instance to either class) if the outcome probabilities are 0.5. This result does not carry over to a general decision-making setting. For example, if the cost of targeting a consumer who does not respond is equal to the benefit generated when the consumer accepts an offer, the decision maker will not be indifferent between actions when the probability of response is 0.5. This is because targeting a consumer is costly while the alternative is not, and the possible outcomes of a targeting decision are different from those following a decision not to target the consumer. Indeed, the decision maker is only indifferent if the consumer's probability of response is 1, as the cost of targeting a consumer can be recovered in expectation only if the consumer responds with certainty.

The above discussion suggests that a method for decision-centric active learning should also apply to special cases of decision-making, such as cost-sensitive learning (with unequal or equal misclassification costs). As such, our framework extends acquisition policies to address decision-making and classification with varying misclassification costs.

Note also that error-centric policies are not necessarily “cost-insensitive.” For example, one method for addressing a CSL problem is to learn an accurate class-probability estimation model to allow one to estimate the expected utilities from various classifications. The error-centric acquisition policy we use is of this sort—error-centric but still applicable to cost-sensitive and decision-making tasks.

References

- Baram, Y., El Yaniv, R., and Luz, K. 2004. Online Choice of Active Learning Algorithms. *Journal of Machine Learning Research*, 5(Mar):255--291.
- Berry, M. J., and Linoff, G. S. 2004. *Data Mining Techniques*. Wiley, John & Sons.
- Blake, C. L. and Merz, C. J. 1998. UCI Repository of machine learning databases. Irvine, CA: University of California, Department of Information and Computer Science, [<http://www.ics.uci.edu/~mlearn/MLRepository.html>].
- Cohn, D., Atlas, L., and Ladner, R. 1994. Improved generalization with active learning. *Machine Learning*, 15:201-221.
- Cortes, C., Jackel, L.D., Solla, S.A. Vapnik, V. and Denker, J.S.. 1994. Learning curves: asymptotic values and rate of convergence. *Advances in Neural Information Processing Systems*, Volume 6, 327–334.
- Domingos, P. 1999. MetaCost: A General Method for Making Classifiers Cost-Sensitive. In *Proceedings of the Fifth International Conference on Knowledge Discovery and Data Mining (KDD-99)*, 155-164.
- Fedorov, V. 1972. Theory of optimal experiments. Academic Press.
- Friedman, J. H. 1997. On bias, variance, 0/1-loss, and the curse-of-dimensionality. *Knowledge Discovery and Data Mining*, 1:55-77.
- Hastie, T., Tibshirani R., and Friedman, J. 2001. *The Elements of Statistical Learning*. New York: Springer-Verlag.
- Heckman, J. J. 1979. Sample Selection Bias as a Specification Error. *Econometrica* 47(1): 153-161.
- Hevner, A., March, S., Park, J., and Ram, S. 2004. Design Science in Information Systems Research. *MIS Quarterly* 28(1): 75-105.
- Kiefer, J. 1959. Optimal experimental designs. *J. R. Stat. Soc., series B* 21: 272-- 304.

- Kornish, Laura J., 2006, "Technology Choice and Timing with Positive Network Effects," *European Journal of Operational Research*, 173(1), 268-282.
- Lewis, D., and Gale, W. 1994. Training text classifiers by uncertainty sampling. In Proc. of Int'l ACM Conf. on Research and Development in Information Retrieval (SIGIR-94), pages 3--12.
- McCardle, K. 1985. Information acquisition and the adoption of new technology. *Management Science*, 31(11): 1372-1389.
- Melville, P., Saar-Tsechansky, M., Provost, F. and Mooney, R.J. 2005. "An Expected Utility Approach to Active Feature-value Acquisition." In *Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM-2005)*, pp. 483-486.
- Mitchell, T. M. 1997. *Machine Learning*. WCB/McGraw Hill.
- Moe, W. and Fader, P. S. 2004. Dynamic Conversion Behavior at e-Commerce Sites. *Management Science* 50 (3), 326-335.
- Mookerjee, V. S. 2001. Debiasing Training Data for Inductive Expert System Construction. *IEEE Transactions on Knowledge and Data Engineering* 13(3), 497-512.
- Pednault, E., Abe, N. Zadrozny, B., Wang, H., Fan, W. and Apte, C. 2002. Sequential Cost Sensitive Decision Making with Reinforcement Learning. In Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 204–213.
- Perlich, C., Provost F., and Simonoff, J. S. 2003. Tree induction vs. logistic regression: A learning-curve analysis. *Journal of Machine Learning Research* 4 (2003) 211-255.
- Provost, F. and Domingos, P. 2003. Tree induction for probability-based ranking. *Machine Learning*, 52:3, 199-215.
- Quinlan, J. R. 1993. *C4.5: Programs for machine learning*. Morgan Kaufman, San Mateo, California.

- Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin American Mathematical Society*, 55:527-535.
- Rosset, S., Neumann, E., Eick, U., Vatik, N. 2003. Lifetime Value Models for Decision Support. *Data Mining and Knowledge Discovery* 7: 321-339.
- Saar-Tsechansky, M. and Provost, F. 2005. Active Learning for Decision Making. Working Paper CeDER-04-06, Stern School of Business, New York University, NY, NY 10012
- Saar-Tsechansky, M., and Provost, F. 2004. Active sampling for class probability estimation and ranking. *Machine Learning* 54(2): 153-178.
- Seung, H. S., Opper, and Smopolinsky, H. 1992. Query by committee. In *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*, 287-294.
- Simon H. A., and Lea, G. 1974. Problem solving and rule induction: A unified view. In L.W. Gregg (ed.), *Knowledge and Cognition*. Chap. 5. Potomac, MD: Erlbaum.
- Von Neumann, J. and Morgenstern, O. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- Wang, K., Zhou, S., Yang, Q. and Yeung, J. M. S. 2003. Mining customer value: from association rules to direct marketing. In *Proceedings of the IEEE International Conference on Data Engineering*.
- West, P. M., Brockett, P. L., and Golden, L. L. 1997. A Comparative Analysis of Neural Networks and Statistical Methods for Predicting Consumer Choice. *Marketing Science*, 16 (4), 370-391.
- Winston, P. H. 1975. Learning structural descriptions from examples. *The Psychology of Computer Vision*, P. H. Winston (ed.), McGraw-Hill, New York.
- Zadrozny, B. and Elkan, C. 2001. Learning and making decisions when costs and probabilities are both unknown. In *Proc. of the 7th Intl. ACM SIGKDD Conf. on Knowledge Discovery and Data Mining*. 204-212.