

Toward optimal allocation of human resources for active learning with application to safe on-line advertising

Josh Attenberg
Polytechnic Institute of NYU
Brooklyn, NY 11201
josh@cis.poly.edu

Foster Provost
NYU Stern School of Business
New York, NY 10012
fprovost@stern.nyu.edu

ABSTRACT

This paper provides an empirical analysis of alternative techniques for incorporating human resources into the data mining process, in particular, for data acquisition for classifier induction in domains exhibiting extreme class imbalance. The approaches analyzed were motivated by the problem of applying data mining to build systems for “safe” online advertising: helping advertisers to control the content adjacent to their advertisements. Objectionable categories such as hate speech, pornography, etc., are quite rare in the natural distribution of web pages faced by most advertisers. However, wholesome nature of most brands necessitates reducing the occurrence of such content as much as possible much more still. The problem faced by data mining in such settings is that traditional methods for procuring labeled training data, including both labeling random samples and active learning, are much less effective under such extreme class imbalance. On the other hand, there is an alternative way to deploy human resources for training-data acquisition: have them “guide” the learning by searching explicitly for training examples of each class. We show that under extreme skew, even basic techniques for guided learning completely dominate smart (active) strategies for applying human resources to select cases for labeling. Thus, it is critical to consider the relative cost of search versus labeling, and we demonstrate the tradeoffs for different relative costs. Finally, we show that in cost/skew settings where the choice between search and active labeling is equivocal, a hybrid strategy can combine the benefits of each.

1. INTRODUCTION

This paper concerns the interaction of humans in the data acquisition phase of the process of building classification models from data. Consider the following example data mining application: classifying web pages for the purpose of *safe advertising*. Advertisers and advertising networks (hereafter, advertisers) would like a rating system that estimates whether a web page or web site displays certain objec-

tionable content. With such a system, advertisers can control the destination of their ads, advertising only on those pages deemed unlikely to display such unacceptable content (depending on the advertiser, objectionable categories include: adult content, kids content, hate speech, malware, etc.).¹ Evaluating each potential advertising opportunity involves classifying the web page with respect to these objectionable categories. The classification system can take into account various evidence, including the URL, the page text, anchor text, DMOZ categories, third-party classifications, position in the network of pages, and so on [14, 1]. For this paper, we will consider only the textual html source for each page, but the ideas generalize to any type of available feature data.

Manually examining every page encountered by such a system would be prohibitively expensive. This is particularly true in safe advertising, where models for new classification categories must be built rapidly to meet the changing demands of each customer and campaign. Furthermore, assuming that these classifications are based on statistical models, predictions will be more or less effective depending on the particular cases used for training and on the amount and distribution of training data used in their construction. For a given budget, some subset of the cases can be examined by humans—potentially at very low cost using a micro-outsourcing system [17] such as Amazon’s Mechanical Turk [2]—to produce training data.

An important question then is: *which cases should be selected for training?* Simply sampling cases uniformly at random is unlikely to be the best strategy, as is evidenced by the rich field of research comprising active learning[15]. Settings with extreme class imbalance, as is frequently the case on the web, further reduce the effectiveness of random sampling since for reasonable labeling budgets, only rarely would such sampling produce a positive training example at all. For instance, we would hope that the distribution of pages that are offered an advertiser for ad placement contain only a tiny fraction of pages containing hate speech. For safe advertising, depending on the category, the base rate of the minority class can be $\frac{1}{10^4}$ to $\frac{1}{10^7}$. Occasionally filters can be provided on the data (for example, based on selected phrases), to reduce the skew by orders of magnitude, however, generally we see base rates of less than $\frac{1}{10^2}$. These filters introduce a new problem: they bias the data—training and test—toward particular “disjuncts” [21] of the

Revised Feb. 5, 2010.

The 2010 Winter Conference on Business Intelligence, March 11-13, 2010, Salt Lake City, Utah.

Working Paper #CeDER-09-07, Stern School of Business, NYU. 12-2009.

This work was conducted while the authors were at AdSafe Media. Foster Provost thanks NEC for a Faculty Fellowship.

¹This site rating system may be best developed and maintained by a third party to avoid conflicts of interest, but that complexity does not affect the development here.

objectionable category. In the very extreme cases, active learning simply finds no minority-class examples. For generality this paper considers lower skews; the trends to the far extreme are clear.

While many active learning techniques have seen a great deal of success in many settings, extreme class imbalance introduces the same complication for active learning: examples of the positive class (e.g., adult content, hate speech) appear too infrequently in the pool of cases considered for labeling. Any strategy for selecting examples automatically that does not solve the classification problem itself, will be likely to select mostly negative examples. As demonstrated below in Section 5, as the minority class becomes more and more scarce, standard active learning strategies have increased difficulty finding instances that improve performance on held-out data.

The contribution of this paper is not the introduction of new, sophisticated data mining techniques. The methods we use for the most part either are existing techniques, or are fairly straightforward. This paper presents an empirical analysis of different strategies for applying human resources to inducing classifiers in domains where the base rate of an important category is very low. The main novelty is in the consideration of alternative ways to apply human resources to the data mining process. The main contributions of the paper are threefold (as follows).²

First, the paper provides an empirical analysis of the performance of traditional active learning techniques in highly skewed settings. This investigation reveals the deficiency of many active learning techniques in the extreme imbalance setting; dependency of the active learner on an uninformative model leads to repeatedly making the same mistakes, often leading to selecting mainly majority-class instances for labeling, while important portions of the minority class remain completely unrepresented in the training data.

Second, the paper contrasts active learning with “guided learning.” Active learning is based on the availability of human resources that can be applied to the *labeling* of specially selected examples. Guided learning applies those resources specifically to *search* for examples satisfying some criteria; for example, humans could be directed to search specifically for positive examples. For example, for safe advertising adult content or hate speech may have a fairly low prevalence among the pages supported by a particular advertiser (or on the web more generally); however, a human with a search engine may be able to find examples fairly quickly. We examine the relative benefit of using humans for labeling cases (in combination with an active learning strategy) versus using humans for searching for cases.

The results are striking. Looking ahead to Figure 3, it is clear that that straightforward guided learning strongly dominates active learning—improving accuracy much faster as examples are added to the training set. The details of all techniques will be covered below. This is important not only as practical guidance. This result shows that the dominant problem in these domains is simply finding minority-class examples, not finding otherwise “informative” examples or examples near the classification boundary. The paper proceeds with a deeper empirical analysis of this phenomenon. For example, does this result still hold if search for examples

²As a minor contribution we introduce the application of classification for safe advertising, which exhibits extremely skewed class distributions.

is more expensive than labeling examples?

Third, the paper presents and evaluates *hybrid* strategies for cost-effective guided learning, that utilize both search and active labeling. The results show that these hybrid strategies can perform better than either pure guided learning or pure active learning, when the setting does not provide clear dominance of one over the other. An ultimate goal for this sort of learning would be to judge the relative benefit-per-unit-cost of each sub-strategy, and allocate resources to labeling or to search accordingly.

The techniques described here are operating in production as part of the technology underlying the rating system of AdSafe Media.³ Human analysts are tasked with labeling and with search for guided learning; they are supported by systems for web-page labeling and for web search. Models are built across various categories of objectionable content, including adult content, hate speech, violence, and others. Guided/active learning procedures feed analysts with search tasks and with examples to be labeled. The resultant models are used to reduce objectionable on-line advertising adjacencies. In practice, we find that mixing guided learning and active learning is preferable to either in isolation.

The remainder of this paper proceeds as follows: Section 2 covers the baseline techniques used for comparison and explains the details and motivation behind guided learning. Section 3 covers prior work on classification and active learning in imbalanced settings. Section 4 presents the experimental framework and datasets used for evaluation. Section 5 covers the results of these experiments. Section 6 covers the behavior of guided learning under different cost settings, and presents and evaluates hybrid guided/active data acquisition strategies. Section 7 provides further discussion of the issues raised by this work, offers concluding remarks, and notes.

2. LEARNING WITH HUMAN INTERVENTION

As discussed, this paper analyzes two different methods for incorporating human resources in the data mining process. Specifically, via labeling carefully chosen examples, or via searching for examples. We assume that the reader is familiar with the notion of active learning for choosing examples for labeling; Settles provides a comprehensive survey [15]. We call the search for examples based on particular criteria “Guided Learning.” Here we describe the particular techniques that we study in this paper.

2.1 Active Learning

For active learning, this study employs two strategies. The first, uncertainty sampling, is by far the most popular active learning strategy, and is closely related to model-specific strategies such as actively selecting instances closest to a separating hyperplane [19].⁴ The second is a variation of the popular Query-by-committee [16] technique, specifically introduced to deal with skewed class distributions [18].

- **Uncertainty Sampling:** instances with the smallest

³<http://www.adsafemedia.com>

⁴For example, for an SVM that produces probability estimates via the common technique of applying a simple logistic regression to the orthogonal distance of an example from the separating hyperplane, uncertainty sampling will select the unlabeled examples closest to the separator.

margin are chosen for inclusion at each fold. Here we calculate margin as $|p(0) - p(1)|$ [11].

- **Boosted Disagreement with QBC:** instances are ordered by a class-weighted disagreement measure, $-\sum_{j \in \{0,1\}} b_j \frac{V(k_j)}{|C|} \log \frac{V(k_j)}{|C|}$, where $V(k_j)$ is the number of votes from a committee of size $|C|$ that an instance belongs to a class k_j . b_j is a weight corresponding to the importance of including the a certain class; a larger value of b_j corresponds to a increased tendency to include examples that are thought to belong to this class. From a window, W of examples with highest disagreement, instances are selected greedily based on the model’s estimated class membership probabilities so that the batch selected from the window has the highest probability of having a balanced class membership.

2.2 Guided Learning

Guided Learning is an alternative technique for utilizing human resources for model development, beyond traditional (active) instance labeling. Here, humans are tasked with *seeking* examples satisfying some criteria. For this paper, the basic guided learning task is straightforward: find examples representing the different classes in some proportion, ρ . These instances are provided as input to classifier induction.

Humans, using tools such as web search engines combined with their own background knowledge on the criteria defining the task, can often find examples with an efficiency far exceeding those selected by a model-based active learner. This is particularly true in the early stages of active learning where the model does not have a refined knowledge of the input space.

Guided learning is motivated by the results of Weiss & Provost [20, 22], who address the question “if only n training examples are to be selected, in what proportion should the classes be represented?” Their results show that the best proportion varies across domains; however, if one wants to maximize the ranking of cases (i.e., the AUC) a proportion of $\rho = 0.5$ is a very good choice. In principle the problem of this paper is different: how to use human resources to *search* for valid examples using all tools available to them—including both active learning and guided learning. Nevertheless, this paper’s analysis could be seen as a follow-on to this prior work; in our experimental setting we simulate guided learning by class-conditional random sampling. We describe the simulation below.

More specifically, a thorough evaluation of a guided learning system in the wild would require a sizable labeled pool of instances, in effect defeating the cost savings of the techniques proposed here. In order to compare and contrast different techniques, all guided learning experiments presented here are performed in the following way: given an initial pool of labeled instances P with some subset of minority and majority instances, P_+ and P_- respectively, along with a selection ratio, ρ , at each batch, the guided learning simulator selects $\rho|b|$ instances from the P_+ uniformly at random and $(1 - \rho)|b|$ instances uniformly at random from P_- , where $|b|$ is the size of the batch selected at each selection epoch. This process proceeds until either pool is exhausted, at which point, the process switches over to purely random sampling from the other class. This simulation is similar to the procedure of Weiss & Provost who assume that examples

can be produced randomly by class.

3. RELATED WORK

As mentioned above, guided learning was motivated by the results of Weiss and Provost [22, 20]; they investigate the influence of class distribution on classifier performance, empirically showing that given a training set of n examples, barring domain-specific information a balanced class distribution tends to offer the best AUC on held out data.⁵ Lomasky et al [13] also investigate the problem of selecting instances for model creation when their class is known a priori. In both of these lines of work, examples are selected randomly from class-conditional distributions, which is directly related to our experimental setting, but not necessarily to actual guided learning (see the discussion of limitations in Section 7. This paper could be seen as following this line of work, providing an analysis of the effectiveness of these strategies for highly skewed data, comparing with active learning strategies, and combining different strategies for applying human resources to data acquisition.

A closely related task was addressed by Weiss & Provost [22, 20] and Lomasky et al. [13]: given that we can obtain examples by class, how do we decide which classes to choose from, either in general, or in an active selection process. This is a complementary task to what we address in this paper. Our analysis assumes simply that examples will be provided in a particular proportion (balanced by class for our experiments). Incorporating techniques for better choosing the class distribution in the training data could improve the guided learning results presented below and is a direction for future work.

There is an extensive body of work investigating strategies for learning in highly skewed settings. This work includes over-sampling the minority class or under-sampling the majority class [5, 12]. A different branch of work investigates the application of non-uniform misclassification costs during training in order to give additional consideration to the class of interest [6].

There has been some work on active learning on skewed data. Tomanek [18] investigates Query By Committee-based approaches to sampling labeled sentences for the task of named entity recognition. The goal of their selection strategy is to encourage class-balanced selections by incorporating class-specific costs. This work assumes that classifiers can often accurately infer which instances belong to the minority class, giving higher weight to instances thought to belong to the minority class and with a high degree of uncertainty. Our work differs from this by extending to extreme cases where initial performance is poor. Additionally, our techniques are more general, able to extend beyond the tasks faced in NLP.

Bloodgood and Shanker [4] use a similar approach to [18], incorporating class specific cost factors to encourage choosing from the minority class in the skewed setting. Here the base rate is estimated on a small random sample. We note that in many realistic settings, random samples may not reveal any minority instances, thereby foiling this technique.

Zhu and Hovy [23] investigate active learning in conjunction with over and under-sampling to alleviate the class imbalance problem. Here active learning is used to choose a

⁵Many practitioners used this as a rule of thumb prior to Weiss & Provost’s research.

set of instances for labeling, with sampling strategies used to improve the class distribution. Our work differs by seeking strategies for acquiring a good class distribution in the data, removing the necessity for performing sub-sampling.

The work of Ertekin et al [9] is the learning on highly imbalanced data sets. Given a large, imbalanced pool of labeled instances, the authors randomly sub-sample instances, choosing to keep only those that are closely positioned to the margin of a SVM classifier. The authors do not address the problem of seeking unlabeled instances in the wild. Furthermore, the margin-based active learning heuristic is very similar to uncertainty sampling, a strategy that we demonstrate to exhibit difficulty in the extremely skewed cases.

We note that many active learning strategies depend to some degree on the quality of the current model—until the model “warms up”, the instance selection is essentially random. This cold-start problem has been examined by Zhu et al [24], work extended by Donmez and Carbonell [7]. This work seeks to find “clusters” of distinct content among the unlabeled instances. While this offers greater potential overcoming the cold-start than many common active learning techniques, it is still unlikely to succeed in the extremely skewed case; there is often so much diversity within the majority that the method will miss any minority instances. Additionally, these complex methods don’t scale well to the data sizes necessary to experience an extreme class skew.

Donmez et al [8] propose a hybrid active learn technique whereby a density-sensitive learning technique is used to overcome the initial deficiencies of uncertainty sampling until the derivative of the learning rate decreases below some threshold. After this point, traditional uncertainty sampling is incorporated to the instance selection. The intuition here is that the density-sensitive technique is better for exploring the space, while uncertainty sampling is better at “fine tuning” the decision boundary.

4. EXPERIMENTAL SETUP

The experiments for this paper are performed on six data sets with similar characteristics; all represent a task of separating examples of one minority class from examples of a diffuse collection of other topics. While all use text as the raw feature data, the techniques illustrated here apply to any other type input. The first two are taken directly from the domain of safe advertising; the others are publicly available surrogate data sets with similar problem structure. Specifically, the data sets are:

1. **Safe-Adult** A proprietary set of 35,000 pages labeled based on the presence of adult content. Positive instances here are deemed unsafe for advertising.
2. **Safe-Guns** A set containing 55,000 pages labeled based on the presence of guns, ammunition, bombs, or other destructive equipment. Often advertisers choose not to be associated with this type of content.

The two previously mentioned datasets represent random sub-samples from much larger datasets for experimental convenience. Safe-Adult has a class skew of approximately 20 : 1 while Safe-Guns has a class skew of roughly 150 : 1.

The next three data sets were taken from urls contained in the topical hierarchical taxonomy of the Open Directory Project [1]. This data set is a result of a crawl of approximately 4,000,000 urls, and instances are assigned class la-

bels based on their membership in top-level DMOZ categories. To eliminate confusion, pages belonging more than one category were eliminated from this experiment. Data sets were further down-sampled in order to induce a greater degree of skew.

3. **DMOZ-Science**: Positive instances belong to the top level category of Science, while the minority instances belong to all other categories. This set has approximately 130,000 instances. While this data set was used for experiments of varying skew, the nominal class distribution for this set is 200 : 1.
4. **DMOZ-News**: Here positive instances are pages found in the News top level DMOZ category. This data set has 100,000 instances with a class skew of 100 : 1.
5. **DMOZ-Games**: Urls samples from the Games category make up the positive category in this data set, sub-sampled to give 100,000 instances with a 100 : 1 class skew.
6. **20-News-Groups**: This data set is derived from the popular 20 News Groups set frequently used in text mining evaluation [3]. The data set is modified from the original data by assigning a positive label to all science related articles, and a negative label otherwise. Positive instances are randomly removed from the data set to give a highly skewed label distribution of roughly 80 to 1.

Classification and probability estimation are performed with logistic regression trained using stochastic gradient descent. The choice of logistic regression was based on this algorithm’s efficiency during training and induction, critical given the massive numbers of experimental runs performed in this work. Preliminary experiments lead indicate that the main effects described in this work are independent of the type of model used; other machine learning algorithms could be substituted with similar results.

All experiments compare the area under the receiver operating characteristic curve (AUC) at various stages in learning [10]. This metric allows a comparison of model performance that is largely insensitive to the class prior in the evaluation set. This is critical in a highly skewed setting where simply choosing the majority label for each instance would yield very high (and misleading) classification accuracy. The results presented are averages computed after a ten fold of cross validation for every experiment. As the baseline against which to comparing human-intervention strategies we use uniform random sampling.

5. ACTIVE VS. GUIDED

Figure 5 compares the four acquisition strategies covered in sections 2.2 and 4 on the six high-skew data sets under consideration. These plots show how the the area under the ROC curve (AUC; vertical axis) improves with additional labeled training data, as acquired by the different strategies. From these plots we can see that even with moderate imbalance in the class distribution, random sampling almost always underperforms the basic active strategies.

The results show that searching for examples of each class in even proportion, *even without any “active” selection at all*, provides substantially more informative data to the modeling process. For 20-News-Groups, the guided learning

achieves very good class separation (AUC in the high .90s) with considerably fewer examples than are required by the active strategies and random sampling. For Safe-Adult we see similarly that guided learning very quickly reaches maximal accuracy.⁶

On 20-News-Groups we see an interesting behavior: uncertainty sampling and boosted disagreement perform quite well initially, followed by plateaus, as they apparently cannot find the instances that will improve the model beyond a certain point; only after they exhaust a large number of seemingly uninformative examples do they choose examples for labeling that again provide improvement over random sampling. This behavior seems to suggest that the minority class is disjunctive, with portions of the class lying within the high-certainty (of majority) regions of the example space. Examples from these disjuncts are only selected when the less-certain instances are exhausted. As a result, little improvement is offered after repeated example selection. In effect, the ability of the underlying model to find examples of the as-of-yet unexplored disjuncts hinders the success of the active learning techniques.

Additional interesting behavior is observed in the learning curve of guided learning in Figure 1(e). Here guided learning quickly achieves a high AUC, followed by a gradual decline. We should note that due to the constraints of this simulation, there are only a limited number of minority instances for the guided learner to select (see below), after which the algorithm simply performs random selection from the majority class. It is possible that there is sufficient label noise in this case that the good data set provided in the initial epochs is simply washed out as more data is added. This provides a strong motivation for the hybrid techniques covered in Section ??.

Boosted disagreement—though designed specifically designed for active learning in skewed settings by favoring uncertain instances belonging to the minority class—seems to suffer from the same problems as uncertainty sampling, with few exceptions. The models forming the committee are simply unaware of much of the minority class; the method has limited ability to find instances that would improve the generalization performance of the system.

These results suggest either that the active learning heuristics are ill-suited for learning these concepts, perhaps just missing certain subclasses, or more probably that the active learning techniques are having difficulty finding positive examples at all, due to the class imbalance. *Finding* examples of a rare class is critical. The guided learning used here has no provisions whatsoever to find difficult-to-classify, or otherwise uncertain, examples.

Figure 2 examines graphically the relative positions of the minority examples through the active learning. The black curve shows the AUC (right vertical axis) of the models learned by uncertainty sampling on the 20-News Groups data set as in Figure 5 (rescaled as follows). At each epoch we sort all instances by their predicted probability of membership in the majority class, $\hat{P}(y = 0|x)$. The blue dots in Figure 2 represent the minority class instances, with the value on the left vertical axis showing their relative position in this sorted list. The x-axis shows the active learning epoch (here each epoch requests 30 new instances from the pool).

⁶With a larger absolute number of instances generalization performance is better, but the qualitative comparison remains the same.





Figure 2: A comparison of the underlying model’s ordering of the pool with the quality of the cross validated AUC

The blue trajectories mostly show examples’ relative positions changing. Minority examples drop down to the very bottom (certain minority) either because they get chosen for labeling, or because labeling some other example caused the model to “realize” that they are minority examples.


We can see that all the way up to the end of the first major plateau, there are (groups of) minority examples all through the sorted probability list, whose relative position does not change much with each epoch (the dribbling down is due to many known majority examples moving to the top of the list). It’s not until the end of the plateau that we see an effect on these minority examples. Even on the final plateau, when the AUC suggests a very good model, there are minority examples (or groups of examples) whose estimated class-membership probability is not affected at all by the active learning: the active learning procedure does not explore their regions of the example space.

The motivation for introducing guided learning was to facilitate cost-effective learning in settings with high to extreme class skew. We now assess the relative benefit of guided learning for different skews. Figure 3 presents the learning curves for three different instance sampling strategies at different base rates, as induced on our DMOZ-Science data set. Specifically, this plot displays the AUC for three different labeling strategies, uniform random sampling, uncertainty sampling, and guided learning with $\rho = 0.5$. We see that as the skew increases, uncertainty sampling and random sampling have increased difficulty selecting minority instances, resulting in poor AUC. In the most skewed setting, random sampling and uncertainty sampling are unable reliably select any minority instances given 10,000 draws from the available pool. This likely results in a complete lack of understanding of the space, effectively giving worse-than-random performance.

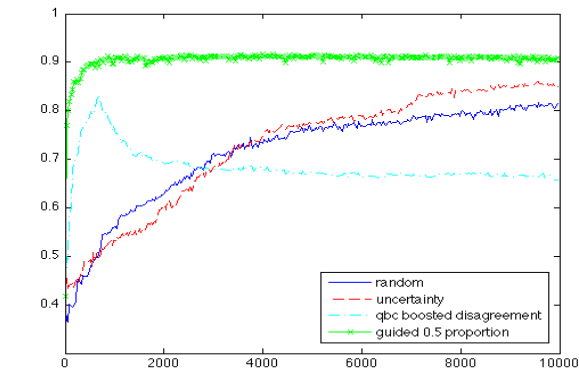
For this experiment, skew was induced by taking a large data set and randomly removing instances belonging to the minority class. This limitation on the number of available minority instances adversely affects the maximum performance that model achieve as the skew increases. We note that this experimental limitation has the effect of handicapping guided learning; on the web, even cases that occur extremely infrequently in relation to the web as a whole still occur in great numbers in the absolute sense. While a guided learning scheme may require greater cost and effort




(a) Safe-Adult




(b) Safe-Guns




(c) DMOZ-Science



(d) DMOZ-Games



(e) DMOZ-News



(f) 20-News-Groups

Figure 1: Comparison of active learning strategies and guided learning. The vertical axis shows the generalization performance of the learned models, measured by the area under the ROC curve (AUC). The horizontal axis shows the number of examples labeled/acquired. (Weighted) uncertainty sampling outperforms random sampling. Guided learning dominates by a large margin.

to get new minority instances as the number of epochs increases, there is unlikely to be a hard cap on performance as seen in this experiment. It is clear that given enough information, a model can reach AUC measurements in excess of 0.9; the performance of guided learning is therefore underestimated. The alternative experimental setting of adding more majority instances to induce such high skews is not practical.

6. COST-SENSITIVE GUIDED LEARNING AND HYBRID ACQUISITION STRATEGIES

The per-instance cost for a guided learning strategy is likely to differ from strictly label-based active learning. Searching for an example of an obscure class may require more effort than simply identifying if a given sample belongs to the class of interest. Alternatively, using tools like web search engines, clear-cut examples may be readily found, whereas




Figure 3: Comparison of Random Sampling, Uncertainty Sampling, and Guided Learning on the same data set with induced skews ranging from 5 : 1 to 10,000 : 1.

labeling would require time-consuming analysis of each case. The relative costs of guided learning and instance labeling vary from setting to setting, and in this section we seek to investigate the behavior in a variety of cost scenarios.

Figures 4 compare various instantiations of our guided learning approach with uncertainty sampling on the three data sets, where the curves show the increase in generalization performance as a function of investment in human effort (labeling or search). The horizontal axis shows the total cost expended by each strategy; to normalize, we fix the cost of labeling to be 1. For active learning, we only report for uncertainty sampling. The different instantiations of guided learning vary the cost of search from 0.5 to 16, doubling each step.




Figure 4: Comparison of guided learning and active learning under different relative costs for search and labeling (20-News-Groups data set). Vertical axis shows AUC; horizontal axis shows total cost, normalized to 1 for acquiring one label. For this data set, guided learning and active learning have equivalent performance-per-unit-cost when search is about 8 times more expensive than labeling.

By construction, the performance-per-unit-cost of guided learning declines gradually as the cost is increased. These results show how one can judge the relative value of applying human resources for search and for labeling. For example, for 20-News-Groups, in terms of performance-per-unit cost uncertainty sampling seems to be approximately equivalent to guided learning when search is approximately 8 times

more expensive than labeling.

Note that practical experience with micro-outsourcing systems have revealed that several hundred minority class web pages can be obtained for model induction at about double the cost-per-page required to label an example pages. Of course, the time required to perform this task tends to be longer than is required to perform page labeling, and the costs and time required to perform this work may vary as the requested number of examples increases.

6.1 Hybrid Strategies

As discussed, guided learning seems to be most appropriate in cases where the class priors are extremely unbalanced, and the cost structure is skewed in the opposite direction: discovering missing instances from the minority class is more expensive than missing examples from the majority class. The results presented up to this point have assumed that one would have to choose either guided learning or active learning. In practice, it might be better to mix strategies. For example, one might bootstrap the active learning process by first searching for good training data, potentially at a higher cost-per-example; alternatively, if one suspects that the active strategies have reached a plateau (as in Figure 1(f)), search may be used to inject additional information.

Given a budget, B , a data set, D , and a cost structure, C , policies for guided/active learning will allocate B to some combination of guided search and instance labeling. The goal of this section is to illustrate hybrid guided/active strategies can be designed to offer superior performance for a given B than either strategy would be capable of in isolation.

While such a hybrid strategy could take many functional forms, here we propose a switching strategy based on the *DUAL* technique used by Donmez et al [8]. We rely on our background knowledge that guided learning excels at finding different examples of the minority class, while many traditional active learning techniques are better at fine-tuning the decision boundary of the base model.

First, the benefit to model generalization of a purely guided selection strategy as a function of the cost of human effort is estimated. When the returns for further guided selection are sufficiently low, we switch to a purely active strategy. We note that almost any conventional active learning strategy may be employed for this phase of our hybrid approach. Due to its simplicity, documented success, and ubiquity, we chose uncertainty sampling as the active learning heuristic

used in this second phase.

More succinctly, given a certain cost structure representing the cost-per-query to an oracle performing guided learning, we perform guided learning by selecting instances from both classes in proportion ρ . After each phase of guided learning, we estimate the performance, A , using this performance estimate to construct a learning curve. When the expected gain for performing additional guided learning as a function of cost is sufficiently low, $\frac{\partial A}{\partial c} \leq t$, we switch from guided learning to a more traditional active learning strategy that requires only choosing examples from the pool for which to request labels.

In order to determine when to switch phases in our proposed switching strategy, we must understand how the performance of a model is changing under a given selection scheme, as a function of that scheme’s cost, $\frac{\partial A}{\partial c}$. This requires careful estimation of the model’s performance at each epoch. To accomplish this, we compute x-validated accuracy of the current model on the available pool of instances. Progress of learning curve is estimated using Local Regression in order to smooth the variances in estimated learning rates at each epoch. More succinctly, the learning rate at any point is estimated by computing the slope of a least-squares linear regressor fit to performance estimates (in this case x-validated accuracy). When the slope of accuracy as a function of cost drops below some threshold, t , we change strategies from guided learning to active learning.

As discussed above, certain cost/skew structures are likely to be best managed by simply performing a guided search. In the case of the 20-News-Groups data set, the cost setting offering a break even point between guided learning and active learning seems to be about 8 : 1, one guided search costs as much as eight label requests.

The learning curves traced by our proposed hybrid technique on the 20-News-Group dataset are presented in Figure 5, using the approximate break even cost of 8 labels per search result. We see that under this cost setting, a switch from guided learning to active learning does indeed improve the learning rate beyond that which is possible from either component technique in isolation. Note that as the cost approaches approximately 2,000, the slope of the learning curve increases drastically as the selection strategy switches from guided learning to active learning. The ability of traditional active learning to refine the class boundary excels given the strong base classifier enabled through guided learning. We note that our hybrid technique is able to reach AUC levels in the high 0.9’s at a far smaller cost than either guided learning or active learning in isolation. While space constraints prevented identical plots from being presented on the other data sets considered, the results are qualitatively similar.

7. CONCLUSION & LIMITATIONS

The main result of this paper—that guided learning can dominate active learning so strongly—raises the possibly contentious question of: when should we be doing active learning at all? Through analysis of the hybrid techniques presented in the previous section, it is clear that such situations exist (for instance, under low class skews and high-cost settings), however, the results presented in this paper (Figure 4) make it clear that the question warrants further investigation.

This is not a trivial point. Research on active learning




Figure 5: Comparison of our proposed switching strategy with pure guided learning and uncertainty sampling

almost always makes one (or both) of two assumptions: (i) that labeling via initial random sampling is going to produce a model sufficiently accurate to do active learning, or (ii) that there is some “cold start” process that provides the system with a small set of labeled examples to prime the process with a model of sufficient accuracy for active learning to be effective. With even moderately high skews, assumption (i) very often does not hold. With even 999 : 1 skew, a labeler would have to label 30,000 examples just to get 30 minority examples.

So most research starts with assumption (ii) that some unspecified process has produced a small labeled training set (often balanced). The results of this paper raise the question: why not just continue with that process? Why do active learning at all? It is possible that the process was an unrepeatably (historic) stroke of good fortune.⁷ Otherwise, these results suggest that we may want to put more research emphasis on this process that procured the initial training set. What is the relative cost as compared to active learning? Under what conditions does it make sense to continue with it, versus switching to active learning, versus applying some combination of both. Moreover, perhaps we should be investing more effort in reducing the cost of search for examples.⁸

A major limitation of the analysis presented in this paper is that the guided learning was simulated as class-conditional random sampling. Real, live guided learning setting is likely to be conducted via a searching system—for our running example, a search engine or a list of resources such as DMOZ. As a result, the instances returned from a guided learn-

⁷For example, in our safe advertising example, an advertiser complained about a particular set of web pages on which her ad appeared, which then become labeled training data.

⁸An initial reaction to the question of quantifying the relative cost of search versus labeling often is: search *has* to be more expensive than labeling, since with search one must find the examples as well as label them. However, with a search engine and a human brain, it may be less costly to envision what a case might look like and find it, than to examine a presented case in detail to be sure. For example, one may be able to search and find examples of all manner of hate speech on the web more efficiently than reading carefully through a borderline web page to determine whether or not it contains some form of hate speech.

ing system are likely to differ significantly from a class-conditioned uniform random selection. This could bias the results above in either way: they could be overly pessimistic if initial instances returned would in some sense be the best, most informative representatives of that particular class. They could be overly optimistic if the internal biases of the search engine or the human user made the selection of certain examples unlikely. Research has shown repeatedly that classifiers' errors are concentrated in the "small disjuncts"—the model's representations of the rare cases [21]. It is unclear just how that relates to the result presented above.

Relatedly, above we mentioned that in practice extreme imbalance often is reduced by orders of magnitude via filters on examples. Such filters can be thought of as codified search procedures for guided learning. These filters suffer from the same potential bias problems, and our experience is that the bias is substantially more extreme than with human searchers, because of the inflexibility. These filters bias the data largely toward particular disjuncts of the interesting category. In practice, this introduces a particularly insidious problem for data mining: the testing data are biased as well!

Thus, it is necessary to provide tools that simultaneously reduce the cost of search, and challenge humans to explore the "far reaches" of the category. This suggests a different integration with active learning: guiding the human search away from cases already "known" by the model to be members of the class, and (somehow) toward as-of-yet unexplored disjuncts.

Our simulation has other important differences from a real-world implementation of guided learning, particularly in the case of document search on the web. First, the extreme size of the web is likely to yield a near-unlimited number of instances belonging to almost any category. Thus, unlike with our experiments, the guided learning curves would not run out of minority class examples (which is often the apparent reason they "knee over" so sharply). On the other hand, many instances in the web's long tail may be increasingly difficult to find. This could mean that the constant-cost model is simplistic; search probably incurs increasing cost as the number of requested instances increases, and as the exploration of the class probes the smaller disjuncts.

This paper does not explore the possibility that guided learning might focus solely on finding minority classes. In cases of extreme skew, one likely could just presume a randomly selected case to be a majority class, and deal with the small amount of resultant noise. This would effectively halve the cost of guided learning as we present it.

There still is much to do to understand the best ways to employ human resources for some combination of search and labeling, to produce the best models per unit investment in training data. We hope that this paper has made useful headway.

8. REFERENCES

- [1] <http://www.dmoz.org/>.
- [2] <https://www.mturk.com/mturk/welcome>.
- [3] A. Asuncion and D. Newman. UCI machine learning repository, 2007.
- [4] M. Bloodgood and K. V. Shanker. Taking into account the differences between actively and passively acquired data: the case of active learning with support vector machines for imbalanced datasets. In *NAACL '09: Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*, pages 137–140, Morristown, NJ, USA, 2009. Association for Computational Linguistics.
- [5] N. V. Chawla, K. W. Bowyer, and P. W. Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.
- [6] P. Domingos. Metacost: A general method for making classifiers cost-sensitive. In *In Proceedings of the Fifth International Conference on Knowledge Discovery and Data Mining*, pages 155–164, 1999.
- [7] P. Donmez and J. Carbonell. Paired Sampling in Density-Sensitive Active Learning. In *Proc. 10th International Symposium on Artificial Intelligence and Mathematics*, 2008.
- [8] P. Donmez, J. G. Carbonell, and P. N. Bennett. Dual strategy active learning. In *ECML '07: Proceedings of the 18th European conference on Machine Learning*, pages 116–127, Berlin, Heidelberg, 2007. Springer-Verlag.
- [9] S. Ertekin, J. Huang, L. Bottou, and L. Giles. Learning on the border: active learning in imbalanced data classification. In *CIKM '07: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 127–136, New York, NY, USA, 2007. ACM.
- [10] T. Fawcett. Roc graphs: Notes and practical considerations for data mining researchers, 2003.
- [11] D. D. Lewis and W. A. Gale. A sequential algorithm for training text classifiers. In *SIGIR '94: Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 3–12, New York, NY, USA, 1994. Springer-Verlag New York, Inc.
- [12] X. Y. Liu, J. Wu, and Z. H. Zhou. Exploratory undersampling for class-imbalance learning. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 39(2):539–550, 2009.
- [13] R. Lomasky, C. E. Brodley, M. Aernecke, D. Walt, and M. Friedl. Active class selection. In *ECML '07: Proceedings of the 18th European conference on Machine Learning*, pages 640–647, Berlin, Heidelberg, 2007. Springer-Verlag.
- [14] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Beyond blacklists: learning to detect malicious web sites from suspicious urls. In *KDD '09: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1245–1254, New York, NY, USA, 2009. ACM.
- [15] B. Settles. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.
- [16] H. S. Seung, M. Opper, and H. Sompolinsky. Query by committee, 1992.
- [17] V. S. Sheng, F. Provost, and P. G. Ipeirotis. Get another label? improving data quality and data mining using multiple, noisy labelers. In *KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 614–622, New York, NY, USA, 2008. ACM.
- [18] K. Tomanek and U. Hahn. Reducing class imbalance during active learning for named entity annotation. In *K-CAP '09: Proceedings of the fifth international conference on Knowledge capture*, pages 105–112, New York, NY, USA, 2009. ACM.
- [19] S. Tong and D. Koller. Support vector machine active learning with applications to text classification. *J. Mach. Learn. Res.*, 2:45–66, 2002.
- [20] G. Weiss and F. Provost. The effect of class distribution on classifier learning, 2001.
- [21] G. M. Weiss. The impact of small disjuncts on classifier learning. In R. Stahlbock, S. F. Crone, and S. Lessmann, editors, *Data Mining*, volume 8, chapter 9, pages 193–226. Springer US, Boston, MA, 2010.
- [22] G. M. Weiss and F. Provost. Learning when training data are costly: the effect of class distribution on tree induction. *J. Artif. Int. Res.*, 19(1):315–354, 2003.
- [23] J. Zhu and E. Hovy. Active learning for word sense disambiguation with methods for addressing the class imbalance problem. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 783–790.
- [24] J. Zhu, H. Wang, T. Yao, and B. K. Tsou. Active learning with sampling by uncertainty and density for word sense disambiguation and text classification. In *COLING '08: Proceedings of the 22nd International Conference on Computational Linguistics*, pages 1137–1144, Morristown, NJ, USA, 2008. Association for Computational Linguistics.