
Estimating Class Membership Probabilities using Classifier Learners

John Langford

Toyota Technological Institute at Chicago
Chicago, IL 60637
jl@tti-c.org

Bianca Zadrozny

IBM T.J. Watson Research Center
Yorktown Heights, NY 10598
zadrozny@us.ibm.com

Abstract

We present an algorithm, “Probing”, which reduces learning an estimator of class probability membership to learning binary classifiers. The reduction comes with a theoretical guarantee: a small error rate for binary classification implies accurate estimation of class membership probabilities. We tested our reduction on several datasets with several classifier learning algorithms. The results show strong performance as compared to other common methods for obtaining class membership probability estimates from classifiers.

1 Introduction

Background. Classifier learning is the problem of inducing a predictor for distinguishing between two (or a few) classes given a set of labeled training examples. There are many reasons for considering classifier learning as the most fundamental learning task. Predicting one bit (or a few bits) is as simple as non-trivial prediction can be. Because of this simplicity, classifier learning is often more amenable to study than other learning problems. On the practical side, several fast learning algorithms (such as SVM [13], boosting [6] on decision trees [14], and neural networks [9]) exist which generally lead to classifiers with excellent predictive performance.

The Problem. On the other hand, there are several real-world applications requiring probabilistic answers instead of discrete class labels. For example, in medical domains, doctors often prefer to make decisions based upon probabilistic predictions of the patient state, rather than being given a discrete chosen action for each patient. Even when the actual probability estimates are not required, a ranking of examples from the most likely member to the least likely member of a class may be desirable. This is the case,

for example, in document categorization, where users might like to see a list of documents ranked by their relevance to a particular topic.

Probability estimates are also important when the classification outputs are not used in isolation but are combined with information from other components in a system. For example, in handwritten character recognition the outputs from the classifier are used as input to a high-level system which incorporates domain information, such as a language model. In this case, probabilities provide a standard language for summarizing information from multiple sources into a single decision. When the components of the system are distributed, bandwidth constraints may make this summarization necessary as well as standard.

Solution Sketch. Our method for answering the needs of these real-world applications can be thought of as a reduction of class probability estimation to classifier learning. The reduction (which we name “Probing”) satisfies strong optimality guarantees: good performance with respect to classification implies good performance with respect to class probability estimation. The essential observation underlying the Probing reduction is that probability estimates (in a Bayesian sense) can be extracted from preferences over bets at different odds ratios. The machinery that applies this observation to classifier learning algorithms is the Probing reduction.

2 Related Work

There are several existing approaches for obtaining class probability estimates from classifiers.

Bagging. One common approach uses bagging [12]. Essentially, the learning algorithm is run on many training sets formed by sampling with replacement from the original dataset, and the estimated probability that any sample x has label $y = 1$ is the proportion of learned classifiers that output 1 for x . The results of the bagging approach may have nothing to do with $\Pr_{(x,y) \sim D}(y = 1|x)$, the probability accord-

ing to the data generation process that $y = 1$. As has been observed before, the probability estimates obtained through bagging are measuring the instability of the classifier learning method over data perturbations [10] rather than $\Pr_{(x,y)\sim D}(y = 1|x)$. Indeed, for larger datasets, bagging with certain learning algorithms might result in the same classifier on every run, causing the bagging probability to be always 0 or 1 even when the fundamental process is much less certain. In contrast, our Probing reduction has a simple guarantee—if the classifiers solve their associated classification problems optimally, then the reported probability is precisely $\Pr_{(x,y)\sim D}(y = 1|x)$. In addition, unlike bagging, Probing preserves independence in the original dataset¹, which is quite important for many learning algorithms (see [17] for further discussion).

Margin Sigmoid. Another approach for extracting probabilities from classifiers is to take an internal representation such as the margin of a support vector machine and transform it into a value $v \in [0, 1]$ to achieve calibration on the training set (for example, using a sigmoid function [11]). The Probing reduction is more general (since it applies to classifiers without an internal margin such as decision trees) and less ad-hoc. In particular, in the Probing reduction a small classification error rate necessarily implies good probability estimates over all distributions generating the data. No guarantee of this quality can be made for any transformation of a margin into a $[0, 1]$ interval, essentially because the margin is not a “sufficient statistic” for probability estimation. See also [2] for a discussion of sparsity versus probability estimation issues.

Direct Prediction. The third standard approach is to simply predict probabilities directly, typically with some form of learned probabilistic model. Our results can be viewed from this approach as (1) providing compatibility between non-probabilistic and probabilistic models and (2) enriching the set of probabilistic models with new and often better performing models.

We discovered that a reduction similar to Probing had been published previously by Halck [7]. The main difference is that Probing reweights the classes appropriately to obtain preferences over bets at different odds ratios (see next section), while Halck’s reduction changes the labels of examples randomly for the same purpose. We provide a theoretical performance guarantee and present a comprehensive experimental evaluation of Probing. Halck’s method may be analyzable and may have good empirical performance, but the theory and evaluation have not been done yet.

¹Resampling with replacement provides samples from the correct distribution in a *dependent* manner.

3 The Probing Reduction

Assume we have examples (x, y) drawn from an unknown distribution D with domain $\mathcal{X} \times \mathcal{Y}$ where \mathcal{X} is the feature space and $\mathcal{Y} = \{0, 1\}$ is the label space.

Basic Observations. The Probing reduction is based upon the observation that a binary classifier which predicts 1 for an example x is implicitly predicting that the probability of class 1 given x is greater than the probability of class 0 given x , that is, $D(y = 1|x) > D(y = 0|x)$. Conversely, a classifier which predicts 0 for x suggests that $D(y = 0|x) > D(y = 1|x)$.

By weighting the training examples such that positive examples are w times more costly to classify incorrectly than negative examples, it is possible to learn classifiers which predict 1 if

$$D(y = 1|x) > \frac{1}{1 + w}$$

and 0 otherwise.

Corollary 1 (Minimal Error Probing) *Let $i_p(y) = \frac{1-p}{p}y + (1 - y)$. Then*

$$\begin{aligned} \arg \min_{c: \mathcal{X} \rightarrow \mathcal{Y}} E_{(x,y)\sim D} i_p(y) I(c(x) \neq y) \\ = I(D(y = 1|x) > p) \end{aligned}$$

except on a set of measure 0.

This corollary only discusses *minimal* error classifiers, so it only applies in the asymptotic limit of a bayes consistent classifier. However, the probing algorithm can be proved sound even for classifiers with nonminimal error rates. We prove this more general theorem in section 4 (theorem 2). The corollary states that a classifier which minimizes a importance weighted loss can be used to determine if the conditional distribution $D(y = 1|x)$ is above or below a threshold for a given value of x , where the threshold depends on the relative weight given to positive and negative examples.

The Algorithm. Now, suppose that we learn a classifier c_p for every choice of p and each of these classifiers has a minimal loss. Then, for an input x , each of the classifiers answers whether $D(y = 1|x)$ is greater than the threshold p . For $p = 0$, we necessarily have $c_p(x) = 1$ because $D(y = 1|x)$ cannot be less than 0. As p grows, the predictions of the subsequent classifiers are monotone in p with one point $p^* = D(y = 1|x)$ where the prediction changes from 1 to 0. Therefore, the Probing reduction reports:

$$D(y = 1|x) = p^*. \tag{1}$$

Three issues remain:

1. Not all classifier learners automatically take the weight w as input. We address this by using the

Algorithm 1 Probing-Learner (classifier learner A , dataset $S = (x, y)^m$, number of iterations t)

1. Let $I = \{[0, 1]\}$
 2. For $i = 1$ to t
 - (a) Let $[a, b] = \text{minimax interval} \in I$
 - (b) Let $p = \text{minimax optimal } p \in [a, b]$
 - (c) Let $I = (I - \{[a, b]\}) \cup \{[a, p], [p, b]\}$
 - (d) Let $S_p = \{(x, y, i_p(y)) : (x, y) \in S\}$
 - (e) Let $c_p = A(S_p)$
 3. Return all c_p
-

Algorithm 2 Probing-Predictor (set of weighted classifiers c_p , input x)

1. Let $n = \sum_p c_p(x)$
 2. Let $[a, b] = \text{the } n\text{th interval}$
 3. Return minimax optimal $p \in [a, b]$.
-

costing reduction [17] from importance weighted classification to binary classification, which applies to any classifier learner.

2. Learning a classifier for every choice of w is too computationally intense. We resolve this by discretizing the set of w and reporting a probability in the discrete interval.
3. Independently learned sub-optimal classifiers may not be consistent, meaning that the sequence of predictions for a given example is not necessarily monotonic. We simply sort the results of the classifiers to make the sequence monotonic (all zeros before all ones). This gives us the solution consistent with the largest possible number of classifiers. Although this may appear to be a hack, the proof of theorem 2 fundamentally relies upon the sort.

Learning. The algorithms for learning and making predictions using the Probing reduction are reported in figures Algorithm 1 and Algorithm 2. The Probing-Learner algorithm takes as input a classifier learning algorithm, a set of training examples and some number of iterations. On each iteration, it refines the discretization of the interval with the largest potential gain from refinement (this is loss dependent, see below for details). For each chosen interval a minimax choice of probability and the corresponding weight (obtained by inverting equation 1) are calculated and a weighted classifier is learned.

Prediction. The Probing-Predictor algorithm takes as input the set of weighted classifiers and an input

x . It sums the predictions of each classifier on x to determine the number of classifiers that predict 1 for x . It then picks the minimax probability in the n th interval (counting from 0). This is equivalent to sorting the classifier predictions to obtain monotonicity and reporting a probability from the interval in which the prediction changes from 1 to 0.

Loss Details. Steps 2(a) and 2(b) in Probing-Learner are dependent on the probabilistic loss function of interest. Here we describe these steps for two of the most commonly used probabilistic loss functions: cross entropy and squared error. In order to determine the minimax interval and the minimax probability for these loss functions, we make the somewhat unjustified but convenient assumption that all classifiers are correct.

1. Cross entropy is given by

$$E_{(x,y) \sim D} I(y=1) \ln \frac{1}{\hat{p}(x)} + I(y=0) \ln \frac{1}{1-\hat{p}(x)}.$$

The minimax optimal $\hat{p} \in [a, b]$ is given by:

$$\begin{aligned} & \min_{\hat{p} \in [a, b]} \max_{p \in [a, b]} [\text{loss on } \hat{p} - \text{loss on } p] \\ & = \min_{\hat{p} \in [a, b]} \max_{p \in [a, b]} \left[p \ln \frac{1}{\hat{p}} + (1-p) \ln \frac{1}{1-\hat{p}} \right. \\ & \quad \left. - p \ln \frac{1}{p} - (1-p) \ln \frac{1}{1-p} \right] \end{aligned}$$

This implies that

$$\hat{p} = \frac{1}{1 + e^{\frac{H(b) - H(a)}{b-a}}},$$

where $H(a)$ is Shannon entropy (in nats) of a coin with bias a .

The minimax interval is the interval which creates the largest potential change in loss when refined on the training set. Suppose we let S_a be the set of training examples whose current probabilistic prediction is $\hat{p} \in [a, b]$. The chosen ² interval is:

$$\arg \max_{[a, b] \in I} |S_a| \left[(1-a) \ln \frac{1-a}{1-\hat{p}} + a \ln \frac{a}{\hat{p}} \right]$$

2. Mean squared error is given by

$$E_{(x,y) \sim D} [(y - D(y|x))^2].$$

The minimax optimal $\hat{p} \in [a, b]$ is given by:

$$\begin{aligned} & \min_{\hat{p} \in [a, b]} \max_{p \in [a, b]} [\text{loss on } \hat{p} - \text{loss on } p] \\ & = \min_{\hat{p} \in [a, b]} \max_{p \in [a, b]} \left[p(1-\hat{p})^2 + (1-p)\hat{p}^2 \right. \\ & \quad \left. - p(1-p)^2 - (1-p)p^2 \right] \end{aligned}$$

This implies that

$$\hat{p} = (a + b)/2.$$

²Note that the definition \hat{p} implies this is equivalent to $\arg \max_{[a, b] \in I} |S_a| \left[(1-b) \ln \frac{1-b}{1-\hat{p}} + b \ln \frac{b}{\hat{p}} \right]$

The minimax interval is similarly given by

$$\arg \max_{[a,b] \in I} |S_a|(b-a).$$

In a transductive setting, using the test set distribution to weight the minimax interval is a better than $|S_a|$.

Multiclass Case. For the multiclass case with $k > 2$ labels, we simply apply the binary method to $k - 1$ binary problems defined by the mapping:

$$(x, y) \rightarrow (x, I(l = y))$$

for $k - 1$ choices of l . This method gives us probability estimates $\hat{p}_1, \dots, \hat{p}_{k-1}$ for $k - 1$ classes, so we can estimate the probability of the last class according to $\hat{p}_k = 1 - \sum_{i=1}^{k-1} \hat{p}_i$.

4 Analysis of the Probing Reduction

The following theorem shows that the probing reduction is inherently robust against classification errors. For the theorem, remember $i_p(y) = \frac{1-p}{p}y + (1-y)$ and let the importance weighted loss of a classifier c under distribution D be

$$l_p(c) = \frac{1}{Z_p} E_{(x,y) \sim D} i_p(y) I(c(x) \neq y)$$

where $Z_p = E_{(x,y) \sim D} i_p(y)$ is a normalization constant.

This theorem analyzes probing in the limit as the number of classifiers approaches ∞ uniformly over the unit interval. Discretization to intervals of size d adds at most $d + d^2/4$ to the expected squared error.

Theorem 2 (Probing Error Transformation) If the classifiers c_p learned under Probing have average relative importance weighted loss

$$\epsilon = E_{p \sim U(0,1)} \left[l_p(c_p) - \min_c l_p(c) \right],$$

then

$$E_{x \sim D} (D(1|x) - \hat{p}(x))^2 \leq 2 \max\{D(1), D(0)\} \epsilon \leq 2\epsilon$$

where $\hat{p}(x)$ is the probing prediction.

This is a strong theorem in the sense that it ties the error in the probability predictions to the *average relative* importance weighted loss of the classifiers. Using the average loss over w results in a more powerful statement than using the maximal loss, because it is easier to obtain a set of classifiers which have small average loss than to obtain a set of classifiers, all with a small loss. Using a loss that is relative to the loss of the Bayes optimal classifier means that the theorem applies even when the fundamental noise rate is large.

Note that when the proportion of positive and negative examples is balanced ($D(1)=D(0)=0.5$) the bound on the error in the probability predictions is exactly the average relative importance weighted loss,

$$E_{x \sim D} (D(1|x) - \hat{p}(x))^2 \leq \epsilon$$

Proof: Let $c_p^* = \min_c l_p(c)$, then

$$\begin{aligned} \epsilon &= E_{p \sim U(0,1)} [l_p(c_p) - l_p(c_p^*)] \\ &= E_p \left[\frac{1}{Z_p} E_{(x,y) \sim D} [i_p(y)(I(c_p(x) \neq y) - I(c_p^*(x) \neq y))] \right] \\ &= E_p \left[\frac{1}{Z_p} E_x [D(1|x) \frac{1-p}{p} (c_p^*(x) - c_p(x)) \right. \\ &\quad \left. + D(0|x)(c_p(x) - c_p^*(x))] \right] \\ &= E_p \left[\frac{1}{Z_p} E_x [D(1|x) \frac{1-p}{p} (c_p^*(x) - c_p(x)) \right. \\ &\quad \left. + (1 - D(1|x))(c_p(x) - c_p^*(x))] \right] \\ &= E_p \left[\frac{1}{Z_p} E_x [(c_p^*(x) - c_p(x)) (\frac{1-p}{p} D(1|x) - 1)] \right]. \end{aligned}$$

Rewriting with $D(1|x) > p \Rightarrow c_p^*(x) = 1$ and $D(1|x) < p \Rightarrow c_p^*(x) = 0$, we have

$$\epsilon = E_p \left[\frac{1}{pZ_p} E_x [I(c_p^*(x) \neq c_p(x)) |D(1|x) - p|] \right].$$

Notice that for all p , $Z_p = D(0) + D(1) \frac{1-p}{p}$. Hence

$$pZ_p = pD(0) + (1-p)D(1) \leq \max\{D(1), D(0)\}.$$

Therefore, we have

$$E_p E_x [I(c_p^* \neq c_p(x)) |D(1|x) - p|] \leq \max\{D(1), D(0)\} \epsilon,$$

which can be re-written as

$$\begin{aligned} E_{x \sim D} \int_0^1 I(c_p^* \neq c_p(x)) |D(1|x) - p| d_p \\ \leq \max\{D(1), D(0)\} \epsilon, \end{aligned} \quad (2)$$

Given a fixed budget of ϵ for binary errors, the scenario which maximizes probability error estimates has all classifiers erring in one direction (either for $p < D(1|x)$ or $p > D(1|x)$) because two errors in different directions are canceled in the sorting phase of Probing-Predictor. Furthermore, errors by classifiers nearer to $D(1|x)$ are preferred to errors far from $D(1|x)$, since the importance weighted loss paid by the adversary increases monotonically with distance from $D(1|x)$. Therefore, when the Probing prediction is $\hat{p}(x)$, all the classifiers c_p for which the corresponding p are between $\hat{p}(x)$ and the correct probability $D(1|x)$ give incorrect predictions whereas the other classifiers give correct predictions. As a result, the integrand of equation 2 is non-zero only for $\hat{p}(x) < p < D(1|x)$ (if the classifiers err in one direction) or $D(1|x) < p < \hat{p}(x)$ (if the classifiers err in the other direction). Thus, we can re-write equation 2 as

$$E_{x \sim D} \left| \int_{\hat{p}(x)}^{D(1|x)} |D(1|x) - p| d_p \right| \leq \max\{D(1), D(0)\} \epsilon.$$

By solving the integral, we get

$$E_{x \sim D}(D(1|x) - \hat{p}(x))^2 \leq 2 \max\{D(1), D(0)\}\epsilon. \blacksquare$$

In this theorem, we measure the probabilistic loss using the true distribution $D(1|x)$. Because this distribution is unknown for most real-world problems, it is important to understand the implications of the theorem to metrics that only use the class labels. We consider these implications for three of these metrics in the following subsections.

4.1 The squared error metric for classification

The loss in the squared error metric for classification is defined as

$$E_{x, y \sim D}(y - \hat{p}(x))^2.$$

Note that

$$\begin{aligned} & E_{x, y \sim D}(y - D(1|x))^2 + (\hat{p}(x) - D(1|x))^2 \\ &= E_{x \sim D} D(1|x) [(1 - D(1|x))^2 + (\hat{p}(x) - D(1|x))^2 \\ &\quad + (1 - D(1|x)) [D(1|x)^2 + (\hat{p}(x) - D(1|x))^2]] \\ &= E_x D(1|x) [1 - 2D(1|x)] + D(1|x)^2 + (\hat{p}(x) - D(1|x))^2 \\ &= E_x D(1|x) + \hat{p}(x)^2 - 2\hat{p}(x)D(1|x) \\ &= E_x (\hat{p}(x) - D(1|x))^2 + D(1|x) - D(1|x)^2 \\ &= E_x (\hat{p}(x) - D(1|x))^2 + C \end{aligned}$$

where C is some problem dependent constant.

Consequently, optimizing the squared error metric for classification implicitly optimizes the squared error in the probability estimate.

4.2 The cross entropy metric for classification

The loss in the cross entropy metric is $E_{x, y \sim D} I(y = 1) \ln \frac{1}{\hat{p}(x)} + I(y = 0) \ln \frac{1}{1 - \hat{p}(x)}$. Cross entropy differs fundamentally from the squared error metric because it is (theoretically) unbounded. Consequently, the optimal strategy of an adversarial binary classifier is to err simultaneously with all classifiers on one side of $D(1|x)$, producing an unbounded loss when the discretization goes to 0. For cross entropy we can only prove a much weaker theorem: minimal error implies $\hat{p}(x)$ is optimal up to the discretization.

In practice, this worst case behavior can be avoided. The asymptote to ∞ is very slow as $\hat{p}(x)$ approaches 0 or 1. Consequently, we can project $\hat{p}(x)$ into the nearest element in the interval $[0.001, 0.999]$ and incur a maximal loss of about $\log(1/0.001)\epsilon + 0.001 = 6.9\epsilon + 0.001$ (when $D(1) = D(0)$), while preserving almost all of the dynamic range.

4.3 Implications for ordering metrics

Relative Ordering. In some situations, we may be interested only in the relative ordering or ranking of

examples given by the class probability estimates. Unfortunately, very small errors in the ordering with respect to $[0, 1]$ can result in very large misorderings with respect to a subset of $[0, 1]$. In particular, suppose that the optimal answer is $\hat{p}(x) = 0.5$, but the answer we provide is $\hat{p}(x) = 0.49999$. If every other element in the ordering is in the interval $(0.49999, 0.5)$, then this small error results in a very large ordering error with respect to the set of elements.

AUC analysis. Area under the ROC curve (AUC) is one commonly used relative ordering metric. Suppose we have one example with label $y = 1$, $D(y = 1|x) = 1$ and Probing estimates $\hat{P}(y = 1|x) = 0$, and n examples with label $y = 0$, $D(Y = 1|x) = 1/n$ where Probing estimates $\hat{P}(y = 1|x) = 1/n$. In this setting, the AUC (as defined in [1]) is 0. Despite this terrible performance, the classifiers need err only once. This is less than any fixed importance weighted loss rate for sufficiently large n .

5 Experimental Results

Here we present the results obtained by applying the Probing reduction to fifteen (binary) datasets: eleven from the UCI machine learning repository [5] (Adult, Australian Credit, Breast, Diabetes, Echocardiogram, Hepatitis, Ionosphere, Kr-vs-kp, Liver, Mushroom, Sick), two from the UCI KDD archive [8] (KDD Cup 98 and COIL) and two from the 2004 KDD Cup (Biology and Physics).

We use four different base classifier learners available within the machine learning tool Weka [15]: the J48 decision tree learner, the SMO support vector machine (SVM) learner (linear kernel), Naive Bayes and logistic regression. For choosing the weights, we use the Probing variant that attempts to optimize cross entropy for 100 iterations. For realizing the importance weighted classification, we use rejection sampling (proportionally to the weights) for the decision tree learner and for logistic regression, and give the weights directly to the learners for SVM and Naive Bayes (see [17]).

We are interested in comparing the performance of the Probing reduction with standard methods for obtaining probability estimates from these classifiers. Decision trees and Naive Bayes produce an internal score in the interval $[0, 1]$ that can be used as a class probability estimate, but they are known to be inaccurate [16]. For this reason, besides testing their performance, we have also tested well-known methods for obtaining better probability estimates from these classifiers: bagging for decision trees (100 classifiers) and calibration using a sigmoid function for Naive Bayes [3]. Since the SVM classifier does not produce probability estimates, we have tested only the standard method for obtaining probabilities from SVM (also available in Weka): calibration using a sigmoid function [11].

For datasets that have a standard training/test split in the UCI repository, we use the standard split. For the other datasets, we randomly split the data into a training set with 66% of the examples and a test set with 33% of the examples.

We use three different metrics for assessing the accuracy of the probability estimates: root mean squared error (RMS), cross entropy (CXE) and area under the ROC curve (AUC). For more details about these metrics and an analysis of their characteristics, see [4]. The results on the test set using CXE, RMS and AUC are shown, respectively, in tables 1, 2 and 3. Note that for RMS and CXE the smaller the value the better the performance, while for AUC the larger the value the better the performance. Note also that AUC only measures how well the probability estimates rank the examples from the most probable to the least probable member of the class, while RMS and CXE also measure the calibration of the estimates.

Decision tree results. The Probing reduction outperforms a single J48 classifier for all datasets but Mushroom and for the majority of the datasets also outperforms Bagging J48, on the three metrics. A cross entropy of ∞ means that the classifier predicted probability one for an example whose label is zero (or vice-versa). This is often the case for the J48 classifier, showing that its internal scores are not reasonable class probability estimates for cross entropy loss. Bagging solves this problem by taking advantage of the instability of the learner over different training sets to average out extreme results. However, in cases where there is no instability (like the KDD Cup 98 dataset) bagging performs the same as the base classifiers. The same would be true if we apply bagging to other learners that are stable such as SVM and Naive Bayes.

SVM Results. The Probing reduction outperforms the SVM with sigmoid calibration for most datasets on all three metrics. The AUC metric results show a greater advantage of Probing over sigmoid calibration than the two other metrics. This is probably due to the fact that the sigmoid calibration maintains the same ranking as given by the SVM margins while Probing can give a better ranking based on the results of the different weighted classifiers.

Naive Bayes results. The Probing reduction outperforms a single Naive Bayes classifier for most datasets on the cross entropy and squared error metrics. On the AUC metric, however, a single Naive Bayes classifier usually performs slightly better than Probing Naive Bayes. This may be due to Naive Bayes estimates ranking well (although with terrible probability estimates), while Probing estimates have less resolution leading to more ties. The results also show that Naive Bayes with sigmoid calibration outperforms Probing Naive Bayes on most datasets in all three metrics. Naive Bayes is not a very accurate classifiers so

Dataset	RMS	CXE	AUC
Adult	ProbJ48	ProbJ48	ProbJ48
Australian	ProbJ48	ProbJ48	ProbJ48
Biology	SigSVM	SigSVM	Prob(J48/Log)
Breast	SigNB	SigNB	ProbNB
COIL	ProbJ48	ProbJ48	ProbJ48
Diabetes	Log	Log	ProbSVM
Echocardio	SigNB	SigNB	Log
Hepatitis	BagJ48	ProbSVM	ProbNB
Ionosphere	BagJ48	BagJ48	BagJ48
KDD-98	ProbJ48	ProbJ48	ProbLog
Kr-vs-kp	BagJ48	BagJ48	BagJ48
Liver	BagJ48	BagJ48	BagJ48
Mushroom	J48/SVM	J48/SVM	J48/SVM
Physics	ProbJ48	ProbJ48	ProbJ48
Sick	BagJ48	BagJ48	ProbJ48

Table 4: Best method per dataset/metric.

the Probing reduction, which relies on that accuracy for different weight settings, is not optimal. On the other hand, the sigmoid approach has a separate fitting stage which can repair bad estimates to some extent. Furthermore, the sigmoid transformation maintains approximately the same ranking of examples as the original Naive Bayes classifier, leading to approximately the same AUC.

Logistic Regression. The Probing reduction performs better than a single Logistic Regression classifier for optimizing cross-entropy. In some cases, Logistic Regression yields a cross entropy of ∞ , while Probing gives a reasonable result. For the two other metrics there is not a clear difference.

To get a better idea of which algorithms are performing best, we have one last table consisting of the method that led to best test performance for each metric and dataset in table 4. Probing wins in 20 out of 45 experiments against the other tested alternatives (J48, Naive Bayes, Sigmoid Naive Bayes, Sigmoid SVM, Bagged Decision Tree and Logistic Regression). It is also interesting to note that the performance of the decision tree, J48, is very good. This may be due³ to the fact that the other learners that we are using are linear (we have not tried different kernels for SVM).

Finally, to give an idea of the speed of convergence of Probing in practice, we plot the cross entropy performance (using J48) against the number of Probing iterations for 5 representative datasets in figure 1. Probing seems to converge to a reasonable probability estimate after 20 or so iterations.

6 Discussion

Probing is a general technique for converting any classifier learner into a class probability estimator.

³It's also possible that the design of J48 (which inherits from C4.5) is overfit to the UCI datasets.

Dataset	J48	BagJ48	ProbJ48	SigSVM	ProbSVM	NB	SigNB	ProbNB	Log	ProbLog
Adult	∞	0.499	0.436	0.472	0.479	0.949	0.521	0.464	∞	0.460
Australian	∞	0.444	0.442	0.487	0.471	∞	0.637	0.656	0.531	0.503
Biology	∞	0.018	0.017	0.016	0.021	∞	0.071	0.057	0.018	0.016
Breast	∞	0.793	0.791	0.837	0.809	0.840	0.753	0.830	1.07	0.850
COIL	0.326	0.323	0.301	0.325	0.311	0.357	0.305	0.350	0.304	0.303
Diabetes	∞	0.806	0.784	0.758	0.775	∞	0.861	0.780	0.746	0.774
Echocardio	1.216	0.924	0.897	0.891	0.837	0.902	0.810	0.852	0.829	0.847
Hepatitis	0.733	0.655	0.665	0.695	0.660	1.416	0.733	0.953	2.45	1.21
Ionosphere	∞	0.210	0.287	0.418	0.416	∞	0.353	0.304	0.303	0.297
KDD-98	0.289	0.289	0.283	0.285	0.285	∞	0.285	0.351	0.283	0.283
Kr-vs-kp	∞	0.026	0.069	0.173	0.128	0.416	0.406	0.418	0.145	0.103
Liver	∞	0.832	0.888	0.885	0.891	0.940	0.917	0.950	0.882	0.892
Mushroom	0.000	0.000	0.006	0.000	0.001	0.191	0.166	0.298	0.000	0.003
Physics	∞	0.763	0.762	0.802	0.805	∞	0.880	0.933	0.792	0.783
Sick	∞	0.070	0.086	0.171	0.203	0.219	0.190	0.183	∞	0.164

Table 1: Cross entropy results. The best results for each classifier (J48, SVM and Naive Bayes) are in bold.

Dataset	J48	BagJ48	ProbJ48	SigSVM	ProbSVM	NB	SigNB	ProbNB	Log	ProbLog
Adult	0.347	0.331	0.310	0.324	0.326	0.335	0.332	0.320	0.321	0.320
Australian	0.328	0.303	0.298	0.320	0.313	0.382	0.358	0.363	0.321	0.322
Biology	0.061	0.048	0.048	0.047	0.052	0.075	0.094	0.057	0.050	0.049
Breast	0.439	0.426	0.426	0.439	0.436	0.433	0.415	0.432	0.457	0.448
COIL	0.237	0.236	0.232	0.236	0.234	0.254	0.233	0.254	0.233	0.233
Diabetes	0.474	0.429	0.428	0.412	0.424	0.460	0.450	0.429	0.410	0.416
Echocardio	0.539	0.470	0.466	0.464	0.447	0.445	0.435	0.449	0.448	0.454
Hepatitis	0.390	0.383	0.368	0.368	0.358	0.432	0.400	0.404	0.440	0.449
Ionosphere	0.303	0.201	0.215	0.295	0.284	0.259	0.251	0.238	0.304	0.297
KDD-98	0.219	0.219	0.218	0.219	0.218	0.233	0.219	0.247	0.218	0.218
Kr-vs-kp	0.093	0.065	0.091	0.179	0.160	0.300	0.295	0.300	0.158	0.142
Liver	0.560	0.439	0.461	0.459	0.461	0.473	0.461	0.479	0.457	0.456
Mushroom	0.000	0.001	0.012	0.000	0.003	0.182	0.170	0.201	0.000	0.008
Physics	0.529	0.425	0.424	0.435	0.434	0.490	0.490	0.477	0.432	0.430
Sick	0.125	0.116	0.119	0.175	0.205	0.199	0.187	0.190	0.175	0.172

Table 2: Squared error results. The best results for each classifier (J48, SVM and Naive Bayes) are in bold.

Dataset	J48	BagJ48	ProbJ48	SigSVM	ProbSVM	NB	SigNB	ProbNB	Log	ProbLog
Adult	0.852	0.891	0.912	0.901	0.901	0.906	0.906	0.905	0.902	0.904
Australian	0.900	0.937	0.944	0.912	0.924	0.920	0.920	0.917	0.918	0.926
Biology	0.885	0.976	0.987	0.986	0.982	0.957	0.922	0.956	0.981	0.987
Breast	0.606	0.602	0.633	0.680	0.659	0.723	0.723	0.724	0.642	0.611
COIL	0.500	0.641	0.720	0.542	0.682	0.709	0.709	0.717	0.706	0.708
Diabetes	0.715	0.796	0.798	0.817	0.824	0.779	0.779	0.805	0.820	0.819
Echocardio	0.543	0.644	0.645	0.704	0.743	0.759	0.759	0.753	0.770	0.755
Hepatitis	0.654	0.670	0.662	0.749	0.748	0.743	0.743	0.757	0.629	0.667
Ionosphere	0.895	0.981	0.966	0.926	0.928	0.960	0.960	0.957	0.905	0.952
KDD-98	0.500	0.500	0.616	0.598	0.602	0.606	0.606	0.604	0.616	0.617
Kr-vs-kp	0.998	1.000	0.999	0.992	0.995	0.953	0.953	0.953	0.995	0.998
Liver	0.607	0.754	0.724	0.711	0.714	0.713	0.713	0.706	0.710	0.711
Mushroom	1.000	1.000	1.000	1.000	1.000	0.998	0.998	0.990	1.000	1.000
Physics	0.674	0.803	0.803	0.785	0.789	0.738	0.738	0.729	0.788	0.791
Sick	0.965	0.985	0.989	0.946	0.940	0.949	0.949	0.952	0.935	0.952

Table 3: Area under the ROC results. The best results for each classifier (J48, SVM and Naive Bayes) are in bold.

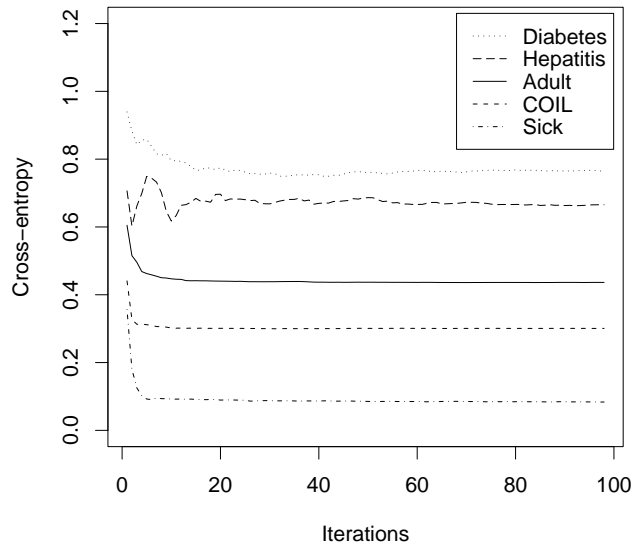


Figure 1: Cross entropy vs. the number of Probing iterations for 5 representative datasets (using J48).

Theorem 2 shows that the probing reduction is well founded: good performance on the created classification problems implies good performance with respect to Probing probability estimates. Experimental results show that Probing achieves strong performance relative to a variety of other (typically more specialized) techniques using a variety of classifier learners. This combination of positive results is unmatched by any alternative method.

In future work, we would like to validate the method on multiclass problems as outlined in section 3. Also, further studying the behavior of different discretization schemes for choosing the weights could lead to enhancements of the Probing estimates for different performance criteria.

References

[1] S. Agarwal, T. Graepel, R. Herbrich, S. Har-Peled and D. Roth (2004). Generalization bounds for the Area under an ROC curve. Technical report UIUCDCS-R-2004-2433, Computer Science Department, University of Illinois at Urbana-Champaign.

[2] P. Bartlett and A. Tewari (2004). Sparseness versus estimating conditional probabilities: Some asymptotic results. In *Proceedings of the 17th Annual Conference on Learning Theory (COLT-2004)*, 564-578. Springer-Verlag.

[3] P. Bennett (2000). Assessing the Calibration of Naive Bayes' Posterior Estimates (2000). Technical Re-

port CMU-CS-00-155, Computer Science Department, Carnegie Mellon University.

[4] R. Caruana, A. Niculescu-Mizil (2004). Data mining in metric space: an empirical analysis of supervised learning performance criteria. In *Proceedings of the Tenth International Conference on Knowledge Discovery and Data Mining (KDD-2004)*, 69-78. ACM Press.

[5] C. L. Blake and C. J. Merz (1998). UCI Repository of machine learning databases [<http://www.ics.uci.edu/~mllearn/MLRepository>]. Irvine, CA: University of California, Department of Information and Computer Science.

[6] Y. Freund and R. E. Schapire (1997). A decision-theoretic generalization of on-line learning and an application to boosting. In *Journal of Computer and System Sciences*, 55:1, 119-139.

[7] O. M. Halck (2002). Using Hard Classifiers to Estimate Conditional Class Probabilities. In *Proceedings of the European Conference on Machine Learning (ECML-2002)*, 124-134, Springer.

[8] S. Hettich and S. D. Bay (1999). The UCI KDD Archive [<http://kdd.ics.uci.edu>]. Irvine, CA: University of California, Department of Information and Computer Science.

[9] Y. LeCun, L. Bottou, G. Orr, and K. Muller (1998). Efficient BackProp. In G. Orr and K. Muller (eds.) *Neural Networks: Tricks of the trade*. Springer-Verlag.

[10] D. Margineantu (2000). On class probability estimates and cost-sensitive evaluation of classifiers. In *Workshop on Cost-Sensitive Learning, The Seventeenth International Conference on Machine Learning (ICML-2000)*.

[11] J. Platt (1999). Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. In A. Smola, P. Bartlett, B. Schölkopf, D. Schuurmans (eds.) *Advances in Large Margin Classifiers*, 61-74. MIT Press.

[12] F. Provost and P. Domingos (2003). Tree Induction for Probability-based Rankings. In *Machine Learning*, 52:3, 199-216. Kluwer.

[13] B. Schölkopf and A. J. Smola (2002). *Learning with Kernels*. MIT Press.

[14] J. R. Quinlan (1993). *C4.5: Programs for Machine Learning*. San Mateo CA: Morgan Kaufmann.

[15] I. H. Witten and E. Frank (1999). *Data mining: practical machine learning tools and techniques with Java implementations* [<http://www.cs.waikato.ac.nz/ml/weka/>]. Morgan Kaufmann.

[16] B. Zadrozny and C. Elkan (2001). Obtaining calibrated probability estimates from decision trees and naive Bayesian classifiers. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML-01)*, 609-616. Morgan Kaufmann.

[17] B. Zadrozny, J. Langford, and N. Abe (2003). Cost Sensitive Learning by Cost-Proportionate Example Weighting. In *Proceedings of the 2003 IEEE International Conference on Data Mining (ICDM-03)*, 435-442. IEEE Press.