

PAC Classification based on PAC Estimates of Label Class Distributions*

Nick Palmer and Paul W. Goldberg

Dept. of Computer Science, University of Warwick, Coventry CV4 7AL, U.K.
(npalmer|pwg@dcs.warwick.ac.uk),
Research group home page: <http://www.dcs.warwick.ac.uk/research/acrg>

Abstract. A standard approach in pattern classification is to estimate the distributions of the label classes, and then to apply the Bayes classifier to the estimates of the distributions in order to classify unlabeled examples. As one might expect, the better our estimates of the label class distributions, the better the resulting classifier will be. In this paper we make this observation precise by identifying risk bounds of a classifier in terms of the quality of the estimates of the label class distributions. We show how PAC learnability relates to estimates of the distributions that have a PAC guarantee on their L_1 distance from the true distribution, and we bound the increase in negative log likelihood risk in terms of PAC bounds on the KL-divergence. We give an inefficient but general-purpose smoothing method for converting an estimated distribution that is good under the L_1 metric into a distribution that is good under the KL-divergence.

keywords. Bayes error, Bayes classifier, plug-in decision function

1 Introduction

We consider a general approach to pattern classification in which elements of each class are first used to train a probabilistic model via some unsupervised learning method. The resulting models for each class are then used to assign discriminant scores to an unlabeled instance, and a label is chosen to be the one associated with the model giving the highest score. For example [3] uses this approach to classify protein sequences, via training a well-known probabilistic suffix tree model of Ron et al. [18] on each sequence class. Indeed, even where an unsupervised technique is mainly being used to gain insight into the process that generated two or more data sets, it is still sometimes instructive to try out the associated classifier, since the misclassification rate provides a quantitative measure of the accuracy of the estimated distributions.

* This work was supported by EPSRC Grant GR/R86188/01. This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778. This publication only reflects the authors' views.

The work of [18] has led to further related algorithms for learning classes of probabilistic finite state automata (PDFAs) in which the objective of learning has been formalized as the estimation of a true underlying distribution (over strings output by the *target PDFA*) with a distribution represented by a hypothesis PDFA. The natural discriminant score to assign to a string, is the probability that the hypothesis would generate that string at random.

As one might expect, the better one’s estimates of label class distributions (the class-conditional densities), the better should be the associated classifier. The contribution of this paper is to make precise that observation. We give bounds on the risk of the associated Bayes classifier¹ in terms of the quality of the estimated distributions.

These results are partly motivated by our interest in the relative merits of estimating a class-conditional distribution using the variation distance, as opposed to the KL-divergence (defined in the next section). In [4] it has been shown how to learn a class of PDFAs using KL-divergence, in time polynomial in a set of parameters that includes the expected length of strings output by the automaton. In [15] we show how to learn this class with respect to variation distance, with a polynomial sample-size bound that is independent of the length of output strings. Furthermore, it can be shown that it is necessary to switch to the weaker criterion of variation distance, in order to achieve this. We show here that this leads to a different—but still useful—performance guarantee for the Bayes classifier.

Abe and Warmuth [2] study the problem of learning probability distributions using the KL-divergence, via classes of probabilistic automata. Their criterion for learnability is that—for an unrestricted input distribution D —the hypothesis PDFA should be almost (i.e. within ϵ) as close as possible to D . Abe, Takeuchi and Warmuth [1] study the negative log-likelihood loss function in the context of learning *stochastic rules*, i.e. rules that associate an element of the domain X to a probability distribution over the range Y . We show here that if two or more label class distributions are learnable in the sense of [2], then the resulting stochastic rule (the conditional distribution over Y given $x \in X$) is learnable in the sense of [1].

We show that if instead the label class distributions are well estimated using the variation distance, then the associated classifier may not have a good negative log likelihood risk, but will have a *misclassification rate* that is close to optimal. This result is for general k -class classification, where distributions may overlap (i.e. the optimum misclassification rate may be positive). We also incorporate variable misclassification penalties (sometimes one might wish a false positive to cost more than a false negative), and show that this more general loss function is still approximately minimized provided that discriminant likelihood scores are rescaled appropriately.

¹ The *Bayes classifier* associated with two or more probability distributions is the function that maps an element x of the domain to the label associated with the probability distribution whose value at x is largest. This is of course a well-known approach for classification, see [7].

As a result we show that PAC-learnability and more generally p-concept² learnability [12], follows from the ability to learn class distributions in the setting of Kearns et al. [11]. Papers such as [5, 14, 8] study the problem of learning various classes of probability distributions with respect to KL-divergence and variation distance, in this setting.

It is well-known (noted in [12]) that learnability with respect to KL-divergence is stronger than learnability with respect to variation distance. Furthermore, the KL-divergence is usually used (for example in [4, 10]) due to the property that when minimized with respect to an sample, the empirical likelihood of that sample is maximized. An algorithm that learns with respect to variation distance can sometimes be converted to one that learns with respect to KL-divergence by a smoothing technique [5], when the domain is $\{0, 1\}^n$, and n is a parameter of the learning problem. In this paper we give a related smoothing rule that applies to the version of the PDFFA learning problem where we seem to “need” to use the variation distance. However, the smoothed distribution does not have an efficient representation, and requires the probabilities used in the target PDFFA to have limited precision.

1.1 Notation and Terminology

In k -class classification, labeled examples are generated by distribution D over $X \times \{1, \dots, k\}$. We consider the problem of predicting the label ℓ associated with $x \in X$, where x is generated by the marginal distribution of D on X , $D|_X$. A non-negative cost is incurred for each classification, based either on a cost matrix (where the cost depends upon both the hypothesized label and the true label) or the negative log-likelihood of the true label being assigned. The aim is to optimize the expected cost given by the occurrence of a randomly generated example. We refer to the expected cost associated with any classifier $f : X \rightarrow \{1, \dots, k\}$, as *risk* (as described by Vapnik [17]), denoted as $R(f)$.

Let D_ℓ be D restricted to points (x, ℓ) , $\ell = \{1, \dots, k\}$. D is a mixture $\sum_{\ell=1}^k g_\ell D_\ell$, where $\sum_{i=1}^k g_i = 1$, and g_ℓ is the *class prior* of class ℓ —the probability that a randomly generated data point has label ℓ .

In Section 2 it is shown that if we have upper bounds on the inaccuracy of the estimated distributions of each class label, then we can derive bounds on the risk associated with the classifiers. Suppose D and D' are probability distributions over the same domain X . We define the L_1 distance as $L_1(D, D') = \int_X |D(x) - D'(x)| dx$. We usually assume that X is a discrete domain, in which case

$$L_1(D, D') = \sum_{x \in X} |D(x) - D'(x)|.$$

The KL-divergence from D to D' is defined as

$$I(D||D') = \sum_{x \in X} D(x) \log \left(\frac{D(x)}{D'(x)} \right).$$

² p-concepts are functions probabilistically mapping elements of the domain to 2 classes.

1.2 Learning Framework

In the PAC-learning framework an algorithm receives labeled samples generated independently according to distribution D over X , where distribution D is unknown, and where labels are generated by an unknown function f from a known class of functions \mathcal{F} . The algorithm must output a hypothesis h from a class of hypotheses \mathcal{H} , such that with probability at least $1 - \delta$, $err_h \leq \epsilon$, where ϵ and δ are parameters. Notice that in this setting, if $f \in \mathcal{H}$, then $err^* = 0$, where err^* is the error associated with the optimal hypothesis.

We use a variation on the framework used in [12] for learning p-concepts, which adopts performance measures from the PAC model, extending this to learn stochastic rules with k classes. Therefore it is the case that $err^* = \inf_{h \in \mathcal{H}} \{err_h\}$. The aim of the learning algorithm in this framework is to output a hypothesis $h \in \mathcal{H}$ such that with probability of at least $1 - \delta$, the error err_h of h satisfies $err_h \leq err^* + \epsilon$.

Our notion of learning distributions is similar to that of Kearns et al. [11].

Definition 1. Let \mathcal{D}_n be a class of distributions. \mathcal{D}_n is said to be efficiently learnable if an algorithm A exists, such that given $\epsilon > 0$ and $\delta > 0$ and access to randomly drawn examples (see below) from any unknown target distribution $D \in \mathcal{D}_n$, A runs in time polynomial in $(\frac{1}{\epsilon})$, $(\frac{1}{\delta})$ and n and returns a probability distribution D' that with probability at least $1 - \delta$ is within L_1 -distance (alternatively KL-divergence) ϵ of D .

We define p-concepts as introduced by Kearns and Shapire [12]. This definition is for 2-class classification, but generalizes in a natural way to more than 2 classes.

Definition 2. A Probabilistic Concept (or p-concept) f on domain X is given by a real-valued function $p_f : X \rightarrow [0, 1]$. An observation of f consists of some $x \in X$ together with a 0/1 label ℓ with $\Pr(\ell = 1) = p_f(x)$.

2 Results

In Section 2.1 we give bounds on the risk associated with a hypothesis, with respect to the accuracy of the approximation of the underlying distribution generating the instances. In Section 2.2 we show that these bounds are close to optimal, and in Section 2.3 we give corollaries showing what these bounds mean for PAC learnability.

We define the accuracy of an approximate distribution in terms of L_1 distance and KL divergence, both of which are commonly used measurements. It is assumed that the class priors of each class label are known.

2.1 Bounds on Increase in Risk

First we examine the case where the accuracy of the hypothesis distribution is such that the distribution for each class label is within L_1 distance ϵ of the true

distribution for that label, for some $0 \leq \epsilon \leq 1$. A cost matrix C specifies the cost associated with any classification, where the cost of classifying a data point which has label i as some label j is denoted as c_{ij} (where $c_{ij} \geq 0$). It is usually the case that $c_{ij} = 0$ for $i = j$. We introduce the following notation:

Given classifier f over discrete domain X , $f : X \rightarrow \{1, \dots, k\}$, the risk of f is given by

$$R(f) = \sum_{x \in X} \sum_{i=1}^k c_{if(x)} \cdot g_i \cdot D_i(x).$$

Let f^* be the Bayes optimal classifier, i.e. the function with the minimal risk, or optimal expected cost, and $f'(x)$ is the function with optimal expected cost with respect to alternative distributions $D'_i, i \in \{1, \dots, k\}$. For $x \in X$,

$$\begin{aligned} f^*(x) &= \arg \min_j \sum_{i=1}^k c_{ij} \cdot g_i \cdot D_i(x) \\ f'(x) &= \arg \min_j \sum_{i=1}^k c_{ij} \cdot g_i \cdot D'_i(x). \end{aligned}$$

Theorem 1. ³ *Let f^* be the Bayes optimal classifier and let f' be the Bayes classifier associated with estimated distributions D'_i . Suppose that for each label $i \in \{1, \dots, k\}$, $L_1(D_i, D'_i) \leq \epsilon/g_i$. Then $R(f') \leq R(f^*) + \epsilon \cdot k \cdot \max_{ij} \{a_{ij}\}$.*

Proof. Let $R_f(x)$ be the contribution from $x \in X$ towards the total expected cost associated with classifier f . For f such that $f(x) = j$,

$$R_f(x) = \sum_{i=1}^k c_{ij} \cdot g_i \cdot D_i(x).$$

Let $\tau_{\ell' - \ell}(x)$ be the increase in risk for labelling x as ℓ' instead of ℓ , so that

$$\begin{aligned} \tau_{\ell' - \ell}(x) &= \sum_{i=1}^k c_{i\ell'} \cdot g_i \cdot D_i(x) - \sum_{i=1}^k c_{i\ell} \cdot g_i \cdot D_i(x) \\ &= \sum_{i=1}^k (c_{i\ell'} - c_{i\ell}) \cdot g_i \cdot D_i(x). \end{aligned} \tag{1}$$

Note that due to the optimality of f^* on D_i , $\forall x \in X : \tau_{f'(x) - f^*(x)}(x) \geq 0$. In a similar way, the expected contribution to the total cost of f' from x must be less than or equal to that of f^* with respect to D'_i – given that f' is chosen to be optimal on the D'_i values. We have:

$$\sum_{i=1}^k c_{if'(x)} \cdot g_i \cdot D'_i(x) \leq \sum_{i=1}^k c_{if^*(x)} \cdot g_i \cdot D'_i(x).$$

³ This result is essentially a generalization of Exercise 2.10 of Devroye et al's textbook [6], from 2 class to multiple classes, and in addition we show here that variable misclassification costs can be incorporated. This is the closest thing we have found to this Theorem that has already appeared, but we suspect that other related results may have appeared. We would welcome any further information or references on this topic. Theorem 2 is another result which we suspect may be known, but likewise we have found no statement of it.

Rearranging, we have

$$\sum_{i=1}^k D'_i(x) \cdot g_i \cdot (c_{if^*(x)} - c_{if'(x)}) \geq 0. \quad (2)$$

From (1) and (2) it can be seen that

$$\begin{aligned} \tau_{f'(x)-f^*(x)}(x) &\leq (D_i - D'_i(x)) \cdot g_i \cdot (c_{if'(x)} - c_{if^*(x)}) \\ &\leq \sum_{i=1}^k |(D_i - D'_i(x))| \cdot g_i \cdot |(c_{if'(x)} - c_{if^*(x)})|. \end{aligned}$$

Let $d_i(x)$ be the difference between the probability densities of D_i and D'_i at $x \in X$, $d_i(x) = |D_i(x) - D'_i(x)|$. Therefore,

$$\tau_{f'(x)-f^*(x)}(x) \leq \sum_{i=1}^k |c_{if'(x)} - c_{if^*(x)}| \cdot g_i \cdot d_i(x) \leq \tau_{f'(x)-f^*(x)}(x) \leq \sum_{i=1}^k \max_j \{c_{ij}\} \cdot g_i \cdot d_i(x).$$

In order to bound the expected cost, it is necessary to sum over the range of $x \in X$:

$$\sum_{x \in X} \tau_{f'(x)-f^*(x)}(x) \leq \sum_{x \in X} \sum_{i=1}^k \max_j \{c_{ij}\} \cdot g_i \cdot d_i(x) = \sum_{i=1}^k \max_j \{c_{ij}\} \cdot g_i \cdot \sum_{x \in X} d_i(x). \quad (3)$$

Since $L_1(D_i, D'_i) \leq \epsilon/g_i$ for all i , ie. $\sum_{x \in X} d_i(x) \leq \epsilon/g_i$, it follows from (3) that

$$\sum_{x \in X} \tau(x) \leq \sum_{i=1}^k \max_j \{c_{ij}\} \cdot g_i \cdot \left(\frac{\epsilon}{g_i}\right).$$

This expression gives an upper bound on expected cost for labelling x as $f'(x)$ instead of $f^*(x)$. By definition,

$$\sum_{x \in X} \tau(x) = R(f') - R(f^*).$$

Therefore it has been shown that

$$R(f') \leq R(f^*) + \epsilon \cdot \sum_{i=1}^k \max_j \{c_{ij}\} \leq R(f^*) + \epsilon \cdot k \cdot \max_{ij} \{c_{ij}\}.$$

□

We next prove a corresponding result in terms of KL-divergence, which uses the negative log-likelihood of the correct label as the cost function. We define $\text{Pr}_i(x)$ to be the probability that a data point at x has label i , such that $\text{Pr}_i(x) = g_i \cdot D_i(x) \left(\sum_{j=1}^k g_j \cdot D_j(x)\right)^{-1}$. Given a function $f : X \rightarrow \mathbf{R}^k$, where

$f(x)$ is a prediction of the probabilities of x having each label $i \in \{1, \dots, k\}$ (so $\sum_{i=1}^k f_i(x) = 1$), the risk associated with f can be expressed as

$$R(f) = \sum_{x \in X} D(x) \sum_{i=1}^k -\log(f_i(x)) \cdot \text{Pr}_i(x). \quad (4)$$

Let $f^* : X \rightarrow \mathbf{R}^k$ output the true class label distribution for an element of X . From Equation (4) it can be seen that

$$R(f^*) = \sum_{x \in X} D(x) \sum_{i=1}^k -\log(\text{Pr}_i(x)) \cdot \text{Pr}_i(x). \quad (5)$$

Theorem 2. For $f : X \rightarrow \mathbf{R}^k$ suppose that $R(f)$ is given by (4). If for each label $i \in \{1, \dots, k\}$, $I(D_i || D'_i) \leq \epsilon/g_i$, then $R(f') \leq R(f^*) + k\epsilon$.

Proof. Let $R_f(x)$ be the contribution at $x \in X$ to the risk associated with classifier f , $R_f(x) = \sum_{i=1}^k -\log(f_i(x)) \cdot \text{Pr}_i(x)$. Therefore $R(f') = \sum_{x \in X} D(x) \cdot R_{f'}(x)$.

We define $\text{Pr}'_i(x)$ to be the estimated probability that a data point at $x \in X$ has label $i \in \{1, \dots, k\}$, from distributions D'_i , such that $\text{Pr}'_i(x) = g_i \cdot D'_i \left(\sum_{j=1}^k g_j \cdot D'_j(x) \right)^{-1}$.

$$R_{f'}(x) = D(x) \cdot \sum_{i=1}^k -\log(\text{Pr}'_i(x)) \cdot \text{Pr}_i(x).$$

Let $\xi(x)$ denote the contribution to additional risk incurred from using f' as opposed to f^* at $x \in X$. From (5) it can be seen that

$$\begin{aligned} \xi(x) &= R_{f'}(x) - D(x) \cdot \sum_{i=1}^k -\log(\text{Pr}_i(x)) \cdot \text{Pr}_i(x) \\ &= D(x) \cdot \sum_{i=1}^k \text{Pr}_i(x) \cdot (\log(\text{Pr}_i(x)) - \log(\text{Pr}'_i(x))) \\ &= D(x) \cdot \sum_{i=1}^k \left(\frac{g_i \cdot D_i(x)}{\sum_{j=1}^k g_j \cdot D_j(x)} \right) \left(\log \left(\frac{g_i \cdot D_i(x)}{\sum_{j=1}^k g_j \cdot D_j(x)} \right) - \log \left(\frac{g_i \cdot D'_i(x)}{\sum_{j=1}^k g_j \cdot D'_j(x)} \right) \right) \\ &= D(x) \cdot \sum_{i=1}^k \left(\left(\frac{g_i \cdot D_i(x)}{\sum_{j=1}^k g_j \cdot D_j(x)} \right) \cdot \left(\log \left(\frac{g_i \cdot D_i(x)}{g_i \cdot D'_i(x)} \right) - \log \left(\frac{\sum_{j=1}^k g_j \cdot D_j(x)}{\sum_{j=1}^k g_j \cdot D'_j(x)} \right) \right) \right). \end{aligned}$$

We define D' such that $D'(x) = \sum_{i=1}^k g_i \cdot D'_i(x)$. Since it is the case that $D(x) = \sum_{i=1}^k g_i \cdot D_i(x)$, $\xi(x)$ can be rewritten as

$$\begin{aligned} \xi(x) &= D(x) \cdot \sum_{i=1}^k \left(\frac{g_i \cdot D_i(x)}{D(x)} \right) \cdot \left(\log \left(\frac{g_i \cdot D_i(x)}{g_i \cdot D'_i(x)} \right) - \log \left(\frac{D(x)}{D'(x)} \right) \right) \\ &= \sum_{i=1}^k \left(g_i \cdot D_i(x) \log \left(\frac{D_i(x)}{D'_i(x)} \right) \right) - D(x) \log \left(\frac{D(x)}{D'(x)} \right). \end{aligned}$$

We define $I(D||D')(x)$ to be the contribution at $x \in X$ to the KL-divergence, such that $I(D||D')(x) = D(x) \log(D(x)/D'(x))$. It follows that

$$\sum_{x \in X} \xi(x) = \sum_{i=1}^k (g_i \cdot I(D_i||D'_i)) - I(D||D'). \quad (6)$$

We know that the KL divergence between D_i and D'_i is bounded by ϵ/g_i for each label $i \in \{1, \dots, k\}$, so (6) can be rewritten as

$$\sum_{x \in X} \xi(x) \leq \sum_{i=1}^k \left(g_i \cdot \left(\frac{\epsilon}{g_i} \right) \right) - I(D||D') \leq k \cdot \epsilon - I(D||D').$$

Due to the fact that the KL-divergence between two distributions is non-negative, an upper bound on the cost can be obtained by letting $I(D||D') = 0$, so $R(f') - R(f^*) \leq k\epsilon$. Therefore it has been proved that $R(\hat{f}) \leq R(f^*) + k\epsilon$. \square

2.2 Lower Bounds

In this section we give lower bounds corresponding to the two upper bounds given in Section 2.

Example 1. Consider a distribution D over domain $X = \{x_0, x_1\}$, from which data is generated with labels 0 and 1 and there is an equal probability of each label being generated ($g_0 = g_1 = \frac{1}{2}$). $D_i(x)$ denotes the probability that a point is generated at $x \in X$ given that it has label i . D_0 and D_1 are distributions over X , such that at $x \in X$, $D(x) = \frac{1}{2}(D_0(x) + D_1(x))$.

Suppose that D'_0 and D'_1 are approximations of D_0 and D_1 , and that $L_1(D_0, D'_0) = \frac{\epsilon}{g_0} = 2\epsilon$ and $L_1(D_1, D'_1) = \frac{\epsilon}{g_1} = 2\epsilon$, where $\epsilon = \epsilon' + \gamma$ (and γ is an arbitrarily small constant).

Given the following distributions, assuming that a misclassification results in a cost of 1 and that a correct classification results in no cost, it can be seen that $R(f^*) = \frac{1}{2} - \epsilon'$:

$$\begin{aligned} D_0(x_0) &= \frac{1}{2} + \epsilon', D_0(x_1) = \frac{1}{2} - \epsilon', \\ D_1(x_0) &= \frac{1}{2} - \epsilon', D_1(x_1) = \frac{1}{2} + \epsilon'. \end{aligned}$$

Now if we have approximations D'_0 and D'_1 as shown below, it can be seen that f' will misclassify for every value of $x \in X$:

$$\begin{aligned} D'_0(x_0) &= \frac{1}{2} - \gamma, D'_0(x_1) = \frac{1}{2} + \gamma, \\ D'_1(x_0) &= \frac{1}{2} + \gamma, D'_1(x_1) = \frac{1}{2} - \gamma. \end{aligned}$$

This results in $R(f') = \frac{1}{2} + \epsilon'$. Therefore $R(f') = R(f^*) + 2\epsilon' = R(f^*) + 2(\epsilon - \gamma)$.

In this example the risk is only 2γ under $R(f^*) + \epsilon \cdot k \cdot \max_j \{a_{ij}\}$, since $k = 2$. A similar example can be used to give upper bounds to the lower bound given in Theorem 2.

Example 2. Consider distributions D_0, D_1, D'_0 and D'_1 over domain $X = \{x_0, x_1\}$ as defined in Example 1. It can be seen that the KL-divergence between each label's distribution and its approximated distribution is

$$I(D_0||D'_0) = I(D_1||D'_1) = \left(\frac{1}{2} + \epsilon'\right) \log\left(\frac{\frac{1}{2} + \epsilon'}{\frac{1}{2} - \gamma}\right) + \left(\frac{1}{2} - \epsilon'\right) \log\left(\frac{\frac{1}{2} - \epsilon'}{\frac{1}{2} + \gamma}\right).$$

The optimal risk, measured in terms of negative log-likelihood, can be expressed as $R(f^*) = -\left(\frac{1}{2} + \epsilon'\right) \log\left(\frac{1}{2} + \epsilon'\right) - \left(\frac{1}{2} - \epsilon'\right) \log\left(\frac{1}{2} - \epsilon'\right)$. The risk incurred by using f' as the discriminant function is $R(f') = -\left(\frac{1}{2} + \epsilon'\right) \log\left(\frac{1}{2} - \gamma\right) - \left(\frac{1}{2} - \epsilon'\right) \log\left(\frac{1}{2} + \gamma\right)$. Therefore,

$$R(f') = R(f^*) + \left(\frac{1}{2} + \epsilon'\right) \log\left(\frac{\frac{1}{2} + \epsilon'}{\frac{1}{2} - \gamma}\right) + \left(\frac{1}{2} - \epsilon'\right) \log\left(\frac{\frac{1}{2} - \epsilon'}{\frac{1}{2} + \gamma}\right) = R(f^*) + \epsilon.$$

2.3 Learning near-optimal classifiers in the PAC sense

We show that the results of Section 2.1 imply learnability within the framework defined in Section 1.2.

The following corollaries refer to algorithms A_{class} and $A_{class'}$. These algorithms generate classifier functions $f' : X \rightarrow \{1, 2, \dots, k\}$, which label data in a k -label classification problem, using L_1 distance and KL -divergence respectively as measurements of accuracy.

Corollary 1 shows (using Theorem 1) that a near optimal classifier can be constructed given that an algorithm exists which approximates a distribution over positive data in polynomial time. We are given cost matrix C , and assume knowledge of the class priors g_i .

Corollary 1. *If an algorithm A_{L_1} approximates distributions within L_1 distance ϵ' with probability at least $1 - \delta'$, in time polynomial in $1/\epsilon'$ and $1/\delta'$, then an algorithm A_{class} exists which (with probability $1 - \delta$) generates a discriminant function f' with an associated risk of at most $R(f^*) + \epsilon$, and A_{class} is polynomial in $1/\delta$ and $1/\epsilon$.*

Proof. A_{class} is a classification algorithm which uses unsupervised learners to fit a distribution to each label $i \in \{1, \dots, k\}$, and then uses the Bayes classifier with respect to these estimated distributions, to label data.

A_{L_1} is a PAC algorithm which learns from a sample of positive data to estimate a distribution over that data. A_{class} generates a sample N of data, and divides N into sets $\{N_1, \dots, N_k\}$, such that N_i contains all members of N with label i . Note that for all labels i , $|N_i| \approx g_i \cdot |N|$.

With a probability of at least $1 - \frac{1}{2}(\delta/k)$, A_{L_1} generates an estimate D' of the distribution D_i over label i , such that $L_1(D_i, D') \leq \epsilon (g_i \cdot k \cdot \max_{ij} \{c_{ij}\})^{-1}$.

Therefore the size of the sample $|N_i|$ must be polynomial in $g_i \cdot k \cdot \max_{ij} \{c_{ij}\} / \epsilon$ and k/δ . For all $i \in \{1, \dots, k\}$ $g_i \leq 1$, so $|N_i|$ is polynomial in $\max_{ij} \{c_{ij}\}$, k , $1/\epsilon$ and $1/\delta$.

When A_{class} combines the distributions returned by the k iterations of A_{L_1} , there is a probability of at least $1 - \delta/2$ that all of the distributions are within $\epsilon (g_i \cdot k \cdot \max_{ij} \{c_{ij}\})^{-1} L_1$ distance of the true distributions (given that each iteration received a sufficiently large sample). We allow a probability of $\delta/2$ that the initial sample N did not contain a good representation of all labels ($\neg \forall i \in \{1, \dots, k\} : |N_i| \approx g_i \cdot |N|$), and as such – one or more iteration of A_{L_1} may not have received a sufficiently large sample to learn the distribution accurately.

Therefore with probability at least $1 - \delta$, all approximated distributions are within $\epsilon (g_i \cdot k \cdot \max_{ij} \{c_{ij}\})^{-1} L_1$ distance of the true distributions. If we use the classifier which is optimal on these approximated distributions, f' , then the increase in risk associated with using f' instead of the Bayes Optimal Classifier, f^* , is at most ϵ . It has been shown that A_{L_1} requires a sample of size polynomial in $1/\epsilon$, $1/\delta$, k and $\max_{ij} \{c_{ij}\}$. It follows that

$$|N| = \sum_{i=1}^k |N_i| = \sum_{i=1}^k p \left(\frac{1}{\epsilon}, \frac{1}{\delta}, k, \max_{ij} \{c_{ij}\} \right) \in O \left(p \left(\frac{1}{\epsilon}, \frac{1}{\delta}, k, \max_{ij} \{c_{ij}\} \right) \right).$$

□

Corollary 2 shows (using Theorem 2) how a near optimal classifier can be constructed given that an algorithm exists which approximates a distribution over positive data in polynomial time.

Corollary 2. *If an algorithm A_{KL} has a probability of at least $1 - \delta$ of approximating distributions within ϵ KL-divergence, in time polynomial in $1/\epsilon$ and $1/\delta$, then an algorithm $A_{class'}$ exists which (with probability $1 - \delta$) generates a function f' that maps $x \in X$ to a conditional distribution over class labels of x , with an associated log-likelihood risk of at most $R(f^*) + \epsilon$, and $A_{class'}$ is polynomial in $1/\delta$ and $1/\epsilon$.*

Proof. $A_{class'}$ is a classification algorithm using the same method as A_{class} in Corollary 1, whereby a sample N is divided into sets $\{N_1, \dots, N_k\}$, and each set is passed to algorithm A_{KL} where a distribution is estimated over the data in the set.

With a probability of at least $1 - \frac{1}{2}(\delta/k)$, A_{KL} generates an estimate D' of the distribution D_i over label i , such that $I(D_i || D') \leq \epsilon (g_i \cdot k)^{-1}$. Therefore the size of the sample $|N_i|$ must be polynomial in $g_i \cdot k / \epsilon$ and k/δ . Since $g_i \leq 1$, $|N_i|$ is polynomial in k/ϵ and k/δ .

When $A_{class'}$ combines the distributions returned by the k iterations of A_{KL} , there is a probability of at least $1 - \delta/2$ that all of the distributions are within $\epsilon (g_i \cdot k)^{-1}$ KL-divergence of the true distributions. We allow a probability of $\delta/2$ that the initial sample N did not contain a good representation of all labels ($\neg \forall i \in \{1, \dots, k\} : |N_i| \approx g_i \cdot |N|$).

Therefore with probability at least $1 - \delta$, all approximated distributions are within $\epsilon(g_i.k)^{-1}$ KL-divergence of the true distributions. If we use the classifier which is optimal on these approximated distributions, f' , then the increase in risk associated with using f' instead of the Bayes Optimal Classifier f^* , is at most ϵ . It has been shown that A_{KL} requires a sample of size polynomial in $1/\epsilon$, $1/\delta$ and k . Let $p(1/\epsilon, 1/\delta)$ be an upper bound on the time and sample size used by A_{KL} . It follows that

$$|N| = \sum_{i=1}^k |N_i| = \sum_{i=1}^k p\left(\frac{1}{\epsilon}, \frac{1}{\delta}\right) \in O\left(k.p\left(\frac{1}{\epsilon}, \frac{1}{\delta}\right)\right).$$

□

2.4 Smoothing: from L_1 distance to KL-divergence

Given a distribution that has accuracy ϵ under the L_1 distance, is there a generic way to “smooth” it so that it has similar accuracy under the KL-divergence? From [5] this can be done for $X = \{0, 1\}^n$, if we are interested in algorithms that are polynomial in n in addition to other parameters. Suppose however that the domain is bit strings of unlimited length. Here we give a related but weaker result in terms of bit strings that are used to represent distributions, as opposed to members of the domain. We define class \mathcal{D} of distributions specified by bit strings, such that each member of \mathcal{D} is a distribution on discrete domain X , represented by a discrete probability scale. Let L_D be the length of the bit string describing distribution D . Note that there are at most 2^{L_D} distributions in \mathcal{D} represented by strings of length L_D .

Lemma 1. *Suppose $D \in \mathcal{D}$ is learnable under L_1 distance in time polynomial in δ , ϵ and L_D . Then \mathcal{D} is learnable under KL-divergence, with polynomial sample size.*

Proof. Let D be a member of class \mathcal{D} , represented by a bit string of length L_D , and let algorithm A be an algorithm which takes an input set S (where $|S|$ is polynomial in ϵ , δ and L_D) of samples generated i.i.d. from distribution D , and with probability at least $1 - \delta$ returns a distribution D_{L_1} , such that $L_1(D, D_{L_1}) \leq \epsilon$.

Let $\xi = \frac{1}{12}(\epsilon^2/L_D)$. We define algorithm A' such that with probability at least $1 - \delta$, A' returns distribution D'_{L_1} , where $L_1(D, D'_{L_1}) \leq \xi$. Algorithm A' runs A with sample S' , where $|S'|$ is polynomial in ξ , δ and L_D (and it should be noted that $|S'|$ is polynomial in ϵ , δ and L_D).

We define D_{L_D} to be the unweighted mixture of all distributions in \mathcal{D} represented by length L_D bit strings, $D_{L_D}(x) = 2^{-L_D} \sum_{D \in \mathcal{D}} D(x)$. We now define distribution D'_{KL} such that $D'_{KL}(x) = (1 - \xi)D'_{L_1}(x) + \xi.D_{L_D}(x)$.

By the definition of D'_{KL} , $L_1(D'_{L_1}, D'_{KL}) \leq 2\xi$. With probability at least $1 - \delta$, $L_1(D, D'_{L_1}) \leq \xi$, and therefore with probability at least $1 - \delta$, $L_1(D, D'_{KL}) \leq 3\xi$.

We define $X_{<} = \{x \in X | D'_{KL}(x) < D(x)\}$. Members of $X_{<}$ contribute positively to $I(D || D'_{KL})$. Therefore

$$\begin{aligned}
I(D||D'_{KL}) &\leq \sum_{x \in X_{<}} D(x) \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right) \\
&= \sum_{x \in X_{<}} (D(x) - D'_{KL}(x)) \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right) \\
&\quad + \sum_{x \in X_{<}} D'_{KL}(x) \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right).
\end{aligned} \tag{7}$$

We have shown that $L_1(D, D'_{KL}) \leq 3\xi$, so $\sum_{x \in X_{<}} (D(x) - D'_{KL}(x)) \leq 3\xi$. Analysing the first term in (7),

$$\sum_{x \in X_{<}} (D(x) - D'_{KL}(x)) \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right) \leq 3\xi \max_{x \in X_{<}} \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right).$$

Note that for all $x \in X$, $D'_{KL}(x) \geq \xi \cdot 2^{-L_D}$. It follows that

$$\max_{x \in X_{<}} \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right) \leq \log(2^{L_D}/\xi) = L_D - \log(\xi).$$

Examining the second term in (7),

$$\sum_{x \in X_{<}} D'_{KL}(x) \left(\frac{\log(D(x))}{\log(D'_{KL}(x))} \right) = \sum_{x \in X_{<}} D'_{KL}(x) \left(\frac{\log(D'_{KL}(x) + h_x)}{\log(D'_{KL}(x))} \right),$$

where $h_x = D(x) - D'_{KL}(x)$, which is a positive quantity for all $x \in X_{<}$. Due to the concavity of the logarithm function, it follows that

$$\begin{aligned}
\sum_{x \in X_{<}} D'_{KL}(x) \left(\frac{\log(D'_{KL}(x) + h_x)}{\log(D'_{KL}(x))} \right) &\leq \sum_{x \in X_{<}} D'_{KL}(x) h_x \left[\frac{d}{dy} (\log(y)) \right]_{y=D'_{KL}(x)} \\
&= \sum_{x \in X_{<}} h_x \leq 3\xi.
\end{aligned}$$

Therefore, $I(D||D'_{KL}) \leq 3\xi(1 + L_D - \log(\xi))$. For values of $\xi \leq \frac{1}{12} (\epsilon^2/L_D)$, it can be seen that $I(D||D'_{KL}) \leq \epsilon$. \square

Corollary 3. *Consider the problem of learning PDFAs having n states, over alphabet Σ , and probabilities represented by bit strings of length ℓ . Using sample size (but not time) polynomial in n , $|\Sigma|$ and ℓ (and the PAC parameters ϵ and δ), a distribution in this class can be estimated within KL distance ϵ .*

The proof follows from the observation that such a PdFA can be represented using a bit string whose length is polynomial in the parameters.

Consequently we can learn the same class of PDFAs under the KL-divergence that can be learned under the L_1 distance in [15], i.e. PDFAs with distinguishable states but no restriction on the expected length of their outputs. However, note that the hypothesis is ‘‘inefficient’’ (a mixture of exponentially many PDFAs).

3 Conclusion

We have shown a close relationship between the error of an estimated input distribution (as measured by L_1 distance or KL-divergence) and the error rate of the resulting classifier. In situations where we believe that input distributions may be accurately estimated, the resulting information about the data may be more useful than just a near-optimal classifier.

A general issue of interest is the question of when one can obtain good classifier from estimated distributions that satisfy weaker goodness-of-approximation criteria than those considered here. Suppose for example that elements of a 2-element domain $\{x_1, x_2\}$ are being labeled by the stochastic rule that assigns labels 0 and 1 to either element of the domain, with equal probability. Then any classifier does no better than random labeling, and so we can use arbitrary distributions D'_0 and D'_1 as estimates of the distributions D_0 and D_1 over examples with label 0 and 1 respectively. In [9] we show that in the basic PAC framework we can sometimes design discriminant functions based on unlabeled data sets, that result in PAC classifiers without any guarantee on how well-estimated is the input distribution. Further work should possibly compromise between the distribution-free setting, and the objective—considered here—of approximating the input distributions in a strong sense.

4 Acknowledgements

We would like to thank Luc Devroye for drawing to our attention the statement of the version of Theorem 1 that appears in [6].

References

1. N. Abe, J. Takeuchi and M. Warmuth. Polynomial learnability of stochastic rules with respect to the KL-divergence and quadratic distance. *IEICE Trans. Inf. and Syst.*, Vol E84-D[3] pp. 299-315 (2001).
2. N. Abe, and M. Warmuth. On the Computational Complexity of Approximating Distributions by Probabilistic Automata. *Machine Learning*, 9, pp. 205-260 (1992).
3. G. Bejerano and G. Yona. Variations on probabilistic suffix trees: statistical modeling and prediction of protein families. *Bioinformatics*, Vol. 17, No. 1, pp. 23-43 (2001).
4. A. Clark and F. Thollard. PAC-learnability of probabilistic deterministic finite state automata. *Journal of Machine Learning Research*, [5] pp. 473-497 (2004).
5. M. Cryan and L. A. Goldberg and P. W. Goldberg. Evolutionary trees can be learnt in polynomial time in the two-state general Markov model. *SIAM Journal on Computing*, 31[2] pp. 375-397 (2001).
6. L. Devroye, L. Györfi and G. Lugosi. *A Probabilistic Theory of Pattern Recognition* Springer (1996).
7. R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons (1973).

8. J. Feldman and R. O'Donnell and R. Servedio. Learning Mixtures of Product Distributions over Discrete Domains. 46th Symposium on Foundations of Computer Science (FOCS), pp. 501-510, (2005).
9. P.W. Goldberg. Some Discriminant-based PAC Algorithms. *Journal of Machine Learning Research*, Vol. 7 (2006), pp. 283-306.
10. K. Hoffgen. Learning and robust learning of product distributions. In *ACM COLT*, pp. 77-83 (1993).
11. M. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. E. Schapire and L. Sellie. On the learnability of discrete distributions. In *Proceedings of STOC*, pp. 273-282 (1994).
12. M. Kearns and R. E. Schapire. Efficient distribution-free learning of probabilistic concepts. *Journal of Computer and System Sciences*, 48[3] pp. 464-497 (1993).
13. M. Kearns, R. E. Schapire and L. M. Sellie. Toward efficient agnostic learning. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, ACM Press, pp. 341-352 (1992).
14. E. Mossel and S. Roch. Learning Nonsingular Phylogenies and Hidden Markov Models. Proceedings of the thirty-seventh annual ACM Symposium on Theory of computing, Baltimore (STOC05), MD, USA, 366-376.
15. N. Palmer and P. W. Goldberg. PAC-learnability of probabilistic deterministic finite state automata in terms of variation distance In *Proceedings of ALT 05*, LNAI 3734, pp. 157-170 (2005).
16. L. G. Valiant. A theory of the learnable. *Journal of the ACM* [27] pp. 1134-1142 (1984).
17. V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, New York, second edition (2000).
18. D. Ron, Y. Singer and N. Tishby. On the Learnability and Usage of Acyclic Probabilistic Finite Automata. *Journal of Computer and System Sciences*, 56(2), pp. 133-152 (1998).