Computable Shell Decomposition Bounds

John Langford
School of Computer Science
Carnegie Mellon University
jcl@cs.anyu.edu
www.cs.cmu.edu/~jcl

David McAllester
AT&T Labs-Research
dmac@research.att.com
www.research.att.com/~dmac

Abstract

Haussler, Kearns, Seung and Tishby introduced the notion of a shell decomposition of the union bound as a means of understanding certain empirical phenomena in learning curves such as phase transitions. Here we use a variant of their ideas to derive an upper bound on the generalization error of a hypothesis computable from its training error and the histogram of training errors for the hypotheses in the class. In most cases this new bound is significantly tighter than traditional bounds computed from the training error and the cardinality or VC dimension of the class. Our results can also be viewed as providing PAC theoretical foundations for a model selection algorithm proposed by Scheffer and Joachims.

1 Introduction

For an arbitrary finite hypothesis class we consider the hypothesis of minimal training error. We give a new upper bound on the generalization error of this hypothesis computable from the training error of the hypothesis and the histogram of the training errors of the other hypotheses in the class. This new bound is typically much tighter than more traditional upper bounds computed from the training error and cardinality or VC dimension of the class.

As a simple example, suppose that we observe that all but one empirical error in a hypothesis space is $1/2$ and one empirical error is 0. Furthermore, suppose that the sample size is large enough (relative to the size of the hypothesis class) that with high confidence we have that, for all hypotheses in the class, the true (generalization) error of a hypothesis is within $1/5$ of its training error. This implies, that with high confidence, hypotheses with training error near $1/2$ have true error in $[3/10, 7/10]$. Intuitively, we would expect that the true error of the hypothesis with minimum empirical error to be very near to 0 rather than simply less than $1/5$ because none of the hypotheses which produced an empirical error of $1/2$ could have a true error close enough to 0 that there exists a significant probability of producing 0 empirical error. The bound presented here validates this intuition. We show that you can ignore hypotheses with training error near $1/2$ in calculating an "effective size" of the class for hypotheses with training error near 0. This new effective class size allows us to calculate a tighter bound on the difference between training error and true error for hypotheses with training error near 0. The new bound is proved using a distribution-dependent application of the union bound similar in spirit to the shell decomposition introduced by Haussler, Kearns, Seung and Tishby [1].

We actually give two upper bounds on generalization error — an uncomputable bound and a computable bound. The uncomputable bound is a function of the unknown distribution of true error rates of the hypotheses in the class. The computable bound is, essentially, the uncomputable bound with the unknown distribution of true errors replaced by the known histogram of training errors. Our main contribution is that this replacement is sound, i.e., the computable version remains, with high confidence, an upper bound on generalization error.

When considering asymptotic properties of learning theory bounds it is important to take limits in which the cardinality (or VC dimension) of the hypothesis class is allowed to grow with the size of the sample. In practice more data typically justifies a larger hypothesis class. For example, the size of a decision tree is generally proportional the amount of training data available. Here we analyze the asymptotic properties of our bounds by considering an infinite sequence of hypothesis classes $\mathcal{H}_m$, one for each sample size m, such that $\frac{\ln |\mathcal{H}_m|}{m}$ approaches a limit larger than zero. This kind of asymptotic analysis provides a clear account of the improvement achieved by bounds that are functions of error rate distributions rather than simply the size (or VC dimension) of the class.

We give a lower bound on generalization error show-

ing that the uncomputable upper bound is asymptotically as tight as possible — any upper bound on generalization error given as a function of the unknown distribution of true error rates must asymptotically be greater than or equal to our uncomputable upper bound. Our lower bound on generalization error also shows that there is essentially no loss in working with an upper bound computed from the true error distribution rather than expectations computed from this distribution as used by Scheffer and Joachims [4].

Asymptotically, the computable bound is simply the uncomputable bound with the unknown distribution of true errors replaced with observed histogram of training errors. Unfortunately, we can show that in limits where $\frac{\ln|\mathcal{H}_m|}{m}$ converges to a value greater than zero, the histogram of training errors need not converge to the distribution of true errors — the histogram of training errors is a "smeared out" version of the distribution of true errors. This smearing loosens the bound even in the large-sample asymptotic limit. We give a precise asymptotic characterization of this smearing effect for the case where distinct hypotheses have independent training errors. In spite of the divergence between the uncomputable and computable bounds, the computable bound is still significantly tighter than classical bounds not involving error distributions.

The computable bound can be used for model selection. In the case of model selection we can assume an infinite sequence of finite model classes $\mathcal{H}_0, \mathcal{H}_1, ...$ where each $\mathcal{H}_j$ is a finite class with $\ln|\mathcal{H}_j|$ growing linearly in j. To perform model selection we find the hypothesis of minimal training error in each class and use the computable bound to bound its generalization error. We can then select, among these, the model with the smallest upper bound on generalization error. Scheffer and Joachims propose (without formal justification) replacing the distribution of true errors with the histogram of training errors. Under this replacement, the model selection algorithm based on our computable upper bound is asymptotically identical to the algorithm proposed by Scheffer and Joachims.

The shell decomposition is a distribution-dependent use of the union bound. Distribution-dependent uses of the union bound have been previously exploited in so-called self-bounding algorithms. Freund [5] defines, for a given learning algorithm *and data distribution*, a set S of hypotheses such that with high probability over the sample, the algorithm will always return a hypothesis in that set. Although S is defined in terms of the unknown data distribution, Freund gives a way of computing a set S' from the given algorithm and the sample such that, with high confidence, S' contains S and hence the "effective size" of the hypothesis class is bounded by $|S'|$. Langford and Blum [7] give a more practical version of this algorithm. Given an algorithm and data distribution they conceptually define a weighting over the possible executions of the algorithm. Although the data distribution is unknown, they give a way of computing a lower bound on the weight of the particular execution of the algorithm generated by the sample at hand. In this paper we consider distribution dependent union bounds defined independently of any particular learning algorithm.

2 Mathematical Preliminaries

For an arbitrary measure on an arbitrary sample space we use the notation $\forall^\delta S\ \Phi[S,\delta]$ to mean that with probability at least $1-\delta$ over the choice of the sample S we have that $\Phi[S,\delta]$ holds. In practice S is the training sample of a learning algorithm. Note that $\forall x\ \forall^\delta S\ \Phi[x,\ S,\ \delta]$ does not imply $\forall^\delta S\ \forall x\ \Phi[x,\ S,\ \delta]$. If X is a finite set, and for all $x \in X$ we have the assertion $\forall \delta > 0\ \forall^\delta S\ \Phi[S,x,\delta]$ then by a standard application of the union bound we have the assertion $\forall \delta > 0\ \forall^\delta S\ \forall x \in X\ \Phi[S,x,\frac{\delta}{|X|}]$. We will call this the quantification rule. If $\forall \delta > 0\ \forall^\delta S\ \Phi[S,\delta]$ and $\forall \delta > 0\ \forall^\delta S\ \Psi[S,\ \delta]$ then by a standard application of the union bound we have $\forall \delta > 0\ \forall^\delta S\ \Phi[S,\frac{\delta}{2}] \wedge \Psi[S,\frac{\delta}{2}]$. We will call this the conjunction rule.

The KL-divergence of p from q, denoted $D(q||p)$, is $q\ln(\frac{q}{p})+(1-q)\ln(\frac{1-q}{1-p})$ with $0\ln(\frac{0}{p})=0$ and $q\ln(\frac{q}{0})=\infty$. Let $\hat{p}$ be the fraction of heads in a sequence S of m tosses of a biased coin where the probability of heads is p. For $\hat{p} \geq p$ we have the following inequality given by Chernoff in 1952 [3].

$$\forall q \in [p,1]: \quad Pr(\hat{p} \geq q) \quad \leq \quad e^{-mD(q||p)} \qquad (1)$$

This bound can be rewritten as follows.

$$\forall \delta > 0\ \forall^\delta S \quad D(\max(\hat{p},p)||p) \leq \frac{\ln(\frac{1}{\delta})}{m} \qquad (2)$$

To derive (2) from (1) note that $Pr(D(\max(\hat{p},p)||p) \geq \frac{\ln(\frac{1}{\delta})}{m})$ equals $Pr(\hat{p} \geq q)$ where $q \geq p$ and $D(q||p) = \frac{\ln(\frac{1}{\delta})}{m}$. By (1) we then have that this probability is no larger than $e^{-mD(q||p)} = \delta$. It is just as easy to derive (1) from (2) so the two statements are equivalent. By duality, i.e., by considering the problem defined by replacing p by $1-p$, we get the following.

$$\forall \delta > 0\ \forall^\delta S \quad D(\min(\hat{p},p)||p) \leq \frac{\ln(\frac{1}{\delta})}{m} \qquad (3)$$

Conjoining (2) and (3) yields the following corollary of (1).

$$\forall \delta > 0\ \forall^\delta S \quad D(\hat{p}||p) \leq \frac{\ln(\frac{2}{\delta})}{m} \qquad (4)$$

Using the inequality $D(q||p) \geq 2(q-p)^2$ one can show that (4) implies the following better known form of the

Chernoff bound.

$$\forall \delta > 0 \; \forall^{\delta} S \quad |p - \hat{p}| \leq \sqrt{\frac{\ln(\frac{2}{\delta})}{2m}} \tag{5}$$

Using the inequality $D(q||p) \geq \frac{(p-q)^2}{2q}$, which holds for $q \leq p$, we can show that (3) implies the following.[1]

$$\forall \delta > 0 \; \forall^{\delta} S \quad p \leq \hat{p} + \sqrt{\frac{2\hat{p}\ln(\frac{1}{\delta})}{m} + \frac{2\ln(\frac{1}{\delta})}{m}} \tag{6}$$

Note that for small values of $\hat{p}$ formula (6) gives a tighter upper bound on p than does (5). The upper bound on p implicit in (4) is somewhat tighter than the minimum of the bounds given by (5) and (6).

We now consider a formal setting for hypothesis learning. We assume a finite set $\mathcal{H}$ of hypotheses and a space $\mathcal{X}$ of instances. We assume that each hypothesis represents a function from $\mathcal{X}$ to $\{0, 1\}$ where we write $h(x)$ for the value of the function represented by hypothesis h when applied to instance x. We also assume a distribution D on pairs $\langle x,\ y\rangle$ with $x \in \mathcal{X}$ and $y \in \{0, 1\}$. For any hypothesis h we define the error rate of h, denoted $e(h)$, to be $P_{\langle x,\ y\rangle \sim D}(h(x) \neq y)$. For a given sample S of m pairs drawn from D we write $\hat{e}(h)$ to denote the fraction of the pairs $\langle x,\ y\rangle$ in S such that $h(x) \neq y$. Quantifying over $h \in \mathcal{H}$ in (4) yields the following second corollary of (1).

$$\forall^{\delta} S \quad \forall h \in \mathcal{H} \quad D(\hat{e}(h)||e(h)) \leq \frac{\ln|\mathcal{H}| + \ln(\frac{2}{\delta})}{m} \tag{7}$$

By consider bounds on $D(q||p)$ we can derive the following more well known corollary of (7).

$$\forall^{\delta} S \quad \forall h \in \mathcal{H} \quad |e(h) - \hat{e}(h)| \leq \sqrt{\frac{\ln|\mathcal{H}| + \ln(\frac{2}{\delta})}{2m}} \tag{8}$$

These two formulas both limit the distance between $\hat{e}(h)$ and $e(h)$. In this paper we will work with (7) rather than (8) because it yields an (uncomputable) upper bound on generalization error that is optimal up to asymptotic equality.

3 The Upper Bound

Our goal now is to improve on (7). Our first step is to divide the hypotheses in $\mathcal{H}$ into m disjoint sets based on their true error rates. More specifically, for $p \in [0, 1]$ define $\lceil\lceil p\rceil\rceil$ to be $\frac{\max(1, \lceil mp\rceil)}{m}$. Note that $\lceil\lceil p\rceil\rceil$ is of the form $\frac{k}{m}$ where either $p = 0$ and $k = 1$ or $p > 0$ and $p \in (\frac{k-1}{m},\ \frac{k}{m}]$. In either case we have $\lceil\lceil p\rceil\rceil \in \{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}$ and if $\lceil\lceil p\rceil\rceil = \frac{k}{m}$ then $p \in [\frac{k-1}{m},\ \frac{k}{m}]$. Now we define $\mathcal{H}(\frac{k}{m})$ to be the set of $h \in \mathcal{H}$ such that $\lceil\lceil e(h)\rceil\rceil = \frac{k}{m}$. We define $s(\frac{k}{m})$ to be $\ln(\max(1, |\mathcal{H}(\frac{k}{m})|))$. We now have the following lemma.

Lemma 3.1 $\forall \delta > 0 \; \forall^{\delta} S \quad \forall h \in \mathcal{H}$

$$D(\hat{e}(h)||e(h)) \leq \frac{s(\lceil\lceil e(h)\rceil\rceil) + \ln(\frac{2m}{\delta})}{m}$$

Proof: Quantifying over $p \in \{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}$ and $h \in \mathcal{H}(p)$ in (4) gives $\forall \delta > 0$, $\forall^{\delta} S$, $\forall p \in \{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}$, $\forall h \in \mathcal{H}(p)$,

$$D(\hat{e}(h)||e(h)) \leq \frac{\ln|\mathcal{H}(p)| + \ln(\frac{2m}{\delta})}{m}$$

But this implies the lemma. □

Lemma 3.1 imposes a constraint, and hence a bound, on $e(h)$. More specifically, we have the following.

$$e(h) \leq \sup_{q : D(\hat{e}(h)||q) \leq \frac{s(\lceil\lceil q\rceil\rceil) + \ln(\frac{2m}{\delta})}{m}} q \tag{9}$$

This is our uncomputable bound. It is uncomputable because the m numbers $s(\frac{1}{m}), \ldots, s(\frac{m}{m})$ are unknown. Ignoring this problem, however, we can see that this bound is typically significantly tighter than (7). More specifically, we can rewrite (7) as follows.

$$e(h) \leq \sup_{q : D(\hat{e}(h)||q) \leq \frac{\ln|\mathcal{H}| + \ln(\frac{2}{\delta})}{m}} q \tag{10}$$

Since $s(\frac{k}{m}) \leq \ln|\mathcal{H}|$, and since $\frac{\ln m}{m}$ is small for large m, we have that (9) is never significantly looser than (10). Now consider a hypothesis h such that the bound on $e(h)$ given by (7), or equivalently, (10), is significantly less than 1/2. Assuming m is large, the bound given by (9) must also be significantly less than 1/2. But for q significantly less than 1/2 we will typically have that $s(\lceil\lceil q\rceil\rceil)$ is significantly smaller than $\ln|\mathcal{H}|$. For example, suppose $\mathcal{H}$ is the set of all decision trees of size $m/10$. For large m, a random decision tree of this size will have error rate near 1/2. The set of decision trees with error rate significantly smaller than 1/2 will be an exponentially small faction of the set of all possible trees. So for q small compared to 1/2 we get that $s(\lceil\lceil q\rceil\rceil)$ is significantly smaller than $\ln|H|$. This will make the bound given by (9) significantly tighter than the bound given by (10).

We now show that the distribution of true errors can be replaced, essentially, by the histogram of training errors. We first introduce the following definitions.

$$\hat{\mathcal{H}}\left(\frac{k}{m}, \delta\right) \equiv \left\{ h \in \mathcal{H} : \left|\hat{e}(h) - \frac{k}{m}\right| \leq \frac{1}{m} + \sqrt{\frac{\ln(\frac{16m^2}{\delta})}{2m - 1}} \right\}$$

$$\hat{s}\left(\frac{k}{m},\ \delta\right) \equiv \ln\left(\max\left(1,\ 2\left|\hat{\mathcal{H}}\left(\frac{k}{m},\ \delta\right)\right|\right)\right)$$

The definition of $\hat{s}\left(\frac{k}{m},\ \delta\right)$ is motivated by the following lemma.

[1] A derivation of this formula can be found in [8] or [9]. To see the need for the last term consider the case where $\hat{p} = 0$.

Lemma 3.2 $\forall\delta > 0, \forall^\delta S, \forall q \in \{\frac{1}{m}, \ldots, \frac{m}{m}\}$,

$$s(q) \leq \hat{s}(q,\ 2\delta)$$

Before proving lemma 3.2 we note that by conjoining (9) and lemma 3.2 we get the following. This is our main result.

Theorem 3.3 $\forall\delta > 0, \forall^\delta S, \forall h \in \mathcal{H}$,

$$e(h) \leq \sup_{q: D(\hat{e}(h)||q) \leq \frac{\hat{s}(\lceil\lceil q\rceil\rceil,\ \delta) + \ln(\frac{4m}{\delta})}{m}} q$$

As for lemma 3.1, the bound implicit in theorem 3.3 is typically significantly tighter than the bound in (7) or its equivalent form (10). The argument for the improved tightness of theorem 3.3 over (10) is similar to the argument for (9). More specifically, consider a hypothesis h for which the bound in (10) is significantly less than 1/2. Since $\hat{s}(\lceil\lceil q\rceil\rceil,\ \delta) \leq \ln|\mathcal{H}|$, the set of values of q satisfying the condition in theorem 3.3 must all be significantly less than 1/2. but for large m we have that $\sqrt{\frac{\ln(16m^2)}{2m-1}}$ is small. So if q is significantly less than 1/2 then all hypotheses in $\hat{\mathcal{H}}(\lceil\lceil q\rceil\rceil, \delta)$ have empirical error rates significantly less than 1/2. But for most hypothesis classes, e.g., decision trees, the set of hypotheses with empirical error rates far from 1/2 should be an exponentially small fraction of the class. Hence we get that $\hat{s}(\lceil\lceil q\rceil\rceil,\ \delta)$ is significantly less than $\ln|\mathcal{H}|$ and theorem 3.3 is tighter than (10).

The remainder of this section is a proof of lemma 3.2. Our departure point for the proof is the following lemma from [6].

Lemma 3.4 (McAllester 99) *For any measure on any hypothesis class we have the following where* $\mathrm{E}_h f(h)$ *denotes the expectation of* $f(h)$ *under the given measure on* h.

$$\forall\delta > 0\ \forall^\delta S \quad \mathrm{E}_h e^{(2m-1)(\hat{e}(h)-e(h))^2} \leq \frac{4m}{\delta}$$

Intuitively, this lemma states that with high confidence over the choice of the sample most hypotheses have empirical error near their true error. This will allow us to prove that $\hat{s}(\lceil\lceil q\rceil\rceil,\ \delta)$ bounds $s(\lceil\lceil q\rceil\rceil)$. More specifically, by considering the uniform distribution on $\mathcal{H}(\frac{k}{m})$, lemma 3.4 implies the following.

$$\mathrm{E}_{h\sim\mathcal{H}\left(\frac{k}{m}\right)}\left(e^{(2m-1)(\hat{e}(h)-e(h))^2}\right) \quad \leq \quad \frac{4m}{\delta}$$

$$Pr_{h\sim\mathcal{H}\left(\frac{k}{m}\right)}\left(e^{(2m-1)(\hat{e}(h)-e(h))^2} \geq \frac{8m}{\delta}\right) \quad \leq \quad \frac{1}{2}$$

$$Pr_{h\sim\mathcal{H}\left(\frac{k}{m}\right)}\left(e^{(2m-1)(\hat{e}(h)-e(h))^2} \leq \frac{8m}{\delta}\right) \quad \geq \quad \frac{1}{2}$$

$$\left|\left\{h \in \mathcal{H}(\frac{k}{m}) :\ |\hat{e}(h) - e(h)| \leq \sqrt{\frac{\ln(\frac{8m}{\delta})}{2m-1}}\right\}\right| \quad \geq \quad \frac{1}{2}|\mathcal{H}(\frac{k}{m})|$$

$$\left|\left\{h \in \mathcal{H}(\frac{k}{m}) :\ |\hat{e}(h) - \frac{k}{m}| \leq \frac{1}{m} + \sqrt{\frac{\ln(\frac{8m}{\delta})}{2m-1}}\right\}\right| \quad \geq \quad \frac{1}{2}|\mathcal{H}(\frac{k}{m})|$$

$$\left|\mathcal{H}(\frac{k}{m})\right| \leq 2\left|\hat{\mathcal{H}}\left(\frac{k}{m},\ 2m\delta\right)\right|$$

Lemma 3.2 now follows by quantification over $q \in \{\frac{1}{m}, \ldots, \frac{m}{m}\}$. $\square$

4 Asymptotic Analysis and Phase Transitions

The bounds given in (9) and theorem 3.3 exhibit phase transitions. More specifically, the bounding expression can be discontinuous in δ and m, e.g., arbitrarily small changes in δ can cause large changes in the bound. To see how this happens consider the following constraint on the quantity q.

$$D(\hat{e}(h)||q) \leq \frac{s(\lceil\lceil q\rceil\rceil) + \ln(\frac{2m}{\delta})}{m} \tag{11}$$

The bound given by (9) is the supremum of the values of q satisfying (11). Assume that m is sufficiently large that we can think of $\frac{s(\lceil\lceil q\rceil\rceil)}{m}$ as a continuous function of q which we will write as $\bar{s}(q)$. We can then rewrite (11) as follows where λ is a quantity not depending on q and $\bar{s}(q)$ does not depend on δ.

$$D(\hat{e}(h)||q) \leq \bar{s}(q) + \lambda \tag{12}$$

For $q \geq \hat{e}(h)$ we that $D(\hat{e}(h)||q)$ is a monotonically increasing function of q. It is reasonable to assume that for $q \leq 1/2$ we also have that $\bar{s}(q)$ is a monotonically increasing function of q. But even under these conditions it is possible that the feasible values of q, i.e., those satisfying (12) can be divided into separated regions. Furthermore, increasing λ can cause a new feasible region to come into existence. When this happens the bound, which is the supremum of the feasible values, can increase discontinuously. At a more intuitive level, consider a large number of high error concepts and smaller number of lower error concepts. At a certain confidence level the high error concepts can be ruled out. But as the confidence requirement becomes more stringent suddenly (and discontinuously) the high error concepts must be considered. A similar discontinuity can occur in sample size. Phase transitions in shell decomposition bounds are discussed in more detail by Haussler et al. [1].

Phase transition complicate asymptotic analysis. But asymptotic analysis illuminates the nature of phase transitions. As mentioned in the introduction, in the asymptotic analysis of learning theorem bounds it is important

that one not hold $\mathcal{H}$ fixed as the sample size increases. If we hold $\mathcal{H}$ fixed then $\lim_{m\to\infty} \frac{\ln|\mathcal{H}|}{m} = 0$. But this is not what one expects for large samples in practice. As the sample size increases larger hypothesis classes can be and typically are used. Intuitively, we expect that even for very large m we have that $\frac{\ln|\mathcal{H}|}{m}$ is far from zero. In this section we formulate an asymptotic analysis by considering a sequence of functions s_1, s_2, s_3, ... where s_m is a function defined on the i rational numbers $\frac{1}{m}$, ..., $\frac{m}{m}$. As mentioned above, for large m it is convenient to think of the function of $\frac{s(\lceil\lceil q\rceil\rceil)}{m}$ as a continuous function of q. Let $\bar{s}(q)$ be a continuous function of $q \in [0,1]$. We say that the sequence s_1, s_2, s_3, ..., converges to $\bar{s}(q)$ if the sequence of functions $\frac{s_1(\lceil\lceil q\rceil\rceil)}{1}$, $\frac{s_2(\lceil\lceil q\rceil\rceil)}{2}$, $\frac{s_3(\lceil\lceil q\rceil\rceil)}{3}$, ..., viewed as functions of $q \in [0,1]$, converge uniformly to the function $\bar{s}(q)$. This means that for any $\epsilon > 0$ there exists a k such that for all $m > k$ we have the following.

$$\forall q \in [0,1] \;\; |\frac{s_m(\lceil\lceil q\rceil\rceil)}{m} - \bar{s}(q)| \leq \epsilon$$

Given these function sequences we define the following two functions of an empirical error rate $\hat{e}$ and a sample size m.

$$\begin{aligned} B_m(\hat{e}) &\equiv \sup_{q: D(\hat{e}||q) \leq \frac{s_m(\lceil\lceil q\rceil\rceil)+\ln(\frac{2m}{\delta})}{m}} q \\ B(\hat{e}) &\equiv \sup_{q: D(\hat{e}||q) \leq \bar{s}(q)} q \end{aligned}$$

The function $B_m(\hat{e})$ corresponds directly to the upper bound in (9). The function $B(\hat{e})$ is intended to be the large m asymptotic limit of $B_m(\hat{e})$. However, phase transitions complicate asymptotic analysis. The bound $B(\hat{e})$ need not be a continuous function of $\hat{e}$. A value of $\hat{e}$ where the bound $B(\hat{e})$ is discontinuous corresponds to a phase transition in the bound. At a phase transition the sequence $B_m(\hat{e})$ need not converge. Away from phase transitions, however, we have the following theorem.

Theorem 4.1 *If the bound $B(\hat{e})$ is continuous at the point $\hat{e}$ (so we are not at a phase transition), and s_1, s_2, s_3, ... converges to $\bar{s}$, then we have the following.*

$$\lim_{m\to\infty} B_m(\hat{e}) = B(\hat{e})$$

Proof: Define the set $F_m(\hat{e})$ as follows.

$$F_m(\hat{e}) \equiv \{q \in [0,1] : \; D(\hat{e}||q) \leq \frac{s_m(\lceil\lceil q\rceil\rceil) + \ln(\frac{2m}{\delta})}{m}\}$$

This gives the following.

$$B_m(\hat{e}) = \sup F_m(\hat{e})$$

Similarly, define $F(\hat{e},\ \epsilon)$ and $B(\hat{e},\ \epsilon)$ as follows.

$$\begin{aligned} F(\hat{e},\ \epsilon) &\equiv \{q \in [0,1] : \; D(\hat{e}||q) \leq \bar{s}(q) + \epsilon\} \\ B(\hat{e},\ \epsilon) &\equiv \sup F(\hat{e},\ \epsilon) \end{aligned}$$

The functions of the form $\frac{s_i(\lceil\lceil q\rceil\rceil)+\ln\frac{2m}{\delta}}{m}$ converge uniformly to the function $\bar{s}(q)$. This implies that for any $\epsilon > 0$ there exists a k such that for all $m > k$ we have the following.

$$F(\hat{e},\ -\epsilon) \subseteq F_m(\hat{e}) \subseteq F(\hat{e},\ \epsilon)$$

But this in turn implies the following.

$$B(\hat{e},\ -\epsilon) \leq B_m(\hat{e}) \leq B(\hat{e},;\epsilon) \tag{13}$$

But it is possible to show that if $B(\hat{e})$ is continuous at $\hat{e}$ then $B(\hat{e},\ \epsilon)$ is continuous at the point $\langle\hat{e},\ 0\rangle$. So we get the following.

$$\lim_{m\to\infty} B_m(\hat{e}) = B(\hat{e},\ 0) = B(\hat{e})$$

□

Theorem 4.1 can be interpreted as saying that for large sample sizes, and for values of $\hat{e}$ other than the special phase transition values, the bound has a well defined value independent of the confidence parameter δ and determined only by a smooth function $\bar{s}(q)$.

A similar theorem can be given for the bound in theorem 3.3. Again we prove a theorem stating that for large m, and at points other than phase transitions, the bound is independent of δ and is determined by a smooth limit curve. Let $U_m(\frac{k}{m})$ be the set of hypotheses with empirical error $\frac{k}{m}$. We assume a sequence of functions u_1, u_2, u_3, ... where u_m is a function on the m rational numbers $\frac{1}{m}$, ..., $\frac{m}{m}$ and $u_m(\frac{k}{m}) = \ln|U_m(\frac{k}{m})|$. Let $\bar{u}(q)$ be a continuous function on $q \in [0,1]$. We say that u_1, u_2, u_3 converges to $\bar{u}$ if the sequence of functions $\frac{u_1(\lceil\lceil q\rceil\rceil)}{1}$, $\frac{u_2(\lceil\lceil q\rceil\rceil)}{2}$, $\frac{u_3(\lceil\lceil q\rceil\rceil)}{3}$ ... converges uniformly to the function $\bar{u}(q)$. Unlike the sets $\hat{\mathcal{H}}(\frac{k}{m},\ \delta)$, the sets $U_m(\frac{k}{m})$ are defined independent of δ. Given the sequence u_1, u_2, u_3, ... converging to $\bar{u}$ we can define $\hat{s}_i(\frac{k}{m},\ \delta)$ as follows.

$$\begin{aligned} \hat{s}_m(\frac{k}{m},\ \delta) &\equiv \ln(\max(1,\ N_m(\frac{k}{m},\ \delta))) \\ N_m(\frac{k}{m},\ \delta) &\equiv \sum_{\frac{k'}{m} \in I_m(\frac{k}{m},\ \delta)} |U_m(\frac{k}{m})| \\ I_m(q,\ \delta) &\equiv \left\{ q' : \; |q' - q| \leq \frac{2}{m} + \sqrt{\frac{\ln(\frac{8m}{\delta})}{2m-1}} \right\} \end{aligned}$$

The basic idea is that $u_m(\frac{k}{m})$ is based on the number of hypotheses with an empirical error rate of exactly $\frac{k}{m}$ while $\hat{s}_m(\frac{k}{m},\ \delta)$ corresponds to the quantity actually used in theorem 3.3, i.e., the number of concepts with empirical error rate within a certain distance of $\frac{k}{m}$ where the distance depends on the confidence parameter δ. Our first observation is that, asymptotically, the numbers $\hat{s}_m(\frac{k}{m},\ \delta)$ are the same as the numbers $u_m(\frac{k}{m})$ — asymptotically the dependence on δ vanishes.

Theorem 4.2 *If the functions $\frac{u_m(\lceil\lceil q\rceil\rceil)}{m}$ converge uniformly to the function $\bar{u}(q)$ as m increases, and $\bar{u}(q)$ is differentiable with bounded derivative, then, for any fixed value of δ, the functions $\frac{\hat{s}_m(\lceil\lceil q\rceil\rceil,\ \delta)}{m}$ also converge uniformly to $\bar{u}(q)$.*

Proof: Since the functions $\frac{u_i(\lceil\lceil q\rceil\rceil)}{m}$ converge uniformly to $\bar{u}(q)$, the width of the interval $I(q,\ \delta)$ goes to zero, and $\bar{u}(q)$ has bounded derivative, for any $\epsilon > 0$ there exists a k such that for $m > k$ we have the following.

$$\forall q \in [0,1]\ e^{m(\bar{u}(q)-\epsilon)} \leq N_m(\lceil\lceil q\rceil\rceil,\ \delta) \leq m e^{m(\bar{u}(q)+\epsilon)}$$

Taking the log of all these quantities and dividing by m gives the following.

$$\forall q \in [0,1]\ \frac{\bar{u}(q) - \epsilon \leq \hat{s}_m(\lceil\lceil q\rceil\rceil,\ \delta)}{m} \leq \bar{u}(q) + \epsilon + \frac{\ln m}{m}$$

Since ϵ can be arbitrarily close to zero, the result follows. □

Given quantities $\frac{\hat{s}_i(\lceil\lceil q\rceil\rceil,\ \delta)}{m}$ that converge uniformly to $\bar{u}(q)$ the remainder of the analysis is identical to that for the asymptotic analysis of (9). We define the following upper bounds.

$$\hat{B}_m(\hat{e}) \equiv \sup_{q: D(\hat{e}||q) \leq \frac{\hat{s}_m(\lceil\lceil q\rceil\rceil,\ \delta) + \ln\left(\frac{4m}{\delta}\right)}{m}} q$$

$$\hat{B}(\hat{e}) \equiv \sup_{q: D(\hat{e}||q) \leq \bar{u}(q)} q$$

Again we say that $\hat{e}$ is at a phase transition if the function $\hat{B}(\hat{e})$ is discontinuous at the value $\hat{e}$. We then get the following whose proof is identical to that of theorem 4.1.

Theorem 4.3 *If the bound $\hat{B}(\hat{e})$ is continuous at the point $\hat{e}$ (so we are not at a phase transition), the sequence u_1, u_2, u_3, ... converges to $\bar{u}$, and $\bar{u}(q)$ has bounded derivative, then we have the following.*

$$\lim_{m\to\infty} \hat{B}_m(\hat{e}) = \hat{B}(\hat{e})$$

5 Asymptotic Optimality of (9)

Let h^* be a hypothesis in $\mathcal{H}$ of minimal empirical error, i.e., an $h \in \mathcal{H}$ minimizing $\hat{e}(h)$. Formula (9) can be viewed as providing an upper bound on $e(h^*)$ as a function of $\hat{e}(h^*)$ and the entropy curve s. In this section we show that for any entropy curve s and value of $\hat{e}(h^*)$ consistent with s, there exists a hypothesis class and data distribution such that the upper upper bound in (9) is realized up to asymptotic equality. Hence, up to asymptotic equality, (9) is the tightest possible bound computable from $\hat{e}(h^*)$ and the m numbers $s(\frac{1}{m})$, ..., $s(\frac{m}{m})$.

Note that classical VC dimensions bounds are optimal up to a constant factor over bounds computable from $\hat{e}(h^*)$ and the class $\mathcal{H}$. The m numbers $s(\frac{1}{m})$, ..., $s(\frac{m}{m})$ depend on both $\mathcal{H}$ and the data distribution. Hence the bound in (9) uses more information and hence can be tighter than classical VC bounds. A similar statement applies to the bound in theorem (3.3) computed from the empirically observable numbers $\hat{s}(\frac{1}{m})$, ..., $\hat{s}(\frac{m}{m})$. In this case, the bound uses more information from the sample than just $\hat{e}(h^*)$. The optimality theorem given here also differs from the traditional lower bound results for VC dimension in that here the lower bounds match the upper bounds up to asymptotic equality rather than a constant factor.

The departure point for our optimality analysis is the following lemma from [2].

Lemma 5.1 (Cover and Thomas) *If $\hat{p}$ is the fraction of heads out of m tosses of a coin where the true probability of heads is p then for $q \geq p$ we have the following.*

$$Pr(\hat{p} \geq q) \geq \frac{1}{m+1} e^{-mD(q||p)}$$

This lower bound on $Pr(\hat{p} \geq q)$ is very close to Chernoff's 1952 upper bound (1). The tightness of (9) is a direct reflection of the tightness (1). To exploit Lemma 5.1 we need to construct hypothesis classes and data distributions where distinct hypotheses have independent training errors. More specifically, we say that a set of hypotheses $\{h_1,\ \ldots,\ h_n\}$ has independent training errors if the random variables $\hat{e}(h_1)$, ..., $\hat{e}(h_n)$ are independent.

We now define $h^*(\frac{k}{m})$ to be the hypothesis of minimal training error in the set $\mathcal{H}(\frac{k}{m})$. For $q \in [0,1]$ we have that lemma 3.1 implies the following bound on $\hat{e}(h^*(\lceil\lceil q\rceil\rceil))$.

$$\hat{e}(h^*(\lceil\lceil q\rceil\rceil)) \geq \inf_{\hat{e}: D\left(\hat{e}||\lceil\lceil q\rceil\rceil - \frac{1}{m}\right) \leq \frac{s(\lceil\lceil q\rceil\rceil) + \ln\left(\frac{2m}{\delta}\right)}{m}} \hat{e}$$

We now have the following upper bound on $\hat{e}(h^*(\lceil\lceil q\rceil\rceil))$ in the case where the hypotheses in $\mathcal{H}(\lceil\lceil q\rceil\rceil)$ are independent.

Lemma 5.2 *If the hypotheses in the class $\mathcal{H}(\lceil\lceil q\rceil\rceil)$ are independent then $\forall\delta > 0$, $\forall^\delta S$, $\forall q \in \{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}$,*

$$\hat{e}(h^*(q)) \leq \inf_{\hat{e}: D(\min(\hat{e},\ q - \frac{1}{m})||q) \leq \frac{s(q) - \ln(m+1) - \ln(\ln(\frac{m}{\delta}))}{m}} \hat{e}$$

The proof of lemma 5.2 is given later in this section. Note that the upper and lower bounds on $\hat{e}(h^*(\lceil\lceil q\rceil\rceil))$ agree up to asymptotic equality — in a large m limit where $\frac{s_m(\lceil\lceil q\rceil\rceil)}{m}$ uniformly approaches a continuous function $\bar{s}(q)$ we have that the upper and lower bound on $\hat{e}(h^*(\lceil\lceil q\rceil\rceil))$ both converge to the following.

$$\hat{e}(h^*(q)) = \inf_{\hat{e}: D(\hat{e}||q) \leq \bar{s}(q)} \hat{e}$$

This asymptotic value of $\hat{e}(h^*(q))$ is a continuous function of q and phase transitions are not an issue. Note that for large m and independent hypotheses we get that $\hat{e}(h^*(q))$ is determined as a function of the true error rate q and $\frac{s(\lceil\lceil q\rceil\rceil)}{m}$.

Lemma 5.2 implies that no uniform bound on $e(h)$ as a function of $\hat{e}(h)$ and $|\mathcal{H}(\frac{1}{m})|$, ..., $|\mathcal{H}(\frac{m}{m})|$ can be asymptotically tighter than (9). To see this consider an arbitrary sequence $|\mathcal{H}(\frac{1}{m})|, \ldots, \mathcal{H}|(\frac{m}{m})|$ of non-negative real numbers. It suffices to show that any such sequence is compatible with a hypothesis class and data distribution such that distinct concepts are independent. Such a class and distribution can be constructed as follows. Consider an instance space of bit strings each of which has N bits and a hypothesis space of N hypotheses h_1, ..., h_N where $h_i(x) = 1$ if bit i in x is 1. We associate each i with a target error rate q_i so that $|\mathcal{H}(\frac{k}{m})|$ concepts have target error rate $\frac{k}{m}$. To ensure that concept h_i has error rate q_i we generate data pairs $\langle x,\ y\rangle$ by first generating y with 0 and 1 having equal probability and generate each bit of x independently so that the probability that the ith bit agrees with y is q_i.

We now turn to the proof of lemma 5.2. By an argument similar to the derivation of (3) from (1) we can prove the following.

$$Pr\left(D(\min(\hat{p},p)||p) \geq \frac{\ln(\frac{1}{\delta}) - \ln(m+1)}{m}\right) \geq \delta \quad (14)$$

We now prove the following rule for existential quantification.

Lemma 5.3 *Existential Quantification Rule: Let X be any finite set. If, for every $x \in X$ and $\delta > 0$ we have $Pr(\Theta[S,x,\delta]) \geq \delta$, and $Pr(\forall x \in X\ \Theta[S,\ x,\ \delta]) = \prod_{x\in X} Pr(\Theta[S,\ x,\ \delta])$, then we have $\forall \delta > 0\ \forall^{\delta} S\ \ \exists x \in X\ \ \Theta[S,x,\frac{\ln(\frac{1}{\delta})}{|X|}]$.*

Proof:

$$\begin{array}{rcl} Pr(\Theta[S,x,\frac{\ln(\frac{1}{\delta})}{|X|}]) & \geq & \frac{\ln(\frac{1}{\delta})}{|X|} \\ Pr(\neg\Theta[S,x,\frac{\ln(\frac{1}{\delta})}{|X|}]) & \leq & 1 - \frac{\ln(\frac{1}{\delta})}{|X|} \\ & \leq & e^{-\frac{\ln(\frac{1}{\delta})}{|X|}} \\ Pr(\forall x \in X \neg\Theta[S,\ x,\frac{\ln(\frac{1}{\delta})}{|X|}]) & \leq & e^{-\ln(\frac{1}{\delta})} = \delta \end{array}$$

□

To prove lemma 5.2 let q be a fixed rational number of the form $\frac{k}{m}$. Assuming independent hypotheses we can applying lemma 5.3 to (14) to get $\forall \delta > 0$, $\forall^{\delta} S$, $\exists h \in \mathcal{H}(\frac{k}{m})$,

$$D(\min(\hat{e}(h),e(h))||e(h)) \geq \frac{s(q) - \ln(m+1) - \ln(\ln(\frac{1}{\delta}))}{m}$$

Let w be the hypothesis in $\mathcal{H}(q)$ satisfying this formula. We now have $\hat{e}(h^*(q)) \leq \hat{e}(w)$ and $q - \frac{1}{m} \leq e(w) \leq q$. These two conditions imply $\forall \delta > 0$, $\forall^{\delta} S$,

$$D(\min(\hat{e}(h^*(q)), q-\frac{1}{m})||q) \geq \frac{s(q) - \ln(m+1) - \ln(\ln\frac{1}{\delta})}{m}$$

This implies the following.

$$\hat{e}(h^*(q)) \leq \inf_{\hat{e}:D(\min(\hat{e},\ q-\frac{1}{m})||q) \leq \frac{s(q)-\ln(m+1)-\ln(\ln(\frac{1}{\delta}))}{m}} \hat{e}$$

Lemma 5.2 now follows by quantification over $q \in \{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}$.

6 Relating $\hat{s}$ and s

In this section we show that in large m limits of the type discussed in section 4 the histogram of empirical errors need not converge to the histogram of true errors. So even in the large m asymptotic limit, the bound given by theorem 3.3 is significantly weaker than the bound given by (9).

To show that $\hat{s}(\lceil\lceil q\rceil\rceil,\ \delta)$ can be asymptotically different from $s(\lceil\lceil q\rceil\rceil)$ we consider the case of independent hypotheses. More specifically, we consider an infinite sequence of hypothesis spaces $\mathcal{H}_1$, $\mathcal{H}_2$, $\mathcal{H}_3$, ... where $\mathcal{H}_i$ is a finite set of hypotheses. Intuitively, $\mathcal{H}_i$ is the hypothesis class to be used on a sample of size i. As the sample size grows so does the size of the hypothesis class. We assume a single fixed data distribution D and assume that each class $\mathcal{H}_m$ has independent hypotheses under data distribution D. For any particular value of m we define $s_m(\frac{k}{m})$ in a manner analogous to $s(\frac{k}{m})$ but for hypothesis class $\mathcal{H}_m$. To ensure a well defined limiting behavior we require that there exists a continuous differentiable non-negative function $\bar{s}(q)$ on $q \in [0,1]$, with bounded derivative, such that the functions $\frac{s_m(\lceil\lceil q\rceil\rceil)}{m}$ converge uniformly to the function $\bar{s}(q)$. Using a construction similar to the construction of an independent hypothesis class in section 5, it is possible to show that for any non-negative bounded derivative function, $\bar{s}(q)$, on $q \in [0,1]$ there exists a data distribution and sequence of hypothesis classes with these properties.

We define a sample sequence $\vec{S}$ to be an infinite sequence of samples $S_1, S_2, S_3, \ldots$, such that S_m consists of m pairs $\langle x,\ y\rangle$ drawn IID from distribution D. For sample sequence $\vec{S}$ and $h \in \mathcal{H}_m$ we define $\hat{e}_m(h)$ and $\hat{s}_m(\frac{k}{m},\ \delta)$ in a manner similar to $\hat{e}(h)$ and $\hat{s}(\frac{k}{m},\ \delta)$ but for sample S_m in the vector $\vec{S}$. The main result of this section is the following.

Theorem 6.1 *Under the conditions described above, and for any $\delta > 0$ and $p \in [0,1]$, we have the following with probability 1 over the choice of the sample vector $\vec{S}$.*

$$\lim_{m\to\infty} \frac{\hat{s}_m(\lceil\lceil p\rceil\rceil,\delta)}{m} = \sup_{q\in[0,1]} \bar{s}(q) - D(p||q)$$

Before proving this theorem we note that the limiting value of $\hat{s}_m(\lceil\lceil p\rceil\rceil,\ \delta)$ is independent of δ. This is consistent with the argument given in section 4 that for large m we have that $\hat{s}_m(\lceil\lceil p\rceil\rceil, \delta)$ is independent of δ and approximately equal to the log of the number of hypothesis with empirical error rate of exactly $\lceil\lceil q\rceil\rceil$. Define $\bar{\hat{s}}(p)$ as follows.

$$\bar{\hat{s}}(p) \equiv \sup_{q\in[0,1]} \bar{s}(q) - D(p||q)$$

Note that we have $\bar{\hat{s}}(p) \geq \bar{s}(p)$ — so we get an asymptotic version of lemma 3.2. But it is not difficult to give examples of functions $\bar{s}(p)$ for which $\bar{\hat{s}}(p)$ is strictly larger than $\bar{s}(p)$.

In proving theorem 6.1 we first define $\mathcal{H}_m(p,\ q)$ for $p, q \in \{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}$ to be the set of all $h \in \mathcal{H}_m(q)$ such that $\hat{e}_m(h) = p$. Intuitively, $\mathcal{H}(p,\ q)$ is the set of concepts with true error rate near q that have empirical error rate p. Ignoring factors that are only polynomial in m, the probability of a hypothesis with true error rate q having empirical error rate p can be written as (approximately) $e^{-mD(p||q)}$. So the expected size of $\mathcal{H}(p,\ q)$ can be written as $|\mathcal{H}_m(q)|e^{-mD(p||q)}$, or alternatively, (approximately) as $e^{m\bar{s}(q)}e^{-mD(p||q)}$ or $e^{m(\bar{s}(q)-D(p||q))}$. More formally, we have the following for any fixed value of p and q.

$$\lim_{m\to\infty} \frac{\ln(\max(1,\ \mathrm{E}(|\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)|)))}{m}$$
$$= \max(0,\ \bar{s}(q) - D(p||q))$$

We now show that the expectation can be eliminated from the above limit. First, consider distinct values of p and q such that $\bar{s}(q) - D(p||q) > 0$. Since p and q are distinct, the probability that a fixed hypothesis in $\mathcal{H}(\lceil\lceil q\rceil\rceil)$ is in $\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)$ declines exponentially in m. Since $\bar{s}(q) - D(p||q) > 0$ the expected size of $\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)$ grows exponentially in m. Since the hypotheses are independent, the distribution of possible values of $|\mathcal{H}(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)|$ becomes essentially a Poisson mass distribution with an expected number of arrivals growing exponentially in m. The probability that $|\mathcal{H}(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)|$ deviates from its expectation by as much as a factor of 2 declines exponentially in m. We say that a sample sequence is safe after k if for all $m > k$ we have that $|\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)|$ is within a factor of 2 of its expectation. Since the probability of being unsafe at m declines exponentially in m, for any δ there exists a k such that with probability at least $1-\delta$ the sample sequence is safe after k. So for any $\delta > 0$ we have that with probability at least $1-\delta$ the sequence is safe after some k. But since this holds for all $\delta > 0$, with probability 1 such a k must exist.

$$\lim_{m\to\infty} \frac{\ln(\max(1, |\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)|))}{m}$$
$$= \bar{s}(q) - D(p||q)$$

We now define $s_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)$ as follows.

$$s_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil) \equiv \ln(\max(1,\ |\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)|))$$

It is also possible to show that with probability 1 we have that $\frac{s_m(\lceil\lceil q\rceil\rceil,\ \lceil\lceil q\rceil\rceil)}{m}$ approaches $\bar{s}(q) - D(p||q)$ and that for distinct p and q with $\bar{s}(q) - D(p||q) \leq 0$ we have that $\frac{s_m(\lceil\lceil q\rceil\rceil,\ \lceil\lceil q\rceil\rceil)}{m}$ approaches 0. Putting these together yields that with probability 1 we have the following.

$$\lim_{m\to\infty} \frac{s_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q\rceil\rceil)}{m} = \max(0,\ \bar{s}(q) - D(p||q)) \tag{15}$$

Since $U_m(p)$ is the set of hypothesis in $\mathcal{H}_m$ with an empirical error (in the mth sample) of exactly p, we have the following equality.

$$U_m(p) = \cup_{q\in\{\frac{1}{m},\ \ldots,\ \frac{m}{m}\}} \mathcal{H}_m(p,\ q)$$

We also define $u_m(p)$ to be $\ln(\max(1,\ |U_m(p)|))$. We will now show that with probability 1 we have that $\frac{u_m(p)}{m}$ approaches $\bar{\hat{s}}(p)$. First, consider a $p \in [0,1]$ such that $\bar{\hat{s}}(p) > 0$. Let Since $\bar{s}(q) - D(q||p)$ is a continuous function, and $[0,1]$ is a compact set, we have that $\sup_{q\in[0,1]} \bar{s}(q) - D(p||q)$ must be realized at some value $q^* \in [0,1]$. Let q^* be such that $\bar{s}(q^*) - D(p||q^*)$ equals $\bar{\hat{s}}(p)$. We have that $u_m(\lceil\lceil p\rceil\rceil) \geq s_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q^*\rceil\rceil)$. This, together with (15), implies the following.

$$\liminf_{m\to\infty} \frac{u_m(\lceil\lceil p\rceil\rceil)}{m} \geq \bar{\hat{s}}(p)$$

We will now say that the sample sequence is safe at m and $\frac{k}{m}$ if $|\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil \frac{k}{m}\rceil\rceil|$ does not exceed twice the expectation of $|\mathcal{H}_m(\lceil\lceil p\rceil\rceil,\ \lceil\lceil q^*\rceil\rceil|$. The probability of not being safe at m and $\frac{k}{m}$ declines exponentially in m at a rate at least as fast as the rate of decline of the probability of not being safe at m and $\lceil\lceil q^*\rceil\rceil$. By the union bound this implies that for a given m the probability that there exists an unsafe $\frac{k}{m}$ also declines exponentially. We say that the sequence is safe after N if it is safe for all m and $\frac{k}{m}$ with $m > N$. The probability of not being being safe after N also declines exponentially with N. By an argument similar to that given above, this implies that with probability 1 over the choice of the sequence there exists a N such that the sequence is safe after N. But if we are safe at m then $|U_m(\lceil\lceil p\rceil\rceil)| \leq 2mE|\mathcal{H}_m(p,\ \lceil\lceil q^*\rceil\rceil)|$. This implies the following.

$$\limsup_{m\to\infty} \frac{u_m(\lceil\lceil p\rceil\rceil)}{m} \leq \bar{\hat{s}}(p)$$

Putting the two bounds together we get the following.

$$\lim_{m\to\infty} \frac{u_m(\lceil\lceil p\rceil\rceil)}{m} = \bar{\hat{s}}(p)$$

Finally we must show the same result for $\hat{s}(\lceil\lceil p\rceil\rceil,\ \delta)$. The proof is similar to the proof of theorem 4.2 except that here we have not established the uniform convergence of the functions $\frac{u_m(\lceil\lceil p\rceil\rceil)}{m}$. But the convergence rate of this quantity at p is bounded by a continuous function of p — the convergence is locally uniform. Since the width of the interval of empirical errors contributing to $\hat{s}_m(\lceil\lceil p\rceil\rceil,\ \delta)$ is going to zero, locally uniform convergence suffices.

7 Future Work

A practical difficulty with the bound implicit in theorem 3.3 is that it is usually impossible to enumerate the elements of an exponentially large hypothesis class and hence impractical to compute the histogram of training errors for the hypotheses in the class. In practice the values of $s(\frac{k}{m})$ might be estimated using some form of Monte-Carlo Markov chain sampling over the hypotheses. For certain hypothesis spaces it might also be possible to directly calculate the empirical error distribution without evaluating every hypothesis.

Here we have emphasized asymptotic properties of our bound but we have not addressed rates of convergence. For finite sample sizes the rate at which the bound converges to its asymptotic behavior can be important. Before mentioning some ways that the convergence rate might be improved, however, we note that near phase transitions standard notions of convergence rate are inappropriate. Near a phase transition the bound is "unstable" — replacing δ by $\delta/2$ can alter the bound significantly. In fact, near a phase transition it is likely that $e(h^*)$ is significantly different for different samples even though $\hat{e}(h^*)$ is highly predictable. Intuitively, we would like a notion of convergence rate that measures the size of the "region of instability" around a phase transition. As the sample size increases the phrase transition becomes sharper and the region of instability smaller. It would be nice to have a formal definition for the region of instability and the rate at which the size of this region goes to zero, i.e., the rate at which phase transitions in the bound become sharp.

The rate of convergence of our bound might be improved in various ways:

- Removing the discretization of true errors.
- Using one-sided bounds.
- Using nonuniform union bounds over discrete values of the form $\frac{k}{m}$.
- Tightening the Chernoff bound using direct calculation of Binomial coefficients.
- Improving Lemma 3.4.

The above ideas may allow one to remove all $\ln(m)$ terms from the statement of the bound.

8 Conclusion

Traditional PAC bounds are stated in terms of the training error and class size or VC dimension. The computable bound given here is typically much tighter because it exploits the additional information in the histogram of training errors. The uncomputable upper bound uses the additional (unavailable) information in the distribution of true errors. Any distribution of true errors can be realized in a case with independent hypotheses. We have shown that in such cases this uncomputable bound is asymptotically equal to actual generalization error. Hence this is the tightest possible bound, up to asymptotic equality, over all bounds expressed as functions of $\hat{e}(h^*)$ and the distribution of true errors. We have also shown that the use of the histogram of empirical errors results in a bound that, while still tighter than traditional bounds, is looser than the uncomputable bound even in the large sample asymptotic limit.

Acknowledgments: Yoav Freund, Avrim Blum, and Tobias Scheffer all provided useful discussion in forming this paper.

References

[1] David Haussler, Michael Kearns, H. Sebastian Seung, and Naftali tishby, "Rigorous learning curve bounds from statistical mechanics", Machine Learning 25, 195-236, 1996.

[2] Thomas & Cover. "Elements of Information Theory", Wiley, 1991.

[3] H. Chernoff. "A measure of asymptotic efficiency for test of a hypothesis based on the sum of observations", Annals of Mathematical Statistics, 23:493-507, 1952.

[4] Tobias Scheffer and Thorsten Joachims, "Expected error analysis for model selection", International Conference on Machine Learning (ICML), 1999.

[5] Yoav Freund, "Self bounding algorithms", COLT, 1998.

[6] David McAllester, "Pac-Bayesian model averaging", COLT, 1999.

[7] John Langford and Avrim Blum, "Microchoice and self-bounding algorithms", COLT, 1999.

[8] Yishay Mansour and David McAllester, "Generalization Bounds for Decision Trees", COLT-2000.

[9] David McAllester and Robert Schapire, "On the Convergence Rate of Good-Turing Estimators", COLT-2000.