15-850(C) Machine Learning Theory
Probabilistic inequalities (draft)

A common question that comes up in machine learning is: "Given some fixed hypothesis h, how much data do I need to see so that I can be confident that its observed error will be near to its true error?" We have already seen this as: "What's the chance that a given hypothesis of true error ϵ or more will have observed error of 0?"

In this case, from first principles we calculated the probability to be $(1-\epsilon)^m$, where m is the number of examples. We then noticed that if there are $|H|$ hypotheses under consideration, and this quantity is $\delta/|H|$, then there is at most a δ chance there exists *any* hypothesis that fools us. Solving for m, we got that $\frac{1}{\epsilon}\left[\ln(|H|) + \ln(\frac{1}{\delta})\right]$ examples suffice.

More generally, suppose we consider a hypothesis with true error p, and let $q = 1 - p$. If we see m examples, then the expected number of mistakes is pm. The *standard deviation* of this quantity is $\sqrt{pqm}$, or equivalently, the standard deviation of the *fraction* of mistakes made is $\sqrt{pq/m}$. We can now use a convenient rule for independent identically distributed Bernoulli trials, which in our terminology is:

$$\Pr[|\text{observed error} - \text{true error}| > 1.96\sigma] < 0.05,$$

where $\sigma = \sqrt{pq/m}$ is the standard deviation. For instance, if we want with 95% confidence for our true and observed errors to differ by only ϵ, then we need to see only $4p(1-p)/\epsilon^2$ examples (approximating 1.96^2 by 4).

The above rule is convenient for testing out a single hypothesis on test data. More generally, we'd like a rule that can be applied for any desired confidence level (like the first case above) but holds even when the observed error is not 0 (like the second case above). For instance, we might want to prove uniform convergence results, or we might just want to test out several hypotheses.[1] Luckily, there are some convenient inequalities for doing this, known as Hoeffding and Chernoff bounds. These bounds state the following. (Actually, they hold for a variety of situations we won't describe here; a good book that discusses how these are derived and many other things is Alon and Spencer's *The Probabilistic Method*.)

Consider a hypothesis with true error rate p (or a coin of bias p) observed on m examples (the coin is flipped m times). Let S be the number of observed errors (the number of heads seen) so S/m is the observed error rate.

Hoeffding bounds state that for any $\epsilon \in [0, 1]$,

1. $\Pr[\frac{S}{m} > p + \epsilon] \leq e^{-2m\epsilon^2}$, and
2. $\Pr[\frac{S}{m} < p - \epsilon] \leq e^{-2m\epsilon^2}$.

[1] Say we test out one hypothesis and find its observed error to be too high using the 0.05 confidence measure. Then we run our algorithm some more, get another hypothesis, test that one with 0.05 confidence, and find it also is too high. Eventually we get one that passes our test. Are we in good shape?

Chernoff bounds state that under the same conditions,

1. $\Pr[\frac{S}{m} > p(1+\epsilon)] \leq e^{-mp\epsilon^2/3}$, and
2. $\Pr[\frac{S}{m} < p(1-\epsilon)] \leq e^{-mp\epsilon^2/2}$.

Let's do a quick example of using Hoeffding bounds. Suppose we'd like to say that with high probability $(1-\delta)$, all hypotheses in H have their observed error within ϵ of their true error. To achieve this, we just need for each hypothesis individually to have confidence parameter $\delta/|H|$. So, we set $2e^{-2m\epsilon^2} \leq \delta/|H|$, or:

$$2m\epsilon^2 \geq \ln(2|H|/\delta).$$

We can now either solve for ϵ as a function of m and δ or solve for m as a function of ϵ and δ. Solving for ϵ, we find that our goal is achieved for ϵ satisfying:

$$\epsilon \geq \sqrt{\frac{\ln(2|H|/\delta)}{2m}}.$$

This says roughly that it suffices to go $\sqrt{2\ln(2|H|/\delta)}$ standard deviations to get our desired confidence. Actually, this is a nice way of thinking about the meaning of Hoeffding bounds: they say that if we want to replace the "0.05" in the earlier equation with some arbitrary small δ', then it suffices to go about $\sqrt{2\ln(2/\delta')}$ standard deviations.

Solving for m we get:

$$m \geq \frac{1}{2\epsilon^2}\ln(2|H|/\delta) = \frac{1}{2\epsilon^2}[\ln(|H|) + \ln(2/\delta)].$$

This is a lot like the first PAC bound, but we are now quadratic in $1/\epsilon$.

How can we get rid of that pesky ϵ^2? We can do that by relaxing our goal, to say that for hypotheses whose true error is greater than ϵ, we are satisfied if their observed error comes just within a *factor of 2*. If you think about it, this is often all we need since we don't care so much about the high-error hypotheses. Now we can use Chernoff bounds.

If $p \geq \epsilon$ and we only need to be within a factor of 2 with confidence δ', we can solve that it suffices to have $m \geq \frac{8}{p}\ln(2/\delta')$ (and we can replace p with ϵ for a weaker bound). On the other hand, if $p < \epsilon$, we can rewrite our additive error goal by saying that we want the observed error to be no more than $p(1+\epsilon/p)$, which by Chernoff bounds implies that $m \geq \frac{3p}{\epsilon^2}\ln(1/\delta')$ examples suffice (and we can replace $\frac{3p}{\epsilon^2}$ with $\frac{3}{\epsilon}$ for a weaker bound). So, setting $\delta' = \delta/|H|$ we get that:

$$m \geq \frac{8}{\epsilon}[\ln(|H|) + \ln(2/\delta)]$$

examples suffice.