Homework 2 Solutions

Problem 1.

We show the equivalence of the one-button and two-button PAC models in two parts. We first show that if there exists an algorithm for efficiently PAC-learning $\mathcal{C}$ by $\mathcal{H}$ in the one-button model, then there exists an algorithm for efficiently PAC-learning $\mathcal{C}$ by $\mathcal{H}$ in the two-button model. We then show that if there exists an algorithm for efficiently PAC-learning $\mathcal{C}$ by $\mathcal{H}$ in the two-button model, then there exists an algorithm for efficiently PAC-learning $\mathcal{C}$ by $\mathcal{H} \cup \{$ TRUE,FALSE $\}$ in the one-button model.

(a) Let $\mathcal{A}_{1B}$ be an algorithm for efficiently PAC-learning $\mathcal{C}$ by $\mathcal{H}$ in the one-button model. Given error parameter ϵ', confidence parameter δ' and access to an example oracle $EX(c,\mathcal{D})$,$\mathcal{A}_{1B}$ will return a hypothesis $h \in \mathcal{H}$ in time polynomial in $1/\epsilon'$ and $1/\delta'$ such that error(h) is at most ϵ' with probability at least $1-\delta'$.

We can use $\mathcal{A}_{1B}$ to solve our learning problem in the two-button PAC model. Run $\mathcal{A}_{1B}$ with error parameter $\epsilon' = \epsilon/2$ and confidence parameter $\delta' = \delta$. When $\mathcal{A}_{1B}$ requests a labelled example, flip a fair coin and return a labelled example from $EX(c,\mathcal{D}^+)$ if heads or from $EX(c,\mathcal{D}^-)$ if tails. Output the hypothesis $h \in \mathcal{H}$ which $\mathcal{A}_{1B}$ eventually returns.

This procedure cleary runs in time polynomial in $1/\epsilon$ and $1/\delta$ (given the parameters input to $\mathcal{A}_{1B}$). I further claim that the hypothesis output will have error rate at most ϵ on both $\mathcal{D}^+$ and $\mathcal{D}^-$ (with probability at least $1-\delta$). Consider any "bad" hypothesis h_{bad} which has error rate greater than ϵ on either $\mathcal{D}^+$ or $\mathcal{D}^-$ (or both). This hypothesis must then have error rate greater than $\epsilon/2$ on the simulated distribution $\mathcal{D}$ according to which labelled examples were drawn for $\mathcal{A}_{1B}$. Hypotheses of this sort are output with probability less than δ (given the parameters input to $\mathcal{A}_{1B}$).

(b) Let $\mathcal{A}_{2B}$ be an algorithm for efficiently PAC-learning $\mathcal{C}$ by $\mathcal{H}$ in the two-button model. Given error parameter ϵ', confidence parameter δ' and access to two example oracles $EX(c,\mathcal{D}^+)$ and $EX(c,\mathcal{D}^-)$, $\mathcal{A}_{2B}$ will return a hypothesis $h \in \mathcal{H}$ in time polynomial in $1/\epsilon'$ and $1/\delta'$ such that h has error rate at most ϵ' on both $\mathcal{D}^+$ and $\mathcal{D}^-$ with probability at least $1-\delta'$.

We use $\mathcal{A}_{2B}$ to solve our learning problem in the one-button PAC model. Run $\mathcal{A}_{2B}$ with error parameter $\epsilon' = \epsilon$ and confidence parameter $\delta' = \delta$. When $\mathcal{A}_{2B}$ requests a positive example, draw up to k labelled examples from $EX(c,\mathcal{D})$ until a positive example is found (k will be determined below). If a positive example is found, then give this to $\mathcal{A}_{2B}$. If no positive example is found after drawing k labelled examples from $EX(c,\mathcal{D})$, then stop and return the hypothesis FALSE. When $\mathcal{A}_{2B}$ requests a negative example, draw up to k labelled examples from $EX(c,\mathcal{D})$ until a negative example is found. If a negative example is found, then give this to $\mathcal{A}_{2B}$. If no negative example is found after drawing k labelled examples from $EX(c,\mathcal{D})$, then stop and return the hypothesis TRUE. If positive and negative examples are always found when requested, then output the hypothesis $h \in \mathcal{H}$ which $\mathcal{A}_{2B}$ eventually returns.

We will now determine the value of k. Intuitively, if we draw a "large" sample and see only positive examples, then the "weight" of negative examples must be small (with high probability), and so the hypothesis TRUE will be accurate. Similarly, if we draw a "large" sample and see only negative examples, then the "weight" of positive examples must be small (with high probability), and so the hypothesis FALSE will be accurate.

Let $p(1/\epsilon, 1/\delta)$ be the number of examples which $\mathcal{A}_{2B}$ requires. Now suppose that TRUE is an ϵ-bad hypothesis. The probability that k consecutive true examples will be drawn from $EX(c,\mathcal{D})$ when a single negative example is requested is at most $(1-\epsilon)^k$. The probability that k consecutive true examples will *ever* be drawn from $EX(c,\mathcal{D})$ when negative examples are requested is at most $p(1/\epsilon, 1/\delta)\,(1-\epsilon)^k$. By choosing k large enough so that this probability is at most δ, we will insure that with probability at least $1-\delta$, TRUE is not output if it is ϵ-bad. A similar argument holds for FALSE. Solving the equation

$$p(1/\epsilon, 1/\delta)\,(1-\epsilon)^k \leq \delta$$

we find that $k = \Omega(\frac{1}{\epsilon}\log\frac{p(1/\epsilon,1/\delta)}{\delta})$ is sufficient. Thus, our total sample complexity is given by

$$\Omega\left(p(1/\epsilon, 1/\delta)\frac{1}{\epsilon}\log\frac{p(1/\epsilon, 1/\delta)}{\delta}\right)$$

which is polynomial as required.

Problem 2.

Here is a solution taken from Haussler, Kearns, Littlestone and Warmuth [1, pages 143–145]. Here, $LE(p, m, r)$ denotes the probability of at most r successes in m independent trials of a Bernoulli random variable with probability of success p. Recall that for $0 \leq \alpha \leq 1$, $LE(p, m, (1-\alpha)mp) \leq e^{-\alpha^2 mp/2}$. Note that the $\mathbf{regular}_2$ condition simply assumes that the concepts TRUE and FALSE are in the hypothesis class in the case of the one-button PAC model, and that the singleton concepts $\{x\}$ $\forall x \in X$ are in the hypothesis class in the case of the two-button PAC model.

References

[1] David Haussler, Michael Kearns, Nick Littlestone, and Manfred K. Warmuth. Equivalence of models for polynomial learnability. *Information and Computation*, 95(2):129–161, December 1991.