15-859(B) Machine Learning Theory

Homework # 4 **Due: March 19, 2014**

Groundrules: Same as before. You should work on the exercises by yourself but may work with others on the problems (just write down who you worked with). Also if you use material from outside sources, say where you got it.

Exercises:

1. [**Kernels and Similarity functions**] Consider an instance space X of strings of text (i.e., an example $x \in X$ is some text string, such as an email message).

   (a) Argue that the similarity function

   $K(x, x')$ = the number of substrings of length 5 that x and x' have in common (a substring is a contiguous sequence of characters)

   is a legal kernel. For instance, if "hello" appears somewhere in x and also appears somewhere in x' then $K(x, x') \geq 1$.

   (b) Argue that the similarity function

   $K(x, x') = 1$ if x and x' have at least one substring of length 5 in common, and $K(x, x') = 0$ otherwise

   is not a legal kernel. In particular, give three examples x, x', x'' such that the values $K(x, x), K(x, x'), K(x, x''), K(x', x'), K(x', x''), K(x'', x'')$ are not consistent with dot-products under any mapping of x, x', x'' into vectors.

Problems:

2. [**On the plausibility of boosting**] Suppose we have a finite hypothesis class H, a finite space of instances X (e.g., $X = \{0, 1\}^n$), and some unknown target function f. Suppose that for any distribution D over X there exists an $h \in H$ with error at most $1/2 - \gamma$. Without going through the full boosting analysis, use the minimax theorem to prove there must exist a function in WeightedMAJ(H) that is correct on all of X by margin at least 2γ. (Here, WeightedMAJ(H) is the class of weighted majority vote functions: functions of the form $f(x) = \mathrm{sgn}[\sum_{h_i \in H} \alpha_i h_i(x)]$ where $h_i(x) \in \{-1, 1\}$, $\alpha_i \geq 0$ and we normalize so that $\sum_i \alpha_i = 1$.) Then use Hoeffding bounds to prove that for any distribution D there must *exist* a hypothesis in $\mathrm{MAJ}_k(H)$ with error at most ϵ for $k = O(\frac{1}{\gamma^2} \log(1/\epsilon))$.

   Note: our boosting results said something even stronger because they gave us a way to efficiently produce the desired hypothesis, given a weak-learning oracle.

3. **[On approximate Nash equilibria]** A two-player general-sum game is like a two-player zero-sum game except that the players do not necessarily have opposite payoffs (it is really more an "interaction" than a "game"). Let us for concreteness focus on games where each player has n actions, and use R to denote the payoff matrix for the row player and C to denote the payoff matrix for the column player. (So if the row-player plays action i and the column-player plays action j, then the row-player gets R_{ij} and the column-player gets C_{ij}. A Nash Equilibrium is a pair of distributions p and q (one for each player) such that neither player has any incentive to deviate from its distribution assuming that the other player doesn't deviate from its distribution either.[1] Formally, a pair of distributions p (for the row player) and q (for the column player) is a Nash equilibrium if the following holds: assuming the column player plays at random from q, the expected payoff to the row player for each row i with $p_i > 0$ is equal to the maximum payoff out of all the rows ($e_i^T Rq = \max_{i'} e_{i'}^T Rq$); and, assuming the row player plays at random from p, the expected payoff to the column player for each column j with $q_j > 0$ is equal to the maximum payoff out of all the columns ($p^T Ce_j = \max_{j'} p^T Ce_{j'}$). (Here, e_i denotes the column-vector with a 1 in position i and 0 everywhere else).

   Now, assume we have a game in which all payoffs are in the range $[0, 1]$. Define a pair of distributions p, q to be an "ϵ-Nash" equilibrium if each player has *at most* ϵ incentive to deviate. That is, the expected payoff to the row player for each row i with $p_i > 0$ is within ϵ of the maximum payoff out of all the rows, and vice-versa for the column player.

   Using the fact that Nash equilibria must exist (proven by Nash in 1950), show that there must exist an ϵ-Nash equilibrium in which each player has positive probability on at most $O(\frac{1}{\epsilon^2}\log n)$ actions (rows or columns).

   Hint: this problem is related to problem 2.

   Note: this fact yields an $n^{O(\frac{1}{\epsilon^2}\log n)}$-time algorithm for finding an ϵ-Nash equilibrium. No PTAS (algorithm running in time polynomial in n for any fixed $\epsilon > 0$) is known, however.

4. **[Project]** Think about what you might want to do as a project. By March 21, either email me a paragraph on what you'd like to do or set up an appointment to talk with me about possible ideas (or both).

[1] Feel free to use the Web to learn more about general-sum games if you haven't seen them before.