15-251: Great Theoretical Ideas In Computer Science

Recitation 14

Important concepts from lecture

- A **Nash equilibrium** is a choice function from players to strategies such that no one player will benefit from changing their strategy.
- The **social cost** of a given solution (strategy choice function) is the sum of the costs to all players of the resulting game outcome.
- The **price of anarchy** is a metric to compare the social costs of "selfish" play (Nash equilibria) and some form of "cooperation" (the social-cost-optimal solution).
- A **consistent hypothesis** with respect to a set S of labelled points is any hypothesis that labels every point in S correctly.
- A **PAC learning algorithm** is one that, for any $\epsilon, \delta > 0$, any distribution D, and any $m_0(\epsilon, \delta)$ training points distributed according to D, has probability at least $1 - \delta$ of arriving at a hypothesis whose error with respect to D is at most ϵ.

The Only Winning Move

Consider the two-player game where each player, without knowledge of the other's choice, chooses a strategy from $S = \{0, 1\}$. Player 1 wins \$100 iff they both choose the same strategy, and Player 2 wins \$100 otherwise.

a) Show that there is no Nash equilibrium.

b) How can we modify S to preserve the nature of the game yet make it so that there will be at least one Nash equilibrium? (Hint: imagine actually playing this game, perhaps repeatedly. How can we make the given model more realistic?)

c) Characterize the set of Nash equilibria after this modification.

Alg-chemy

In lecture we saw that, if we're given an algorithm that always finds a consistent hypothesis, we can construct a PAC learning algorithm by finding a hypothesis consistent with some $m_0(\epsilon, \delta)$ points sampled from the input distribution.

Let's try going the other way: given a set S of labelled points, some $\delta > 0$, and a PAC learning algorithm A that may not always output a consistent hypothesis, devise a procedure to find a hypothesis that is consistent with S with probability at least $1 - \delta$.

Intersection classes

Let C_1 and C_2 be two concept classes. Define the "intersection class"

$$C = \{c \mid \exists c_1 \in C_1, c_2 \in C_2.\ \forall x \in X.\ c(x) = + \iff c_1(x) = c_2(x) = +\}$$

which is to say that every concept $c \in C$ is the intersection of some $c_1 \in C_1$ and $c_2 \in C_2$. Recall that for any set of examples S and any concept class C', $\pi_{C'}(S)$ is the number of ways of labeling examples in S using concepts from C'. Let $\pi_{C'}(m)$ be the max of $\pi_{C'}(S)$ over all m-sized example sets S, and show that $\pi_C(m) \leq \pi_{C_1}(m) \cdot \pi_{C_2}(m)$.