CS 683 Spring 2001 – Problem Set 4

Due in class Thursday, April 26

(1) In part, we motivated our discussion of random walks in terms of card shuffling; so let's try to prove a theorem about card shuffling. Suppose you are given n cards in an arbitrary order, and you want to scramble their order by a sequence of the following very simple steps: in one step, you do nothing with probability 1/2; otherwise, you pick two cards uniformly at random and interchange their positions in the order.

Prove there is a two-variable polynomial function $p(\cdot, \cdot)$ with the following property: For every ordering π of the cards, the probability that the order of the cards is π after $p(n, \varepsilon^{-1})$ steps of this process is between $(1-\varepsilon)/n!$ and $(1+\varepsilon)/n!$.

(2) When we discussed the permanent in class, we focused on "dense" bipartite graphs in order to guarantee that the number of near-perfect matchings would not exceed the number of perfect matchings by more than a polynomial factor. Show that such a guarantee cannot be made for general bipartite graphs. Specifically, for a graph G, let $npm(G)$ denote the number of its near-perfect matchings, and $pm(G)$ denote the number of its perfect matchings. Show that there is a function $f(\cdot)$ that grows (asymptotically) faster than any polynomial, so that for arbitrarily large n there exist n-node bipartite graphs for which $pm(G) \geq 1$ and $npm(G) \geq f(n) \cdot pm(G)$.

(3) Suppose you are given a black box that takes an input graph G, and returns an independent set of G drawn uniformly at random from among all independent sets of G. Design an algorithm, using this black box, that takes an input graph G, together with parameters $\delta > 0$ and $\varepsilon > 0$, and does the following: with probability at least $1 - \delta$ it approximates the number of independent sets in G to within a factor of $1 \pm \varepsilon$. The running time of the algorithm should be polynomial in δ^{-1}, ε^{-1} and the number of nodes of the graph; a call to the black box should be considered to take a single computational step.

(4) Suppose you are given two convex sets K_1 and K_2 in $\mathbf{R}^d$, each of volume 1, and each specified implicitly by the following type of black box: in a single step, the black box for K_i will provide you with a point drawn uniformly at random from K_i. These black boxes represent the only way you can gain information about the convex sets; you want to use the boxes to determine whether the sets intersect one another.

Design an algorithm that does the following: for a given parameter $\varepsilon > 0$, it should either report that $K_1 \cap K_2$ is non-empty; or it should report that the volume of $K_1 \cap K_2$ is at most ε. (In the latter case, the sets are "almost" disjoint; note that if $K_1 \cap K_2$ is non-empty, but has volume $\leq \varepsilon$, then either of the two answers is fine.) The algorithm should be correct with probability at least $1 - \delta$, for a given parameter $\delta > 0$. Its running time should be polynomial in d, ε^{-1}, and δ^{-1}; a call to the black box should be considered to take a single computational step.

(5) Suppose you are given a set P of n points in the plane, specified by their coordinates; each point $p \in P$ also has a secret *label* $\lambda(p)$ that is either "red" or "blue." You are promised that this red/blue distinction is determined as follows: there is a secret line L in the plane (not passing through any of the points) such that all points of P on one side of L are red, and all points of P on the other side are blue.

In addition to standard geometric operations on the points, you're allowed to interact with the point set by "querying" a particular point $p \in P$; in this way, you learn the label of P. Your goal is to perform some queries, and then approximately predict the labels of all points.

In particular, design a (potentially randomized) algorithm that does the following. For given parameters $\varepsilon > 0$ and $\delta > 0$, it first queries a number of points that is bounded by a polynomial in ε^{-1} and δ^{-1}, and independent of n. Then, in time polynomial in n, ε^{-1}, and δ^{-1}, it produces a labeling ℓ of all points in P. We say that ℓ "mis-classifies" $p \in P$ if $\ell(p) \neq \lambda(p)$. With probability at least $1-\delta$, the labeling ℓ produced by the algorithm should have the property that it mis-classifies at most εn of the points.

(6) (a) Let $\mathcal{C}$ denote the collection of all convex subsets of the plane $\mathbf{R}^2$. Show that the (infinite) set system $(\mathbf{R}^2, \mathcal{C})$ does not have a finite VC-dimension. Also, show that for any function $f(\cdot)$, there exists a finite set of points P in the plane such that no subset of P of size $\leq f(\varepsilon)$ forms an ε-net for P with respect to the set system $(\mathbf{R}^2, \mathcal{C})$.

(b) Our standard definition of an ε-net for a finite set P focused only on subsets of P. We've seen in (a) that with respect to the set system $(\mathbf{R}^2, \mathcal{C})$ of all convex sets in the plane, we cannot guarantee the existence of small ε-nets in this standard sense.

We can, however, restore a strong type of guarantee if we allow an ε-net for a set P to contain points that do not belong to P. Specifically, show that the following is true. There is a function $g(\cdot)$ so that for every $\varepsilon > 0$, and every finite set P of points in the plane, there is a set N of $\leq g(\varepsilon)$ points in the plane (not necessarily belonging to P) with the following property: if K is a convex subset of $\mathbf{R}^2$ and $|K \cap P| \geq \varepsilon |P|$, then $K \cap N$ is non-empty.