Problem Set 8

Randomized Algorithms

Due Tuesday, November 9

1. You are given a set of vectors $V = v_1, \ldots, v_m \in \mathbb{R}^n$, with all coordinates non-negative. The goal of this problem is to "sparsify" these vectors: construct $w_1, \ldots, w_k \in \mathbb{R}$ and $u_1, \ldots, u_k \in V$ such that

$$\overline{u} := \sum_{i=1}^{k} w_i u_i$$

and

$$\overline{v} := \sum_{i=1}^{m} v_i$$

satisfy

$$\overline{u}_j = (1 \pm \epsilon)\overline{v}_j \qquad \forall j. \tag{1}$$

That is to say, we can find a small weighted subset of vectors u that represents V in the sense that its sum is close in every coordinate. (This can be seen as spectral sparsification in the setting of diagonal matrices.)

We would like to solve this using the random sampling techniques developed in class: choose a probability distribution $p_1, \ldots, p_m \in \mathbb{R}$, and draw k independent samples (w, u), where each (w, u) equals $(\frac{1}{kp_i}, v_i)$ for $i \sim p$.

(a) For any particular choice of V and p, apply a Chernoff/Bernstein bound to get a bound for the k required for (1). The bound will depend on V and p.

(b) Now try to optimize p. For a *fixed* coordinate j, show a choice of p that achieves (1) with high probability, for $k = \Omega_\epsilon(\log n)$. Call this choice $p^{(j)}$.

(c) Now consider $p^* = \frac{1}{n}\sum_j p^{(j)}$. Show that if $k = \Omega(n \log n)$, this p^* will achieve (1) for *all* coordinates j with high probability.

What is the dependence on ϵ?

(d) For many inputs V, your result should work with many fewer samples (e.g., $k \ll n$). When does this happen? Give a data-dependent bound for the necessary k as a function of V.

(e) (Optional) Show a different, deterministic, algorithm that uses $k = n$ for all V.