HOMEWORK 1

Due: April 17th 2026

Instructions : Please turn in the solutions to any five of the problems. We recommend that you work out all the problems for understanding the lecture material.

In your solutions, you can use the following structure theorem that we proved in the lectures without proof:

Fact (ARV Structure Theorem). *Let $g = (g_1, g_2, \ldots, g_n)$ be jointly gaussian random variables with $\Omega(1) \leq \mathbb{E}[g_i^2] \leq 1$ for every i and $D(i,j) = \frac{1}{4}\mathbb{E}[(g_i - g_j)^2]$ such that $D(i,j) \leq D(i,k) + D(k,j)$ for every i,j,k. Let $A_0 = \{i \mid g_i \leq -1\}$ and $B_0 = \{i \mid g_i \geq 1\}$. Fix a Δ, let $E(H) = \{(i,j) \mid D(i,j) \leq \Delta\}$, and let σ be an ordering on the edges in $E(H)$. Then, for every constant $c' > 0$, there's a small enough constant $c > 0$, such that for $\Delta = c/\sqrt{\log n})$, with probability > 0 the sets $A := A_0 \setminus V(M_\sigma)$ and $B := B_0 \setminus V(M_\sigma)$ satisfy $|A||B| \geq \Omega(n^2)$ and $D(i,j) \geq \Delta$ for all $i \in A$, $j \in B$, where M_σ is the maximal matching of $A_0 \times B_0 \cap E(H)$ outputted by the greedy algorithm where the edges are processed in order σ.*

Exercises

Exercise 1 (Basic SoS Proofs).

1. *Let $f(x)$ be a degree d univariate polynomial in x that is non-negative over $\mathbb{R}$. Show that f is sum-of-squares of degree $\leq d/2$ polynomials. What is the bit-complexity of the SoS representation for f when f has rational coefficients?*
2. *Let A be symmetric and $f(x) = x^\top Ax$. Prove that $\|A\|_2\|x\|_2^2 - f(x)$ is sum-of-squares polynomial in x of degree 2.*
3. *Show that $\frac{1}{k}\sum_{i\leq k} a_i^2 - \frac{1}{k^2}(\sum_{i\leq k} a_i)^2$ is a sum-of-squares polynomial in indeterminates $a_1, a_2, \ldots, a_k$.*

Exercise 2 (Calculus Free Analysis of Max-Cut Rounding). *Let $\mu : \{-1,1\}^n \to \mathbb{R}$ be a pseudo-distribution of degree 2. Let G be a graph with n vertices and m edges E so that $\tilde{\mathbb{E}}_\mu f_G(x) = (1-\delta)m$ for $f_G = \frac{1}{4}\sum_{\{i,j\}\in E}(x_i - x_j)^2$.*

1. *For small enough constant $\delta > 0$, show that the rounding algorithm we discussed outputs a probability distribution μ' over $\{-1,1\}^n$ such that $\mathbb{E}_{\mu'} f_G(x) \geq (1 - O(\sqrt{\delta}))m$.*[1]
2. *Now, consider the following algorithm for computing maximum cut in a graph G: 1) compute a pseudo-distribution μ maximizing $\tilde{\mathbb{E}}_\mu f_G$, 2)*

[1] Hint: you may want to use Jensen's inequality at some point.

take a random cut and the cut chosen according to μ' and output the better of the two. Show that this algorithm achieves (in expectation) an approximation ratio $0.5 + \beta$ for some fixed constant $\beta > 0$.

Exercise 3 (Cycle Graphs and Max-Cut Integrality Gap). *Let C_n be the cycle graph on n vertices labeled by integers in $[n]$ where $\{i, j\}$ is an edge if $i - j \in 1, -1 \; mod(n)$. Let L be the Laplacian matrix of C_n satisfying $L(i, i) = 2$ for every i, $L(i, j) = -1$ if $i - j \in 1, -1 \; mod(n)$ and $L(i, j) = 0$ otherwise. For every $0 \leq k \leq n/2$, let x_k and y_k be vectors with coordinates $x_k(i) = cos(2\pi ik/n)$ and $y_k(i) = sin(2\pi ik/n)$.*

1. *Prove that x_k, y_k are eigenvectors of L with eigenvalues $2 - 2cos(2\pi k/n)$.*
2. *Prove that the diagonal entries of the matrix $M_k = x_k x_k^\top + y_k y_k^\top$ equal 1.*
3. *Let n be odd. Prove that there's a degree 2 pseudo-distribution μ_k on $\{-1, 1\}^n$ with $\tilde{\mathbb{E}}[x] = 0$ and $\tilde{\mathbb{E}}[xx^\top] = M_k$. Prove further that for $k = (n-1)/2$, $\tilde{\mathbb{E}}_{\mu_k}[f_G(x)] \geq (1 - O(1/n^2))n$.*

Exercise 4 (Cycle is not an integrality gap for degree 6 SoS). *Let $\mu : \{-1, 1\}^n \to \mathbb{R}$ be a pseudo-distribution of degree 6.*

1. *Prove the squared triangle inequality $\tilde{\mathbb{E}}_\mu (x_i - x_j)^2 \leq \tilde{\mathbb{E}}_\mu (x_i - x_k)^2 + \tilde{\mathbb{E}}_\mu (x_k - x_j)^2$.*
2. *Show that for f_{C_n} defined in the exercise above and any degree 6 pseudo-distribution μ, $\tilde{\mathbb{E}}_\mu F_{C_n}(x) \leq n - 1$.* [2,3]

Exercise 5 (Rounding for Max Bisection). *The Max Bisection problem*[4] *is a variant of max-cut where the input is a graph $G = G(V, E)$ with $2n$ vertices and the goal is to output a cut $(S, \bar{S})$ such that $|S| = n$ and $|E(S, \bar{S})| = |\{\{i, j\} \mid i \in S, j \notin S\}|$ is maximized. Let μ be a pseudo-distribution on $\{-1, 1\}^n$ such that $\tilde{\mathbb{E}}[x] = 0$ and $\frac{1}{4n^2} \sum_{i,j \leq 2n} |\tilde{\mathbb{E}}[x_i x_j]| \leq \delta$. For small enough δ, give a rounding algorithm that outputs a sample from a probability distribution μ' over $\{x \mid x \in \{-1, 1\}^{2n}, \sum_i x_i = 0\}$ and $\frac{1}{4} \mathbb{E}_{\mu'} \sum_{\{i,j\} \in E} (x_i - x_j)^2 \geq \alpha \frac{1}{4} \mathbb{E}_\mu \sum_{\{i,j\} \in E} (x_i - x_j)^2$ for some constant $\alpha > 0.878 - O(\sqrt{\delta})$.* [5]

Exercise 6 (Polynomial Time Algorithm for Min-Cut). *Let G be a undirected weighted graph on vertex set $[n]$ with non-negative integer weights $w : E \to \mathbb{N} \cup \{0\}$. Let s, t be two special vertices in G. Let μ be a pseudo-distribution of degree 4 on $\{-1, 1\}^n$ such that $\tilde{\mathbb{E}}[(x_s - x_t)^2] = 4$. Give a polynomial time algorithm to that outputs a sample x' from a distribution μ' on $\{-1, 1\}^n$ such that $\sum_{\{i,j\} \in E} w_{\{i,j\}} \mathbb{E}_{x' \sim \mu'} (x'_i - x'_j)^2 \leq$*

[2] While a degree 6 SoS proof of the squared triangle inequality is a one-liner, there's a slightly clever argument that shows a degree 4 SoS proof for the squared triangle inequality. As a result, we obtain that degree 4 SoS exactly "solves" max-cut in C_n (even though degree 2 cannot!).

[3] Hint: Work out the case for $n = 3$ first!

[4] Max bisection is an example of a problem wth "global constraint" - here, the constraint is that the cut be exactly balanced. Analyzing problems with global constraints such as these can be a bit tricky. This exercise shows that one can round pseudo-distributions for max bisection with low average pairwise correlations. We will discuss a method that uses $poly(1/\delta)$-degree SoS to compute degree 2 pseudo-distributions with low average pairwise correlations later on in this class. That will give us an algorithm with a non-trivial approximation guarantee for the max bisection problem.

[5] Hint: Prove that Gaussian rounding produces an approximately balanced cut with a large enough constant probability. "Fix" this cut to be perfectly balanced.

$\sum_{\{i,j\}\in E} w_{\{i,j\}} \tilde{\mathbb{E}}[(x_i - x_j)^2]$. Describe a randomized algorithm that uses the transformation $\mu \to \mu'$ constructed above to find a s-t cut in G in polynomial time with an expected weighted size at most that of the weighted min-s-t-cut in G.

Exercise 7 (Reduction to the Large Average Distance Case). *Let μ be a degree 4 pseudo-distribution on $\{-1,1\}^n$. Let $G([n], E)$ be a d-regular graph on n vertices identified as elements of $[n]$. Let $D : [n] \times [n]$ be defined by $D(i,j) = \frac{1}{4}\tilde{\mathbb{E}}_\mu[(x_i - x_j)^2]$. Suppose $\frac{1}{n^2}\sum_{i,j\le n} D(i,j) = \delta$.*

1. ***Heavy Cluster Case:** Suppose there exists an i^* such that $A = \{j \mid D(i^*, j) \le \delta/4\}$ satisfies $|A| \ge 0.1n$. Prove that $|A|\sum_{j\le n} D(j, A) \ge \Omega(n^2)\,\mathbb{E}_{i,j\in[n]} D(i,j)$.[6] Let μ' be the distribution obtained by choosing $t \in [0,1]$ uniformly at random and letting $x' \in \{\pm1\}^n$ where $x_i = 1$ iff $D(j, A) \le t$. Show that $\sum_{i,j\in[n]} \mathbb{E}_{\mu'}(x_i - x_j)^2 = \Omega(1)\sum_{i,j\in[n]} \tilde{\mathbb{E}}_\mu(x_i - x_j)^2$. Derive that*

$$\frac{\sum_{\{i,j\}\in E} \frac{1}{4}\mathbb{E}_{\mu'}(x_i - x_j)^2}{d/n\sum_{i,j\in[n]} \frac{1}{4}\mathbb{E}_{\mu'}(x_i - x_j)^2} \le O(1)\cdot \frac{\frac{1}{4}\sum_{\{i,j\}\in E} \frac{1}{4}\tilde{\mathbb{E}}_\mu(x_i - x_j)^2}{d/n\sum_{i,j\in[n]} \frac{1}{4}\tilde{\mathbb{E}}_\mu(x_i - x_j)^2}. \quad (1)$$

2. *Suppose that no i^* satisfying the condition in part (1) exists (i.e. there's no heavy cluster). Show that there is a Gaussian random vector g' and a set $U \subseteq [n]$, $|U| \ge 0.4n$ such that (1) $\mathbb{E}[g'_j] = 0$, (2) $\mathbb{E}[g'^2_j] \ge \frac{1}{8}$ and $\mathbb{E}[g'^2_j] \le 1$ for every $j \in U$, (3) $\mathbb{E}(g'_j - g'_k)^2 = D(j,k)/2\delta$ and (4) $\sum_{j,k\in U}\mathbb{E}[(g'_j - g'_k)^2] = \Omega(n^2)$. Use g' to find a A, B such that $|A||B| = \Omega(n^2)$ and $D(i,j) \ge \Omega(1/\sqrt{\log n})\delta$ for every $i \in A$, $j \in B$.* [7]

[6] Hint: Use triangle inequality to prove that for every i, j $D(i,j) \le (D(i,A) + D(j,A) + \delta/2)$.

[7] Hint for defining g': First, prove that there is an i_0 such that $\mathbb{E}_j D(i_0, j) \le \delta$. Let U be the set of all j such that $\frac{\delta}{4} \le D(i_0, j) \le 2\delta$. And for g, the Gaussian vector with mean 0 and $\mathbb{E}\, gg^\top = \tilde{\mathbb{E}}_\mu\, xx^\top$, define g' by setting $g'_j = (g_j - g_{i_0})/2\sqrt{2\delta}$.

Exercise 8 (Tightness of ARV Structure Theorem). *We will design a pseudo-distribution of degree 4 on $\{\pm1\}^n$ - it will actually be simply an* actual *distribution - for which, there cannot be $\Omega(n)$-size, $\gg 1/\sqrt{logn}$ separated sets.*

In the proof, you will use the Poincaré inequality [8] *for the Boolean hypercube.* For any $f : \{-1,1\}^k \to \mathbb{R}$, let $D_\ell f(z) = f(z) - f(z^{(\ell)})$ where $z^{(\ell)}$ is obtained by flipping the bit in the ℓ-th coordinate of z. Let $\mathrm{Inf}(f) = \mathbb{E}_{z\sim\{-1,1\}^k}\sum_{\ell=1}^{k}(D_\ell f(z))^2$ where $\mathbb{E}$ is over the uniform distribution on $\{-1,1\}^k$. Then, $\mathrm{Var}(f) = \mathbb{E}_{z,z'\sim\{-1,1\}^k}(f(z) - f(z'))^2 \le \mathrm{Inf}(f)$. *The Poincaré inequality can be interpreted as saying that Lipschitz functions on the hypercube have bounded variance.*

[8] This inequality can be proven with fairly elementary Fourier analysis on the Boolean hypercube. Such a proof appears on Page 51 of https://www.cs.tau.ac.il/~amnon/Classes/2016-PRG/Analysis-Of-Boolean-Functions.pdf.

Choose $n = 2^k$ for some k and associate each $1 \le i \le n$ with a point $z_i \in \{\pm1\}^k$ arbitrarily. Let μ be the probability distribution on $\{\pm1\}^n$ that picks a uniformly random $\ell \in [k]$ and outputs x where for every

$1 \le i \le n$, $x_i = z_i(\ell)$ - *the ℓ-th coordinate of z_i. Note that this distribution is supported on exactly $k = \log_2 n$ points.*

1. *Prove that* $\frac{1}{n^2} \sum_{i,j \le n} \frac{1}{4} \tilde{\mathbb{E}}_\mu (x_i - x_j)^2 = \frac{1}{2}$.

2. *Let $A, B \subseteq \{-1,1\}^k$ such that for every $z \in A$, $z' \in B$, $H(z,z') \ge k\Delta$ where H is the hamming distance and equals the number of coordinates at which the two strings in the argument to H differ. Prove that there's a function $f : \{-1,1\}^k \to [0,1]$ such that for every $z \in A$, $f(z) = 1$, for every $z \in B$, $f(z) = 0$ and for every z, z' such that $H(z,z') \le 1$, $|f(z) - f(z')| \le 1/k\Delta$.*

3. *For f constructed in the previous part, prove that $\mathbb{E}_{z,z' \in \{-1,1\}^k}(f(z) - f(z'))^2 \ge (|A|/2^k)(|B|/2^k)$. Apply Poincaré inequality on f to conclude that $|A| \cdot |B| \le 2^{2k} 1/(k\Delta^2)$.*

4. *Conclude that if $A, B \subseteq [n]$ such that for every $i \in A$, $j \in B$, $D(i,j) = \frac{1}{4} \tilde{\mathbb{E}}_\mu (x_i - x_j)^2 \ge \Delta$ then, $|A||B| \le n^2/(k\Delta^2)$.*

Exercise 9 (Better than 2-Approximation for Vertex Cover)**.** *Recall that a vertex cover in a graph G is a set of vertices that include at least one end point of every edge in G. In the minimum vertex cover problem, our goal is to compute a vertex cover of smallest possible size. You'll likely have studied a simple 2-approximate polynomial time algorithm for computing the minimum vertex cover in a graph. In this problem, we will design a $(2 - \Omega(1/\sqrt{\log n}))$-approximation algorithm for the minimum vertex cover.*

Let G be a d-regular, undirected, unweighted graph with edge set E and vertex set $[n]$ with minimum vertex cover of size OPT(G). Let μ be a degree 4 pseudo-distribution on $\{-1,1\}^n$ that, in addition, satisfies "edge-covering constraints" $\tilde{\mathbb{E}}_\mu (1 - x_i)(1 - x_j) = 0$ for every $\{i,j\} \in E$.[9]

[9] In the next class, we will define and study pseudo-distributions satisfying constraints that makes such constraints conceptually clearer.

1. *Show that there exists a pseudodistribution μ where additionally $\sum_i \tilde{\mathbb{E}}_\mu [(\frac{1}{2} + \frac{1}{2} x_i)] \le OPT(G)$. Give a 2-approximation algorithm by rounding μ.*

2. ***Small VC from Large Independent Set among Balanced Vertices:*** *We will now design a $<$ 2-approx. algorithm. Fix an ϵ to be chosen later. Let $S_1 = \{i \mid \tilde{\mathbb{E}}_\mu x_i > \epsilon\}$. And let $S_2 = \{i \mid |\tilde{\mathbb{E}}[x_i]| \le \epsilon\}$. Suppose $I \subseteq S_2$ is an independent set (i.e. no edge of G has both endpoints in I) of size $\Omega(|S_2|)$. Prove that $S_1 \cup S_2 \setminus I$ a vertex cover in G of size $(2 - \epsilon) \sum_{i \le n} (1/2 + \frac{1}{2} \tilde{\mathbb{E}}_\mu x_i)$.* [10]

[10] Hint: An edge can have at most one end point in I. Use the edge covering constraints to whether there can be an edge $\{i,j\}$ such that $i, j \notin S_1 \cup S_2$ or an edge where one end point is in I and the other in $[n] \setminus S_1 \cup S_2$.

3. ***"Antipodal" Concatenation:*** *Define the following pseudo-distribution μ_1 on $\{\pm 1\}^{2n}$ where we add a "dummy index" x'_i for every original x_i (let V' be the set of all dummy variables) and for any S, $\tilde{\mathbb{E}}[\prod_{j \in T} x_j \prod_{i \in S} x'_i] = (-1)^{|S|} \tilde{\mathbb{E}}[\prod_{j \in T} x_j \prod_{i \in S} x_i]$*[11]

[11] In an actual distribution, this operation corresponds to taking a sample x on $\{\pm 1\}^n$ from μ and outputting the concatenation of $(x, -x)$ as a sample from μ_1.

Prove that μ_1 is a degree 4 pseudo-distribution on $\{\pm 1\}^n$ satisfying $\frac{1}{4}\sum_{i,j \in S_2 \cup S_2'} \tilde{\mathbb{E}}[(x_i - x_j)^2] \geq |S_2|^2$ where $S_2' = \{i' \mid i \in S_2\}$.

4. ***Constructing Large Independent Set in S_2:*** *Let g be a Gaussian with same first and second moments as that of μ_1. Let $D(i,j) = \frac{1}{4}\tilde{\mathbb{E}}_{\mu_1}[(x_i - x_j)^2] = \frac{1}{4}\mathbb{E}(g_i - g_j)^2]$ for every i, j. Show that there exist sets $A \subseteq S_2 \cup S_2'$ and $B \subseteq S_2 \cup S_2'$ such that $|A||B| \geq \Omega(|S_2|^2)$, $D(i,j) \geq \Delta = \Omega(1/\sqrt{\log n})$ for all $i \in A$, $j \in B$, and if $i \in A$ then $i' \in B$, and if $i \in B$ then $i' \in A$.*

5. *Let $\varepsilon = \Delta/4$. Show that for any $\{i,j\} \in E \cap S_2 \times S_2$, it holds that $\{i,j\} \not\subseteq A$ and $\{i,j\} \not\subseteq B$. Thus, $A \setminus V'$ and $B \setminus V'$ are independent sets in G. Prove that at least one of $A \cap [n], B \cap [n]$ is of size $\Omega(|S_2|)$ to finish the proof.*