Homework 3 Solution

Release date: Oct. 16, Due date: Oct. 23.

Guidelines: Submit your self-grade in pdf format on Gradescope by **5pm on Oct. 23**. Please compare your solutions with the posted ones and assign yourself points honestly:

- Full credit: essentially correct
- 50% credit: good progress but errors/incomplete
- 25% credit: some relevant ideas only
- No credit: no real attempt

If you take partial credit, briefly note why.

Q 1 (Sampling on bipartite expanders, 3 pts)**.** Let $G = (L, R, E)$ be a biregular bipartite graph with normalized biadjacency matrix B whose second largest singular value is λ. Prove the following concentration result. For any subset $S \subset L$ s.t. $\frac{|S|}{|L|} = \varepsilon$ and any $C > 0$,

$$\Pr_{v \sim \text{Uniform}(R)} \left[\frac{|N(v) \cap S|}{|N(v)|} > (1 + C) \cdot \varepsilon \right] < \frac{\lambda^2}{\varepsilon C^2} \ .$$

From this concentration result, we can deduce that when $\lambda^2 \ll \varepsilon$, most of the $v \in R$ has $\approx \varepsilon$-fraction of its neighbors inside S. So in other words $N(v)$ (where $v \sim \text{Uniform}(R)$) samples L well.

Hint: Expander mixing lemma.

Solution

Let $f = \vec{1}_S : L \to \{0, 1\}$ denote the indicator function of the set S. For each $v \in R$, define

$$X_v = \frac{|N(v) \cap S|}{|N(v)|} = (Bf)(v),$$

where B is the normalized bipartite biadjacency matrix. Equip these function spaces with the inner product

$$\langle g, h \rangle_{\text{Uniform}(R)} = \mathbb{E}_{v \sim \text{Uniform}(R)}[g(v)h(v)],$$

and the induced norm $\|g\|^2_{\text{Uniform}(R)} = \mathbb{E}[g(v)^2]$.

Expectation. Since B preserves the constant function $\vec{1}$, we have

$$\mathbb{E}_{v \sim \text{Uniform}(R)}[X_v] = \langle Bf, \vec{1} \rangle_{\text{Uniform}(R)} = \langle f, \vec{1} \rangle_{\text{Uniform}(L)} = \varepsilon.$$

Variance. By definition,

$$\mathrm{Var}(X_v) = \mathbb{E}[(X_v - \varepsilon)^2] = \|B(f - \varepsilon\vec{1})\|^2_{\mathrm{Uniform}(R)}.$$

Let $g = f - \varepsilon\vec{1}$, so that $\langle g, \vec{1}\rangle_{\mathrm{Uniform}(L)} = 0$. The singular values of B are $1 = \sigma_1 \geq \sigma_2 \geq \cdots$, with the top singular function $\vec{1}$, then for all $g \perp \vec{1}$,

$$\|Bg\|_{\mathrm{Uniform}(R)} \leq \lambda\|g\|_{\mathrm{Uniform}(L)}.$$

Therefore,

$$\mathrm{Var}(X_v) \leq \lambda^2\|f - \varepsilon\vec{1}\|^2_{\mathrm{Uniform}(L)} = \lambda^2(\varepsilon - \varepsilon^2) < \lambda^2\varepsilon.$$

Concentration. By Chebyshev's inequality, for any $C > 0$,

$$\Pr_{v\sim\mathrm{Uniform}(R)}[X_v > (1+C)\varepsilon] = \Pr_v[|X_v - \varepsilon| > C\varepsilon] \leq \frac{\mathrm{Var}(X_v)}{(C\varepsilon)^2} < \frac{\lambda^2}{\varepsilon C^2}.$$

Hence,

$$\Pr_{v\sim\mathrm{Uniform}(R)}\left[\frac{|N(v)\cap S|}{|N(v)|} > (1+C)\varepsilon\right] < \frac{\lambda^2}{\varepsilon C^2}.$$

□

Q 2 (Expansion and correlation, 3 pts)**.** Let $G = (V = [n], E)$ be a graph on n vertices. Let π_1 be G's associated edge distribution. As we have seen in the lectures, given G and π_1 we can define a random walk on G with transition matrix P and from there get a stationary distribution π_0 over V.

Now define the distribution $\mu_G : V^2 \to \mathbb{R}$ as

$$\mu_G((i,j)) = \begin{cases} \pi_1(\{i,j\})/2 & \text{if } \{i,j\} \in E \\ 0 & o.w. \end{cases}$$

Since μ_G is the joint distribution of 2 variables X and Y, their correlation coefficient is defined as

$$\rho = \frac{\mathbb{E}_{\mu_G}[XY] - \mathbb{E}_{\mu_G}[X]\cdot\mathbb{E}_{\mu_G}[Y]}{\sqrt{\mathrm{Var}_{\mu_G}(X)\cdot\mathrm{Var}_{\mu_G}(Y)}}.$$

A variational definition of $|\rho|$ is

$$|\rho| = \max_{f,g \text{ not constant functions}} \frac{|\mathbb{E}_{\mu_G}[f(X)g(Y)] - \mathbb{E}_{\mu_G}[f(X)]\cdot\mathbb{E}_{\mu_G}[g(Y)]|}{\sqrt{\mathrm{Var}_{\mu_G}(f(X))\cdot\mathrm{Var}_{\mu_G}(g(Y))}}.$$

Prove that $|\rho| = |\lambda|_2(P)$ where the eigenvalues of P are defined with respect to the inner product $\langle u, w\rangle_{\pi_0} = \mathbb{E}_{v\sim\pi_0}[u(v)\cdot w(v)]$.

Solution

Since P is the transition matrix of the random walk on G, it has real eigenvalues

$$1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq -1,$$

where the constant function $\vec{1}$ is the eigenvector for $\lambda_1 = 1$.

For any functions $f, g : V \to \mathbb{R}$, note that

$$\mathbb{E}_{\mu_G}[f(X)g(Y)] = \sum_{i,j} \mu_G(i,j)f(i)g(j) = \sum_i \pi_0(i)f(i)\sum_j P(i,j)g(j) = \langle f, Pg\rangle_{\pi_0}.$$

Thus

$$\mathbb{E}_{\mu_G}[f(X)g(Y)] - \mathbb{E}_{\mu_G}[f(X)]\mathbb{E}_{\mu_G}[g(Y)] = \langle f, Pg\rangle_{\pi_0} - \langle f, \vec{1}\rangle_{\pi_0}\langle \vec{1}, g\rangle_{\pi_0}.$$

Let $\bar{f} = f - \mathbb{E}_{\pi_0}[f]$ and $\bar{g} = g - \mathbb{E}_{\pi_0}[g]$ be the mean-0 functions. Since $P\vec{1} = \vec{1}$, we get

$$\mathbb{E}_{\mu_G}[f(X)g(Y)] - \mathbb{E}_{\mu_G}[f(X)]\mathbb{E}_{\mu_G}[g(Y)] = \langle \bar{f}, P\bar{g}\rangle_{\pi_0}.$$

Moreover,

$$\mathrm{Var}_{\mu_G}(f(X)) = \|\bar{f}\|_{\pi_0}^2, \quad \mathrm{Var}_{\mu_G}(g(Y)) = \|\bar{g}\|_{\pi_0}^2.$$

Plugging these into the variational definition gives

$$|\rho| = \sup_{\bar{f},\bar{g}\neq 0,\ \mathbb{E}[\bar{f}]=\mathbb{E}[\bar{g}]=0} \frac{|\langle \bar{f}, P\bar{g}\rangle_{\pi_0}|}{\|\bar{f}\|_{\pi_0}\,\|\bar{g}\|_{\pi_0}}.$$

This is exactly the operator norm of P restricted to the subspace

$$\vec{1}^{\perp} = \{h \in \mathbb{R}^V : \langle h, \vec{1}\rangle_{\pi_0} = 0\}.$$

Therefore,

$$|\rho| = \max(\lambda_2(P), |\lambda_n(P)|) = |\lambda|_2(P).$$

□

Q 3. (High dimensional expansion and correlation, 6 pts) The previous problem says that expansion of a graph G can be viewed as bounded correlation between random variables from the corresponding distribution μ_G. This can be generalized to simplicial complexes. Given a d-dimensional simplicial complex X with the distribution π_d over its d-faces. Now we define the distribution $\mu_X : X(0)^{d+1} \to \mathbb{R}$ as

$$\mu_X((Y_0, \ldots, Y_d)) = \begin{cases} \pi_d(\{Y_0, \ldots, Y_d\})/(d+1)! & \text{if } \{Y_0, \ldots, Y_d\} \in X(d) \\ 0 & o.w. \end{cases}$$

Prove the following statements

1. (1.5 pts) The marginal distribution over the first two variables $\mu_X|_{Y_0,Y_1}$ is equal to the distribution μ_G where G is the 1-skeleton of X equipped with the edge distribution π_1 induced by π_d.
2. (1.5 pts) If we condition on $Y_0 = v$ for some $v \in X(0)$, then the conditional distribution $(\mu_X|Y_0 = v)$ is equal to μ_{X_v} which is the distribution associated with the link of v (the $(d-1)$-face distribution π_v in this complex is also induced by π_d).
3. (3 pts) Let P_v be the random walk on the 1-skeleton of X_v (with transition probability induced by π_d). Let $\nu = (\mu_X \mid Y_0 = v)|_{Y_1,Y_2}$ be the marginal distribution of the first two variables Y_1 and Y_2 in the conditional distribution $(\mu_X \mid Y_0 = v)$. Prove that $|\rho(\nu)| = |\lambda|_2\,(P_v)$.

So high dimensional expansion of X can be viewed as bounded correlation between two random variables from the corresponding μ_X and also from the conditional distributions $(\mu_X|Y_0 = v_0, \ldots, Y_t = v_t)$ for all $t \leq d-2$ and $\{v_0, \ldots, v_t\} \in X(t)$.

Solution

1. Fix vertices $a, b \in X(0)$. The marginal probability

$$\mu_X|_{Y_0,Y_1}((Y_0, Y_1) = (a, b)) = \sum_{y_2,\cdots,y_d} \mu_X((a, b, y_2, \ldots, y_d))$$

sums over all ordered $(d+1)$-tuples whose first two coordinates are (a, b). Each unordered d-face $\tau \in X(d)$ that contains the unordered pair $\{a, b\}$ contributes to this sum exactly the number of orderings of the remaining $d-1$ vertices, i.e. $(d-1)!$ orderings. Hence

$$\mu_X|_{Y_0,Y_1}((a, b)) = \sum_{\tau \in X(d):\{a,b\} \subset \tau} \frac{\pi_d(\tau)}{(d+1)!} \cdot (d-1)! = \sum_{\tau \supset \{a,b\}} \frac{\pi_d(\tau)}{d(d+1)}.$$

On the other hand the edge distribution π_1 induced from π_d is defined by: pick a random d-face $\tau \sim \pi_d$ and then pick an uniform unordered edge of τ. Thus for an unordered edge $\{a, b\}$

$$\pi_1(\{a, b\}) = \sum_{\tau \supset \{a,b\}} \frac{\pi_d(\tau)}{\binom{d+1}{2}} = \sum_{\tau \supset \{a,b\}} \frac{2\pi_d(\tau)}{d(d+1)}.$$

Comparing the two displays we get

$$\mu_X|_{Y_0,Y_1}((a, b)) = \frac{1}{2}\pi_1(\{a, b\}) = \mu_G((a, b)).$$

2. Fix $v \in X(0)$. For an ordered $(d-1)$-tuple $(y_1, \ldots, y_d)$ of vertices we compute the conditional probability

$$(\mu_X \mid Y_0 = v)(y_1, \ldots, y_d) = \frac{\mu_X((v, y_1, \ldots, y_d))}{\Pr[Y_0 = v]}.$$

If $\{v, y_1, \ldots, y_d\} \notin X(d)$ the numerator is 0. If $\tau = \{v, y_1, \ldots, y_d\} \in X(d)$ then the numerator equals $\pi_d(\tau)/(d+1)!$ (there is exactly one ordering of τ with v in the first coordinate and the remaining coordinates fixed). The denominator is

$$\Pr[Y_0 = v] = \sum_{\tau \ni v} \frac{\pi_d(\tau)}{(d+1)!} \cdot d! = \frac{1}{d+1} \sum_{\tau \ni v} \pi_d(\tau).$$

Thus for $\tau = \{v, y_1, \ldots, y_d\} \in X(d)$,

$$(\mu_X \mid Y_0 = v)(y_1, \ldots, y_d) = \frac{\frac{\pi_d(\tau)}{(d+1)!}}{\frac{1}{d+1} \sum_{\tau' \ni v} \pi_d(\tau')} = \frac{\pi_d(\tau)}{d! \sum_{\tau' \ni v} \pi_d(\tau')}.$$

Define the *link distribution* π_v on $(d-1)$-faces of the link of v by

$$\pi_v(S) = \frac{\pi_d(S \cup \{v\})}{\sum_{\tau' \ni v} \pi_d(\tau')}, \qquad S \in X_v(d-1).$$

Then the distribution μ_{X_v} associated to the $(d-1)$-complex X_v assigns to the ordered tuple $(y_1, \dots, y_d)$ the probability

$$\mu_{X_v}(y_1, \dots, y_d) = \frac{\pi_v(\{y_1, \dots, y_d\})}{d!} = \frac{\pi_d(\tau)}{d! \sum_{\tau' \ni v} \pi_d(\tau')},$$

which matches the expression for $(\mu_X \mid Y_0 = v)(y_1, \dots, y_d)$. Hence

$$(\mu_X \mid Y_0 = v) = \mu_{X_v}.$$

3. By part (2) ν is exactly the marginal of μ_{X_v} on its first two coordinates. As in part (1) the distribution ν can be described as follows: choose a $(d-1)$-face S of the link according to π_v, then choose an ordering of S uniformly and output its first two coordinates. Equivalently ν is the distribution on ordered pairs obtained by sampling an edge of the 1-skeleton of the link according to the induced edge distribution and then ordering the endpoints uniformly.

   Let α be the marginal distribution of Y_1 under ν: $\alpha(u) = \Pr[Y_1 = u]$. Define the linear operator $T : L^2(\alpha) \to L^2(\alpha)$ by

   $$(Tf)(u) := \mathbb{E}[f(Y_2) \mid Y_1 = u] = \sum_w \frac{\nu(u, w)}{\alpha(u)} f(w).$$

   By construction T is exactly the transition operator P_v of the random walk on the 1-skeleton of the link X_v. Moreover ν satisfies $\nu(u, w) = \alpha(u) P_v(u, w)$, and because ν arises by symmetrically ordering edges we have detailed balance: $\alpha(u) P_v(u, w) = \alpha(w) P_v(w, u)$. Thus the spectrum of P_v is in $[-1, 1]$. The constant function $\vec{1}$ is an eigenfunction with eigenvalue 1; denote the second largest eigenvalue in absolute value by $|\lambda|_2 (P_v)$.

   Now applying the statement from Q2 to P_v, ν, and α, we conclude that

   $$|\rho(\nu)| = |\lambda|_2 (P_v).$$

$\square$

Q 4. (Challenge question , 0 pt) In lecture we proved that any d-dimensional two-sided γ-expander satisfies $\left\| M_i^+ - U_{i-1} D_i \right\|_{op} \leq \gamma$ for $0 \leq i < d$. We now prove the following (almost) converse: If X satisfies $\left\| M_i^+ - U_{i-1} D_i \right\|_{op} \leq \gamma$ for $0 \leq i < d$, X is also a d-dimensional two-sided $3d \cdot \gamma$-expander.

Hint: If the 1-skeleton of a link has a bad eigenfunction, what is the natural way to extend it to the entire complex? Also consider using the decompositions of M_i^+ and $U_{i-1} D_i$ that we used to show the original statement.

Solution from Theorem 5.5 of https://arxiv.org/abs/1804.08155

Step 1: extending a local eigenfunction

Fix an index $i \in \{1, \dots, d-1\}$ and a face $\sigma \in X(i-1)$. Let X_σ denote the link of σ, and let P_σ be the transition operator of the simple random walk on the 1-skeleton of X_σ, with stationary distribution π_σ.

Consider any eigenfunction $h : X_\sigma(0) \to \mathbb{R}$ of X_σ with $\mathbb{E}_{\pi_\sigma}[h] = 0$.

We extend h to a function $f : X(i) \to \mathbb{R}$ by setting

$$f(\tau) = \begin{cases} h(v), & \text{if } \tau = \sigma \cup \{v\} \text{ for some } v \in X_\sigma(0), \\ 0, & \text{if } \sigma \not\subset \tau. \end{cases}$$

That is, f is supported on the i-faces that contain σ, and within that support f corresponds exactly to h on the link vertices.

Note that by construction $\|f\|_{\pi(i)}^2 = (i+1) \cdot \pi_{i-1}(\sigma) \cdot \|h\|_{\pi_\sigma}^2$.

Step 2: Inner products with respect to M_i^+ and $U_{i-1}D_i$

Both operators (the down-then-up operator) preserve the subspace V_σ. As shown in lecture, M_i^+ and $U_{i-1}D_i$ actions on f as follows

Specifically, we have that

$$\left\langle M_i^+ f, f \right\rangle = \mathbb{E}_{s \sim \pi(i-1)} \left\langle f|_s, P_s f|_s \right\rangle_{\pi_s} = \pi_{i-1}(\sigma) \left\langle h, P_\sigma h \right\rangle_{\pi_\sigma}.$$

For $U_{i-1}D_i$ we aim to show

$$\langle U_{i-1} D_i f, f \rangle \le \frac{i+1}{i} \gamma \, \|f\|_{\pi(i)}^2.$$

Consider the inner product $\langle D_i f, D_i f \rangle$:

$$\begin{aligned} \langle D_i f, D_i f \rangle &= \mathbb{E}_{s \sim \pi(i-1), t_1, t_2 \sim \pi_s(0)^2}[f(t_1) f(t_2)] \\ &= \pi_{i-1}(\sigma) \, \mathbb{E}_{t_1, t_2 \sim \pi_\sigma(0)^2}[f(t_1) f(t_2)] + (1 - \pi_{i-1}(\sigma)) \, \mathbb{E}_{s \sim \pi(i-1), t_1, t_2 \sim \pi_s(0)^2}[f(t_1) f(t_2) \mid s \neq \sigma]. \end{aligned} \tag{1}$$

The first term is zero, because

$$\mathbb{E}_{t_1, t_2 \sim \pi_\sigma(0)^2}[f(t_1) f(t_2)] = \mathbb{E}_{\pi_\sigma}[h]^2 = 0.$$

We now analyze the second term. For any $s \neq \sigma$, there is at most one i-face $t = s \cup \sigma$ such that $f(t) \neq 0$. Therefore, $f(t_1) f(t_2) \neq 0$ only when $t_1 = t_2 = s \cup \sigma$. For each $t_1 \in X(i)$, define the event E_{t_1} to be $[t_2 = t_1]$. Then equation (1) becomes

$$(1 - \pi_{i-1}(\sigma)) \, \mathbb{E}_{t_1 \sim \pi(i)} \Big[f^2(t_1) \Pr_{s \subset t_1, t_2 \supset s}[E_{t_1} \mid s \neq \sigma] \Big].$$

If we prove that for every $t_1 \in X(i)$,

$$\Pr_{s \subset t_1, t_2 \supset s}[E_{t_1} \mid s \neq \sigma] \le \frac{i+1}{i} \gamma,$$

then it follows that

$$(1) = (1 - \pi_{i-1}(\sigma)) \, \mathbb{E}_{t_1 \sim \pi(i)} [f^2(t_1) \Pr_{s \subset t_1, t_2 \supset s}[E_{t_1}]] \le \frac{i+1}{i} \gamma \, \|f\|_{\pi(i)}^2. \tag{2}$$

Thus, we are left to prove that for all $t_1 \in X(i)$,

$$\Pr_{s \subset t_1, t_2 \supset s}[E_{t_1} \mid s \neq \sigma] \le \frac{i+1}{i} \gamma.$$

We first bound the unconditioned probability $\Pr_{s\subset t_1, t_2\supset s}[E_{t_1}] = \Pr_{s\subset t_1, t_2\supset s}[t_2 = t_1]$. Fix an $t_1 \in X(i)$, and let $\mathbf{1}_{t_1} : X(i) \to \{0,1\}$ denote its indicator. Notice that $U_{i-1}D_i\mathbf{1}_{t_1}(t_1) = \Pr[E_{t_1}]$, and so

$$\langle U_{i-1}D_i\mathbf{1}_{t_1}, \mathbf{1}_{t_1}\rangle = \pi_i(t_1)\, U_{i-1}D_i\mathbf{1}_{t_1}(t_1) = \pi_i(t_1)\, \Pr[E_{t_1}].$$

Because the operator M_i^+ has no self-loops, $\langle M_i^+\mathbf{1}_{t_1}, \mathbf{1}_{t_1}\rangle = 0$. So have

$$\langle U_{i-1}D_i\mathbf{1}_{t_1}, \mathbf{1}_{t_1}\rangle = \langle (U_{i-1}D_i - M_i^+)\mathbf{1}_{t_1}, \mathbf{1}_{t_1}\rangle \le \gamma\, \pi_i(t_1).$$

Hence $\Pr[E_{t_1}] \le \gamma$.

Now consider $t_1 \in X(i)$ containing σ, and let s be a random $(i-1)$-face contained in t_1. The probability that $s \neq \sigma$ is $\frac{i}{i+1}$, so

$$\Pr[E_{t_1} \mid s \neq \sigma] \le \frac{i+1}{i}\, \Pr[E_{t_1}] \le \frac{i+1}{i}\, \gamma,$$

and (2) holds. Finally we plug in the bounds to get that,

$$\begin{aligned}
\frac{\left|\langle h, P_\sigma h\rangle_{\pi_\sigma}\right|}{\|h\|_{\pi_\sigma}^2} &= \frac{(i+1)\cdot\left|\left\langle M_i^+ f, f\right\rangle\right|}{\|f\|_{\pi(i)}^2} \\
&\le \frac{(i+1)\cdot\left(\left|\left\langle (M_i^+ - U_{i-1}D_i)f, f\right\rangle\right| + \left|\langle U_{i-1}D_i f, f\rangle\right|\right)}{\|f\|_{\pi(i)}^2} \\
&\le \frac{i+1}{\|f\|_{\pi(i)}^2}\left(\gamma\, \|f\|_{\pi(i)}^2 + \gamma\frac{i+1}{i}\, \|f\|_{\pi(i)}^2\right) \\
&\le 3(i+1)\gamma\, \|f\|_{\pi(i)}^2
\end{aligned}$$

Hence every link X_σ has

$$|\lambda|_2(P_\sigma) \le 3d\gamma.$$

□