CS265, Fall 2020

Class 19: Agenda, Questions, and Links

1 Warm-Up

(There's no PollEv link for this one).

Suppose that X is a k-source on $\{0,1\}^n$. Let $N = 2^n$. Let $\sigma \in \mathbb{R}^N$ be the "vectorized" version of the pmf of X. That is,

$$\sigma_i = \Pr[X = i] \qquad \forall i \in \{0, \ldots, N-1\},$$

where we associate a number $i < N$ with its binary expansion in $\{0,1\}^n$.

1. Why is $\|\sigma\|_\infty \leq 2^{-k}$?
2. Argue that $\|\sigma\|_2 \leq 2^{-k/2}$.

   ***Hint**: Use the fact that for any vector x, $\|x\|_2^2 \leq \|x\|_\infty \|x\|_1$ (why is this true?).*

2 Announcements

- HW8 (THE LAST ONE!) is out now, due Friday!
- Wednesday will be our last class! Mary will give a research talk during class.
- Please fill out course evaluations on Axess!

3 Questions?

Any questions from the minilectures and/or the quiz and/or the warm-up? (Expanders, extractors?)

- Go into small groups and ask each other your questions.
- Go to `https://pollev.com/cs265` and ask your questions/comments there, or else upvote others' questions.

4 Extractors *from* expanders

Recall the definition of a (k, ε)-extractor:

Definition 1. *A function* $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ *is a* (k, ε) *extractor if, for all* k*-sources* X *on* $\{0,1\}^n$, $\|\mathsf{Ext}(X, U_d) - U_m\| \leq \varepsilon$.

Above, $\|\cdot\|$ is the total variation distance, and U_d refers to the uniform distribution on d bits.

Suppose that $G = (V, E)$ is an (undirected, unweighted) degree-D expander graph with $|V| = N$, and with expansion parameters $\lambda(G) \leq 1/2$. Recall that

$$\lambda(G) = \max\{\lambda_2, |\lambda_N|\},$$

where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ are the eigenvalues of A, where A is the *normalized adjacency matrix* of G. (aka, A_{ij} is $1/D$ if $\{i, j\} \in E$ and is zero otherwise).

At this point, there will be some slides illustrating a construction of an extractor. A description is below for reference.

Let $\varepsilon > 0$. Let $N = 2^n$ and fix some arbitrary bijection between $\{0,1\}^n$ and V, where V is the vertex set of G above. Fix any $k \leq n$. Let $d = \log(D) \cdot \ell$, where

$$\ell = (n - k)/2 + \log(1/\varepsilon),$$

.

Consider the following function $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^n$.
On input $x, s \in \{0,1\}^n \times \{0,1\}^d$:

- Treat $x \in \{0,1\}^n$ as an element of V.
- Treat $s \in \{0,1\}^d$ as a string of ℓ numbers in $\{0, 1, \ldots, D\}$. (That is, break up s into ℓ chunks, each $\log(D)$ bits long). Say these numbers are $a_1, a_2, \ldots, a_\ell$.
- Consider the following walk on G: let $x^{(0)} = x$. For $i = 1, 2, \ldots, \ell$, get from $x^{(i-1)}$ to $x^{(i)}$ by choosing the a_i'th neighbor of $x^{(i-1)}$.
- output $x^{(\ell)} \in V$, which we treat as an element of $\{0,1\}^n$.

That is, we use the source x to tell us where to start a walk, we use the seed s to tell us how to take a random walk (notice that if $s \sim U_d$ is uniform, then we really are taking a random walk), and after we've walked ℓ steps, we output whatever vertex we happen to be on.

In the rest of this section, you'll show that Ext is a (k, ε) extractor.

Group Work

Important: as you make progress on the question(s), one person in each room should record your progress on http://PollEv.com/cs265**.**

1. Let σ be the "vectorized" pmf of X (as in the warm-up). Explain why the distri-

bution of $\mathsf{Ext}(X, U_d)$ is given by $A^\ell \cdot \sigma$. (Recall that A is the normalized adjacency matrix of G).

At this point, fill out the pollEv!

2. Let $\pi = \frac{1}{N}\mathbf{1}$ be the vector that corresponds to the uniform distribution. Explain why

$$\|U_n - \mathsf{Ext}(X, U_d)\| = \|\pi - A^\ell \cdot \sigma\| \leq \frac{\sqrt{N}}{2}\lambda(G)^\ell \|\pi - \sigma\|_2.$$

***Hint**: Mimic a computation that we did in the Expanders minilecture to show that random walks mix quickly when $\lambda(G)$ is small.*

At this point, fill out the pollEv!

3. Argue that $\|\pi - \sigma\|_2 \leq 2 \cdot 2^{-k/2}$.

***Hint**: Use the warm-up!*

4. Conclude that $\|U_n - \mathsf{Ext}(X, U_d)\| \leq \varepsilon$ and thus Ext is a (k, ε) extractor.

At this point, fill out the pollEv!

5 Using extractors for derandomization (if time)

If we still have time left, let's make rigorous something that we waved our hands about in the extractors minilecture, which is that as long as the seed length is small enough, we can essentially ignore it.

Let Ext be the extractor that you built in the previous section, so it's a (k, ε)-extractor with input length n and output length n, and seed length $O(n - k + \log(1/\varepsilon))$. Suppose that $k = n - \log n$, so the seed length is $d = O(\log(n/\varepsilon))$.

Group Work

Important: as you make progress on the question(s), one person in each room should record your progress on http://PollEv.com/cs265**.**

1. Suppose that $\mathcal{A}(y; r)$ is an (efficient) randomized algorithm that takes input y and random coins r, and outputs 0 or 1. (That is, we think of $\mathcal{A}$ as being a deterministic function of y and r, and all of the random coins that A will ever flip are encoded in r; say that $\mathcal{A}$ is trying to compute $f(y)$ for some 0/1-valued function f).

   Let $\gamma > 0$ be some small constant. Suppose that for any input y, $\mathcal{A}(y; r)$ is correct with probability $1 - \gamma$, and incorrect with probability γ:

$$\Pr_{r \sim U_n}[\mathcal{A}(y; r) = f(y)] = 1 - \gamma.$$

   With $k = n - \log n$ as above, show that for all inputs y, there is a k-source X so

that

$$\Pr_X[\mathcal{A}(y; X) \neq f(y)] = 1.$$

That is, $\mathcal{A}$ fails with probability 1 when it tries to use the k-source X for its random bits. (Here, it's okay to assume that n is sufficiently large, while γ stays constant.)

2. Now suppose we have access to both a k-source X and also to the extractor described above. Consider the following algorithm:

   def $\mathcal{A}'(y)$:

   - Draw $x \sim X$
   - For all $s \in \{0,1\}^d$:
     - Run $\mathcal{A}(y; \mathsf{Ext}(x, s))$
   - Output the majority vote over all runs of $\mathcal{A}$.

   First, convince yourself that this is efficient if $d = O(\log(n/\varepsilon))$ (which is the case if we use our extractor above with $k = n - \log n$).

   What is the probability that $\mathcal{A}'(y) = f(y)$? You can try to answer this now without help, or else you can work through the following parts:

   (a) Fix y. Let $B = \{r \in \{0,1\}^n : \mathcal{A}(y, r) \neq f(y)\}$. Argue that $\Pr_{r \sim U_n}[r \in B] = \gamma$.

   (b) Show that $\Pr_{X,U_d}[\mathsf{Ext}(X, U_d) \in B] \leq \gamma + \varepsilon$.
   ***Hint**: Use the fact that $\|U_n - \mathsf{Ext}(X, U_d)\| \leq \varepsilon$, and also use the definition of total variation distance that $\|W - Z\| = \max_{\mathcal{E}} |\Pr_W[\mathcal{E}] - \Pr_Z[\mathcal{E}]|$, where the maximum is over all events $\mathcal{E}$.*

   (c) Conclude that

   $$\mathbb{E}_{x \sim X}\left[\Pr_{U_d}[\mathsf{Ext}(x, U_d) \in B]\right] \geq \gamma + \varepsilon$$

   and apply Markov's inequality to bound

   $$\Pr_{x \sim X}\left[\Pr[\mathsf{Ext}(x, U_d) \in B] \geq \frac{1}{2}\right] \leq \text{----}$$

   (d) Argue that the thing that you put in the blank above can also go in the blank below:

   $$\Pr_{x \sim X}[\mathcal{A}(y, \mathsf{Ext}(x, U_d)) \neq f(y)] \leq \text{-----}$$

   At this point, please fill out the pollEverywhere!