CS265, Fall 2025

Class 11: Agenda and Questions

1 Announcements

- HW4 due Friday!
- HW5 out now (due 2/28)
- No class Monday (Presidents' Day)

2 Warm-Up

Group Work

In class, we said that the "second moment method" was to using the fact that

$$\Pr[X = 0] \leq \frac{\mathrm{Var}[X]}{(\mathbb{E}X)^2}$$

for all real-valued random variables X. There's another version (we'll prove a stronger form of this on HW5) that says that for any non-negative random variable X,

$$\Pr[X > 0] \geq \frac{(\mathbb{E}X)^2}{\mathbb{E}[X^2]}.$$

Is one of these stronger than the other? Are they incomparable?

Group Work: Solutions

For non-negative random variables, the second is stronger. To see this, notice that the second one says that

$$\Pr[X = 0] = 1 - \Pr[X > 0] \leq \frac{\mathbb{E}[X^2] - (\mathbb{E}X)^2}{\mathbb{E}[X^2]} = \frac{\mathrm{Var}(X)}{\mathbb{E}[X^2]}.$$

This looks like the first one, except that the denominator is $\mathbb{E}[X^2]$, which is at least as big as $(\mathbb{E}X)^2$.

On the other hand, the first one works for all random variables, not necessarily non-negative.

3 Recap/Questions?

Any questions from the minilectures and/or the quiz (second moment method and LLL)?

4 Practice with the LLL

Recall the *k-SAT* problem. There are n variables $x_1, \ldots, x_n$. We consider clauses that looks like $(x_{i_1} \vee x_{i_2} \vee \overline{x_{i_3}} \vee \cdots \vee x_{i_k})$; that is, a clause is the OR of k literals.

Group Work

Let φ be a k-CNF formula, so that each clause has k distinct variables in it. Apply the LLL to get a statement like the following:

Suppose that each variable is in at most t clauses of φ. Then φ is satisfiable.

(You should try to get t to be as large as possible. It's not hard to see that the statement above is true if, say, $t = 1$, but you should get a value of t that grows with k.)

Group Work: Solutions

We claim that if each variable is in at most $t \leq 2^{k-2}k$ clauses, then the formula is satisfiable. To see this, consider a uniformly random assignment to $x_1, \ldots x_n$ (ie setting each x_i to be TRUE/FALSE independently with probability 1/2). Define events $A_1, \ldots, A_m$ where A_i is the indicator random variable of the ith clause NOT being satisfied. For a clause with k variables to not be satisfied, all k variables must take the "bad" assignment, and hence:

$$\Pr[A_i] = 1/2^k.$$

To apply the LLL, we now need to reason about the dependencies. To that end, we claim that A_i is mutually independent of the set of clauses whose variable are disjoint from the variables in clause i, namely the set

$$S_i = \{A_j : vbl(clause_i) \cap vbl(clause_j) = \emptyset\}.$$

Indeed, no matter the assignment to variables that occur in the clauses in S_i, since none appear in the ith clause, they can't alter the probability of A_i. Now we simply count up the number of events not in the set S_i: namely

$$[\#j \text{ such that } vbl(clause_i) \cap vbl(clause_j) \neq \emptyset] \leq kt,$$

since there are k variables in clause i, and each of them is in at most t other clauses.

To conclude, each A_i is mutually independent of all but $d = kt$ other events, and hence by the LLL with $d = kt$ and $p = 2^{-k}$, we have that $\Pr[\cap_i not(A_i)] \geq (1-2p)^m > 0$ provided $dp \leq 1/4$, hence we want $kt \cdot 2^{-k} \leq 1/4$, which implies that want $t \leq 2^{k-2}/k$.

To put some concrete numbers in here, if $k = 10$, then as long as each variable appears in at most $2^{10-8}/8 = 25.6$ clauses, then the formula is always satisfiable, no matter the number of variables of clauses!!! Of course, now the big question on your mind should be "*Its great that we know such formulas are satisfiable, but how do we FIND a satisfying assignment efficiently?*" We'll get to this in the next set of minilectures, on the "Constructive LLL"!!

4.1 More Practice with LLL and Mutual Independence

Here's an example where the mutual independence requirement is a bit trickier to think about. Consider a set of m equations over variables $x_1, \ldots, x_n$:

$$\begin{aligned}
\sum_{j=1}^{n} a_j^{(1)} x_j &\equiv b^{(1)} \mod 17 \\
\sum_{j=1}^{n} a_j^{(2)} x_j &\equiv b^{(2)} \mod 17 \\
&\vdots \\
\sum_{j=1}^{n} a_j^{(m)} x_j &\equiv b^{(m)} \mod 17
\end{aligned}$$

where:

- For all $j = 1, \ldots, n$ and all $r = 1, \ldots, m$, the coefficients $a_j^{(r)} \in \{0, 1, 2, \ldots, 16\}$ are not all zero; and
- for all $r = 1, \ldots, m$, $b^{(r)} \in \{0, 1, \ldots, 16\}$.

Suppose that each variable x_j appears in at most 4 of the m equations. (That is, for each j, $a_j^{(r)} = 0$ for all but four values of r.)

Group Work

With the setup above, prove that there exists an assignment to the variables such that *none* of the equations are satisfied.

***Hint**: Recall that because 17 is prime, for any $a \in \{1, \ldots, 16\}$ and any $b \in \{0, \ldots, 16\}$, the equation $ax \equiv b \mod 17$ has a unique solution for $x \in \{0, \ldots, 16\}$.*

***Hint**: It might be helpful to go back to the definition of mutual independence when arguing*

about the value of d when applying the LLL. Remember that A is mutually independent of events $\{B_1, \ldots, B_\ell\}$, if for any set $J \subseteq \{1, 2, \ldots, \ell\}$, $\Pr[A| \cap_{j \in J} B_j] = \Pr[A]$.

Group Work: Solutions

Let's assign each x_i a random value in $\{0, 1, \ldots, 16\}$. We will define our bad events as follows: Let A_r be the event that the r'th equation is satisfied.

The probability that this occurs is $1/17$. Indeed, by assumption there is some j so that $a_j^{(r)} \neq 0$. Conditioned on the values of x_i for $i \neq j$, the r'th equation reads:

$$a_j^{(r)} x_j \equiv b^{(r)} - \sum_{i \neq j} a_i^{(r)} x_i \mod 17$$

where we've moved everything deterministic (including the stuff we've conditioned on) to the right hand side of the equation. By the hint, there is exactly one value of $x_j \in \{0, \ldots, 16\}$ that will satisfy this equation, and so the probability that we hit that x_j is exactly $1/17$. Since this holds no matter what values we've conditioned on for the $x_i, i \neq j$, the probability that A_r occurs is also $1/17$.[a]

Now we need to determine the "d" parameter in the LLL.

Let's first try the sort of argument we tried for the k-SAT example above. (This won't work!) Each equation contains at most n variables, and each variable occurs in at most 4 other variables, so there are at most $4n$ other equations that share any variable with the r'th equation. Clearly any equation that doesn't share any variables is independent, so we can take $d = 4n$. That's correct, but this is not what we wanted! We'd get $dp = 4n/17$ which will be much larger than $1/4$ for any $n \geq 2$.

Instead, let's dig in a bit to the definition of mutual independence. For a given r, we want to show that there is some set $S \subseteq [m]$ of size at most d so that for any $J \subseteq [m] \setminus S$, $\Pr[A_r| \cap_{\ell \in J} A_\ell] = \Pr[A_r]$. Suppose that j^* is such that $a_{j^*}^{(r)} \neq 0$. (This j^* exists by assumption). By assumption, there are at most three other equations so that x_{j^*} appears in them. Let S be this set of at most three r's (so we are aiming for $d = 3$). Now consider any set $J \subseteq [m] \setminus S$. We'll first write it down a bit less formally, and then write it down really formally just to be sure we know what's going on.

$$\Pr[A_r| \cap_{\ell \in J} A_\ell].$$

Now A_r occurs if $a_{j^*}^{(r)} x_{j^*} + \sum_{j \neq j^*} a_j^{(r)} x_j = b^{(r)}$. However, by construction x_{j^*} doesn't appear in any of the equations that determine the events A_ℓ for $\ell \in J$. So, even after conditioning on $\cap_{\ell \in J} A_\ell$, x_{j^*} is still uniformly random. But this means that $a_{j^*}^{(r)} x_{j^*}$ is uniformly random, which means that the conditional probability of A_r is $1/17$. But

that's the same as the unconditional probability of A_r occurring! Thus,

$$\Pr[A_r | \cap_{\ell \in J} A_\ell] = \Pr[A_r],$$

which is what we wanted to show.

To do this more formally, just as we did above, we'll condition on all of the values of x_j other than x_{j^*}. (That is, set $x_j = y_j$ for some arbitrary y_j for all $j \neq j^*$). Then for all $\ell \in J$, A_ℓ is a deterministic event, since all of the variables that appear in the ℓ'th equation have already been fixed. Thus:

$$\Pr[A_r | \cap_{\ell \in J} A_\ell, x_j = y_j \forall j \neq j^*] = \Pr[A_r | x_j = y_j \forall j \neq j^*] = 1/17,$$

where the last equality follows from the same reasoning we used to bound $\Pr[A_r] = 1/17$. Now we have

$$\begin{aligned}
\Pr[A_r | \cap_{\ell \in J} A_\ell] &= \sum_{\vec{y}} \Pr[A_r | \cap_{\ell \in J} A_\ell, x_j = y_j \forall j \neq j^*] \Pr[x_j = y_j \forall j \neq j^*] \\
&= \sum_{\vec{y}} \frac{1}{17} \Pr[x_j = y_j \forall j \neq j^*] \\
&= 1/17 \\
&= \Pr[A_r].
\end{aligned}$$

where above the sum is over all values $y_j \in \{0, \ldots, 16\}$ for all $j \neq j^*$.

Thus, $d \leq 3$, so we have

$$dp = 3 \cdot \frac{1}{17} \leq 1/4$$

and we can apply the first version of the LLL from the lecture notes.

[a] Formally, we'd write

$$\Pr[A_r] = \sum_{y_i : i \neq j} \Pr[x_i = y_i \forall i \neq j] \cdot \Pr[A_r | x_i = y_i \forall i \neq j].$$

We just argued that each $\Pr[A_r | x_i = y_i \forall i \neq j] = 1/17$, so this says

$$\Pr[A_r] = \sum_{y_i : i \neq j} \Pr[x_i = y_i \forall i \neq j] \cdot 1/17 = 1/17.$$

5 (If time) Practice with the second moment method

In a graph $G = (V, E)$, say that a vertex v is **isolated** if it has no neighboring vertices.

Group Work

Let $G \sim G_{n,p}$ be a random graph where each edge is present independently with probability p, where $p = \frac{c \ln n}{n}$ for some constant $0 < c < 1$.

1. Use the Second Moment Method to show that, with probability at least $1 - o(1)$, there is some isolated vertex in G.

   For this exercise, feel free to use the approximation $e^{-x} \approx 1 - x$ when x is small as an equality without worrying about it.

   Feel free to use either the second moment method we saw in the mini-lecture videos, or the alternate form from the warm-up. (Either will work).

   ***Hint**: Consider the random variable X that is the number of isolated vertices in G.*

   ***Hint**: When computing the variance of X, you may want to consider the following question: given two distinct vertices u, v of G, what is the probability that* both *u and v are isolated?*

Group Work: Solutions

We first compute $\mathbb{E}X$:

$$\begin{aligned}
\mathbb{E}X &= \sum_{v \in V} \Pr[v \text{ is isolated}] \\
&= \sum_{v \in V} (1-p)^{n-1} \\
&\approx \sum_{v \in V} \exp(-p(n-1)) \\
&= n \exp(-c \ln n (1 - 1/n)) \\
&= n^{1-c(1-1/n)} \\
&= n^{1-c+c/n}
\end{aligned}$$

Next, let's compute $\mathbb{E}[X^2]$.

$$\begin{aligned}
\mathbb{E}X^2 &= \sum_{v,w \in V} \Pr[v \text{ isolated AND } w \text{ isolated}] \\
&= \sum_{v \in V} \Pr[v \text{ isolated}] + \sum_{v \neq w \in V} \Pr[v \text{ and } w \text{ both isolated}]
\end{aligned}$$

where above we have broken up the sum into the terms where $v = w$ and the terms where $v \neq w$. Since there are $2n - 1$ edges total coming out of two vertices $v \neq w$, we have

$$\Pr[v, w \text{ both isolated}] = (1-p)^{2n-1} \approx \exp(-p(2n-1)) = n^{-2c+c/n}$$

with a similar derivation as above. Then we can continue our computation:

$$\mathbb{E}X^2 \leq n \cdot n^{-c(1-1/n)} + n^2 \cdot n^{-2c+c/n} = n^{1-c+c/n} + n^{2(1-c)+c/n}.$$

Now we can construct the expression in the 2nd moment method:

$$\begin{aligned}
\Pr[X = 0] &\leq \frac{\mathbb{E}[X^2] - (\mathbb{E}X)^2}{(\mathbb{E}X)^2} \\
&= \frac{n^{1-c+c/n} + n^{2(1-c)+c/n} - n^{2(1-c)+2c/n}}{n^{2(1-c)+2c/n}} \\
&\leq \frac{n^{1-c+c/n}}{n^{2(1-c)+2c/n}} \\
&= \frac{1}{n^{(1-c)+c/n}} \\
&= 1/\text{poly}(n) = o(1)
\end{aligned}$$

since $c \in (0, 1)$. Thus, for large enough n, this probability is really small. We conclude that with probability at least $1 - o(1)$, there is at least one isolated vertex.