Lecture 18

Prof. Nick Harvey *University of British Columbia*

Today's lecture introduces a completely new topic: the Lovász Local Lemma (LLL). This is an important method for analyzing events that are not independent, but have some restricted sort of dependencies. It is not as widely applicable as many of the other the techniques we have seen so far, but from time to time one encounters scenarios in which the LLL is the only technique that works.

1 The Lovász Local Lemma

Very often when designing randomized algorithms, we create a discrete probability space in which there are "bad events" $\mathcal{E}_1, \ldots, \mathcal{E}_n$ that we do not want to occur. For example, in the congestion minimization problem, $\mathcal{E}_i$ could be the event that edge i has too much congestion. Often our analysis aims to show that we can avoid these bad events with high probability, i.e., $\Pr\left[\bigcap_i \overline{\mathcal{E}_i}\right] \gg 0$.

Today let's consider the weaker goal of showing that $\Pr\left[\bigcap_i \overline{\mathcal{E}_i}\right] > 0$. There are two cases in which this goal is particularly simple.

- **Mutually independent events.** Suppose that the events are mutually independent. Then
$$\Pr\left[\textstyle\bigcap_i \overline{\mathcal{E}_i}\right] \;=\; \prod_i \Pr\left[\overline{\mathcal{E}_i}\right] \;>\; 0,$$
assuming only that $\Pr[\mathcal{E}_i] < 1$ for every i.
- **Union bound works.** Suppose that $\sum_i \Pr[\mathcal{E}_i] < 1$. Then, by a union bound,
$$\Pr\left[\textstyle\bigcap_i \overline{\mathcal{E}_i}\right] \;=\; 1 - \Pr\left[\textstyle\bigcup_i \mathcal{E}_i\right] \;\geq\; 1 - \sum_i \Pr[\mathcal{E}_i] \;>\; 0.$$

If neither of these scenarios applies, then there are few general-purpose techniques that we can try. The Lovász Local Lemma (LLL) is one of the few, and it has had spectacular applications for many problems.

Roughly speaking, the LLL is applicable in scenarios where the $\mathcal{E}_i$'s are not mutually independent, but they can have some sort of limited dependencies. Formally, a ***dependency graph*** for events $\mathcal{E}_1, \ldots, \mathcal{E}_n$ is defined as follows. The vertex set is $\{1, \ldots, n\}$. The neighbors of vertex i (excluding i itself) are denoted $\Gamma(i)$, and we also define $\Gamma^+(i) = \Gamma(i) \cup \{i\}$. The event $\mathcal{E}_i$ must be independent from the joint distribution on $\{\, \mathcal{E}_j \,:\, j \notin \Gamma^+(i) \,\}$, the events that are not neighbors of i.

This last condition means that
$$\Pr[\mathcal{E}_i] \;=\; \Pr\left[\mathcal{E}_i \;\middle|\; \textstyle\bigcap_{j \in J} \mathcal{E}_j\right] \qquad \text{for all } J \subseteq [n] \setminus \Gamma^+(i).$$

So, regardless of whether some of the events outside $\Gamma^+(i)$ occur, the probability of $\mathcal{E}_i$ occurring is unaffected.

Theorem 1 (The Symmetric LLL) *Suppose that there is a dependency graph of maximum degree d. If $\Pr[\mathcal{E}_i] \leq p$ for every i and*
$$pe(d+1) \;\leq\; 1 \tag{SLL}$$
then $\Pr\left[\bigcap_{i=1}^n \overline{\mathcal{E}_i}\right] \geq \left(\frac{d}{d+1}\right)^n > 0$.

2 Application: k-SAT

Instead, we will illustrate the LLL by considering a concrete application of it in showing satisfiability of k-CNF Boolean formulas. Recall that a k-CNF formula is a Boolean formula, involving any finite number of variables, where the formula is a *conjunction* ("and") of any number of clauses, each of which is a *disjunction* ("or") of exactly k *literals* (a variable or its negation). Let us assume that, for each clause, the variables appearing in it are all distinct.

For example, here is a 3-CNF formula with three variables and eight clauses.

$$\begin{aligned}\phi(a,b,c) \;\; = \;\; & (a \cup b \cup c) \cap (a \cup b \cup \overline{c}) \cap (a \cup \overline{b} \cup c) \cap (a \cup \overline{b} \cup \overline{c}) \cap \\ & (\overline{a} \cup b \cup c) \cap (\overline{a} \cup b \cup \overline{c}) \cap (\overline{a} \cup \overline{b} \cup c) \cap (\overline{a} \cup \overline{b} \cup \overline{c})\end{aligned}$$

This formula is obviously unsatisfiable. One can easily generalize this construction to get an unsatisfiable k-CNF formula with k variables and 2^k clauses. Our next theorem says: the reason this formula is unsatisfiable is that we allowed each variable to appear in too many clauses.

Theorem 2 *Let ϕ be a k-CNF formula where each variable appears in at most $2^k/ek$ clauses. Then ϕ is satisfiable.*

Proof: Consider the probability space in which each variable is independently set to true or false with equal probability. A clause is not satisfied if every literal appearing in that clause is false. (For example, $\overline{a} \cup b \cup \overline{c}$ is unsatisfied if a is true, b is false, and c is true.) This happens with probability 2^{-k}, since the clause involves k distinct variables.

Let $\mathcal{E}_i$ be the event that the i^{th} clause is unsatisfied. We have just argued that

$$\Pr[\,\mathcal{E}_i\,] \;\leq\; 2^{-k} \;=:\; p.$$

Consider the graph defined on $[n]$ in which there is an edge $\{i,j\}$ if some variable appears in both clause i and clause j. It is easy to see that this is a dependency graph: whether clause i is satisfied is independent from the clauses sharing no variables with clause i.

Each variable in clause i appears in at most $2^k/ek - 1$ other clauses. The number of neighbors of clause i is at most k times larger, since it contains k variables. That is,

$$|\Gamma(i)| \;\leq\; 2^k/e - 1 \;=:\; d.$$

Since $pe(d+1) \leq 1$, condition (SLL) is satisfied and $\Pr\left[\,\bigcap_{i=1}^n \overline{\mathcal{E}_i}\,\right] > 0$. This shows that ϕ is satisfiable. $\Box$

3 Symmetric LLL: a proof sketch

In this section we give a sketchy proof of the symmetric LLL that tries to explain the sequence of ideas that leads to the proof. If this sketch is not to your taste, a correct and concise proof is given in Section 4.1. Instead of assuming (SLL), it will be convenient to assume the strengthened hypothesis

$$4pd \;\leq\; 1. \tag{1}$$

To make the notation more meaningful, let B_i denote the "bad" event $\mathcal{E}_i$ and let G_i be the "good" event $\overline{B_i}$. Note that independence from B_i is equivalent to independence from G_i. For any set $S \subseteq [n]$, let $G_S = \cap_{i \in S} G_i$.

The objective of the proof is to show that $\Pr\left[\, G_{[n]} \,\right] > 0$, i.e., with positive probability, all good events occur simultaneously. Note that the union bound gives the easy lower bound

$$\Pr[\, G_S \,] \;=\; 1 - \Pr[\, \cup_{i \in S} B_i \,] \;\geq\; 1 - \sum_{i \in S} \Pr[\, B_i \,] \;\geq\; 1 - |S|p, \tag{2}$$

but in our scenario, this is too weak — we could have $|S|p \gg 1$. Instead, a natural idea is to use the "chain rule" to break apart this large conjunction:

$$\Pr\left[\, G_{[n]} \,\right] \;=\; \prod_{k=1}^{n} \Pr\left[\, G_k \;\mid\; G_{[k-1]} \,\right]. \tag{3}$$

We just need to show that each factor in this product is strictly positive.

Idea 1: Rather than showing that each factor is positive, it turns out to be convenient to negate the event and prove an upper bound. Specifically, we want to show that

$$\Pr[\, B_k \;\mid\; G_S \,] \;\leq\; \alpha p \qquad \forall S \subseteq [n] \tag{Hope}$$

for some α to be chosen later. In order for this conditional probability to be well defined, we need that $\Pr[\, G_S \,] > 0$. Let's ignore that for now as our whole purpose is to prove the stronger fact that $\Pr\left[\, G_{[n]} \,\right] > 0$. Trivially,

$$\Pr[\, B_k \;\mid\; G_S \,] \;=\; \Pr\left[\, B_k \;\mid\; G_{S \cap \Gamma(k)} \cap G_{S \setminus \Gamma(k)} \,\right].$$

Idea 2: We would like to somehow use the fact that B_k is independent of $G_{S \setminus \Gamma(k)}$ (due to the dependency graph). But it is difficult to use that fact due to the additional conditioning on $G_{S \cap \Gamma(k)}$. So as a first step we use the definition of conditional probability to write

$$\Pr[\, B_k \;\mid\; G_S \,] \;=\; \frac{\Pr\left[\, B_k \cap G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right]}{\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right]}.$$

Idea 3: The next idea is to drop the "$\cap\, G_{S \cap \Gamma(k)}$" event yielding the following upper bound. We would hope that still this gives a good bound since $\Gamma(k)$ is small and $G_{S \cap \Gamma(k)}$ is a very likely event.

$$\Pr[\, B_k \;\mid\; G_S \,] \;\leq\; \frac{\Pr\left[\, B_k \;\mid\; G_{S \setminus \Gamma(k)} \,\right]}{\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right]} \;=\; \frac{\Pr[\, B_k \,]}{\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right]} \;\leq\; \frac{p}{\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right]}$$

The equality here uses our second idea, that B_k is independent of the events outside of $\Gamma(k)$.

Now, to prove (Hope), it suffices to prove that $\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right] \geq 1/\alpha$. The good news is that the conjunction $G_{S \cap \Gamma(k)}$ involves few events, so we are in good shape to use the union bound as in (2):

$$\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right] \;\geq\; 1 - \sum_{i \in S \cap \Gamma(k)} \Pr\left[\, B_i \;\mid\; G_{S \setminus \Gamma(k)} \,\right]. \tag{4}$$

If $S \cap \Gamma(k) = \emptyset$ then this quantity is 1 as the sum is empty. Otherwise, $|S \setminus \Gamma(k)| < |S|$, so we can use induction on $|S|$. We have

$$\begin{aligned}
\Pr\left[\, G_{S \cap \Gamma(k)} \;\mid\; G_{S \setminus \Gamma(k)} \,\right] \;&\geq\; 1 - \sum_{i \in \Gamma(k)} \Pr\left[\, B_i \;\mid\; G_{S \setminus \Gamma(k)} \,\right] && \text{(by (4))} \\
&\geq\; 1 - d\alpha p && \text{(inductively using (Hope))} \\
&\geq\; 1 - \alpha/4 && \text{(by the strengthened hypothesis (1))} \\
&=\; 1/\alpha
\end{aligned}$$

if we now choose $\alpha = 2$. This shows that (Hope) is satisfied.

4 General LLL

The Symmetric Local Lemma is useful, but often somewhat restrictive. It works best when all events are equally probable, and when all neighborhood sizes are the same. We might want to consider scenarios in which some events are quite likely and some are quite rare. The likely events would need to depend on few other events, and the rare events could perhaps depend on many other events. There is a general form of the local lemma that can handle such scenarios.

Theorem 3 (General LLL) *Suppose that there is a dependency graph and an $x \in (0,1)^n$ satisfying*

$$\Pr[\,\mathcal{E}_i\,] \;\leq\; x_i \cdot \prod_{j \in \Gamma(i)} (1 - x_j) \qquad \forall i. \tag{GLL}$$

Then $\Pr\left[\,\bigcap_{i=1}^n \overline{\mathcal{E}_i}\,\right] \geq \prod_{i=1}^n (1 - x_i) > 0$.

This form of the local lemma is confusing at first because it's not obvious what these x_i values should be. In order to satisfy (GLL), on the right-hand side we want x_i to be big and each x_j to be small. Due to that tension, care is needed in finding the right x_i.

4.1 A Concise Proof of Theorem 3

We simultaneously prove by induction on $|S|$ that, for all $S \subseteq [n]$,

$$\Pr[\,G_S\,] > 0 \tag{5a}$$
$$\Pr[\,B_k \mid G_S\,] \leq x_k. \tag{5b}$$

In the case $S = \emptyset$, (5a) is trivial and (5b) follows directly from (GLL).

By relabeling, we may assume that $S = \{1, \ldots, s\}$, so

$$\Pr[\,G_S\,] \;=\; \prod_{j=1}^{s} \Pr\left[\,G_j \mid G_{\{1,\ldots,j-1\}}\,\right] \;=\; \prod_{j=1}^{s} \left(1 - \Pr\left[\,B_j \mid G_{\{1,\ldots,j-1\}}\,\right]\right) \;\geq\; \prod_{j=1}^{s} (1 - x_j),$$

where we inductively use (5a) to ensure that the conditional probabilities are well-defined, and we inductively use (5b) to provide the inequality. This proves (5a).

Next consider (5b). If $S \cap \Gamma(k) = \emptyset$ then $\Pr[\,B_k \mid G_S\,] = \Pr[\,B_k\,]$, so (5b) follows directly from (GLL). Otherwise, relabeling so that $S \cap \Gamma(k) = \{1, \ldots, t\}$, we have

$$\begin{aligned}
\Pr[\,B_k \mid G_S\,] \;&=\; \Pr\left[\,B_k \mid G_{S\cap\Gamma(k)} \cap G_{S\setminus\Gamma(k)}\,\right] \;=\; \frac{\Pr\left[\,B_k \cap G_{S\cap\Gamma(k)} \mid G_{S\setminus\Gamma(k)}\,\right]}{\Pr\left[\,G_{S\cap\Gamma(k)} \mid G_{S\setminus\Gamma(k)}\,\right]} \\
&\leq\; \frac{\Pr\left[\,B_k \mid G_{S\setminus\Gamma(k)}\,\right]}{\Pr\left[\,G_{S\cap\Gamma(k)} \mid G_{S\setminus\Gamma(k)}\,\right]} \;=\; \frac{\Pr[\,B_k\,]}{\prod_{j=1}^{t} \Pr\left[\,G_j \mid G_{(S\setminus\Gamma(k))\cup\{1,\ldots,j-1\}}\,\right]} \\
&\leq\; \frac{x_k \cdot \prod_{i\in\Gamma(k)}(1 - x_i)}{\prod_{j=1}^{t}\left(1 - \Pr\left[\,B_j \mid G_{(S\setminus\Gamma(k))\cup\{1,\ldots,j-1\}}\,\right]\right)} \;\leq\; \frac{x_k \cdot \prod_{i\in\Gamma(k)}(1 - x_i)}{\prod_{j=1}^{t}\left(1 - x_j\right)}.
\end{aligned}$$

The second inequality uses (GLL) and the third uses induction. The last expression is clearly at most x_k, proving (5b).

4.2 General implies Symmetric

Finally, let us conclude by noting that Theorem 3 implies Theorem 1. To see this, consider an instance satisfying (SLL). Set $x_i = 1/(d+1)$. The RHS of (GLL) is

$$
\begin{aligned}
x_i \cdot \prod_{j \in \Gamma(i)} (1 - x_i) \;&\geq\; \frac{1}{d+1} \cdot \big(1 - \frac{1}{d+1}\big)^d \\
&\geq\; \frac{1}{e(d+1)} \qquad \text{(by calculus)} \\
&\geq\; p \qquad \text{(by (SLL))}.
\end{aligned}
$$

This shows that (GLL) is satisfied, so Theorem 3 implies

$$
\Pr\left[\bigcap_{i=1}^{n} \overline{\mathcal{E}_i}\right] \;\geq\; \prod_{i=1}^{n} (1 - x_i) \;=\; \prod_{i=1}^{n} \Big(\frac{d}{d+1}\Big)^n,
$$

which is the conclusion of Theorem 1.