

Lecture 3. The Lopsided Lovász Local Lemma

Next, we turn to a generalized form of the LLL, known as the Lopsided Lovász Local Lemma (LLLL).

3.1 The Lopsided LLL

We note that in the proof of the LLL we used the equation

$$\mathbb{P} \left[\mathcal{E}_a \cap \bigcap_{i \in S \setminus \Gamma^+(a)} \overline{\mathcal{E}_i} \right] = \mathbb{P}[\mathcal{E}_a] \cdot \mathbb{P} \left[\bigcap_{i \in S \setminus \Gamma^+(a)} \overline{\mathcal{E}_i} \right],$$

coming from \mathcal{E}_a being independent of the events in $S \setminus \Gamma^+(a)$ (see the last equation in the chain of inequalities below (2.4)). However, in the proof we only need an inequality

$$\mathbb{P} \left[\mathcal{E}_a \cap \bigcap_{i \in S \setminus \Gamma^+(a)} \overline{\mathcal{E}_i} \right] \leq \mathbb{P}[\mathcal{E}_a] \cdot \mathbb{P} \left[\bigcap_{i \in S \setminus \Gamma^+(a)} \overline{\mathcal{E}_i} \right].$$

Note that this inequality is equivalent to

$$\mathbb{P} \left[\mathcal{E}_a \cap \bigcup_{i \in S \setminus \Gamma^+(a)} \mathcal{E}_i \right] \geq \mathbb{P}[\mathcal{E}_a] \cdot \mathbb{P} \left[\bigcup_{i \in S \setminus \Gamma^+(a)} \mathcal{E}_i \right]$$

(since both the LHS and the RHS of these two inequalities add up to $\mathbb{P}[\mathcal{E}_a]$).

This motivates the notion of a negative dependency graph — a graph where each event \mathcal{E}_a is positively correlated with its non-neighboring events in the following sense.

Definition 3.1 *G is a negative dependency (di)graph if for every event \mathcal{E}_a and every subset $J \subset V \setminus \Gamma^+(a)$ we have*

$$\mathbb{P} \left[\mathcal{E}_a \cap \bigcup_{b \in J} \mathcal{E}_b \right] \geq \mathbb{P}[\mathcal{E}_a] \cdot \mathbb{P} \left[\bigcup_{b \in J} \mathcal{E}_b \right].$$

The proof of the LLL works the same way for a negative dependency (di)graph as for a dependency graph, as explained above. We obtain the following generalization.

Theorem 3.2 (Lopsided Lovász Local Lemma (LLLL)) *Suppose G is a negative dependency (di)graph for events $\mathcal{E}_1, \dots, \mathcal{E}_n$, and there exist $x_1, \dots, x_n \in (0, 1)$ such that for every $i = 1, \dots, n$ we have*

$$\mathbb{P}[\mathcal{E}_i] \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j).$$

Then

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

In fact, we can go one step further and assume directly that for every set of non-neighbors $J \subseteq V \setminus \Gamma^+(i)$, $\mathbb{P} \left[\mathcal{E}_i \mid \bigcap_{j \in J} \overline{\mathcal{E}_j} \right] \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j)$. Again, the proof goes through without any change. We obtain the following lemma that we call the *Extended Lopsided LLL*.

Theorem 3.3 (*Extended Lopsided LLL*) *If there exist $x_1, \dots, x_n \in (0, 1)$ such that for all $i \in V(G)$ and $J \subseteq V \setminus \Gamma^+(i)$ the following inequality holds*

$$\mathbb{P} \left[\mathcal{E}_i \mid \bigcap_{j \in J} \overline{\mathcal{E}_j} \right] \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j)$$

then

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

The Lopsided LLL is typically useful for applications on probability spaces with more structure than independent random variables. In the following, we present two such applications. (It turns out the second one in fact does not require lopsidedness, even though the probability space is non-trivial.)

3.2 Latin transversals

Definition 3.4 *A Latin square is an $n \times n$ matrix, whose entries are numbers from $[n]$ so that each row and column is a permutation of $[n]$.*

For example,

1	2	3
2	3	1
3	1	2

Definition 3.5 *Let $A \in \mathbb{Z}^{n \times n}$ be a matrix. A Latin transversal of A is a permutation $\pi \in \mathbb{S}_n$ such that $A_{i,\pi(i)} \neq A_{j,\pi(j)}$ if $i \neq j$.*

Conjecture 1 *Every Latin square of odd order has a Latin transversal.*

For example, $\pi = id$ is a Latin transversal for the following matrix.

①	2	3
2	③	1
3	1	②

The result we are going to prove in this section is the following:

Theorem 3.6 (Erdős-Spencer) *For any $A \in \mathbb{Z}^{n \times n}$ such that every integer appears as an entry of A at most $\frac{n}{4e}$ times, there is a Latin transversal.*

Proof: Let $\Omega = \mathbb{S}_n$, with the uniform probability measure. The "bad" events that need to be avoided will correspond to pairs of entries $A_{i_1, j_1} = A_{i_2, j_2}$ where $i_1 \neq i_2$ and $j_1 \neq j_2$. Therefore, define

$$\mathcal{E}_{i_1, j_1, i_2, j_2} = \{\pi \in \mathbb{S}_n \mid \pi(i_1) = j_1, \pi(i_2) = j_2\}$$

where $i_1 \neq i_2$ and $j_1 \neq j_2$ (for the moment, we ignore the matrix A). Then,

$$p := \mathbb{P}(\mathcal{E}_{i_1, j_1, i_2, j_2}) = \frac{1}{n(n-1)}.$$

Let V denote all possible 4-tuples (i_1, j_1, i_2, j_2) such that $i_1 \neq i_2$ and $j_1 \neq j_2$. A dependency graph $\mathcal{G} = (V, E)$ can be defined as follows:

$$((i_1, j_1, i_2, j_2), (i'_1, j'_1, i'_2, j'_2)) \in E(\mathcal{G})$$

iff

$$\begin{cases} \{i_1, i_2\} \cap \{i'_1, i'_2\} \neq \emptyset, \\ \text{or} \\ \{j_1, j_2\} \cap \{j'_1, j'_2\} \neq \emptyset. \end{cases}$$

Claim 3.7 \mathcal{G} is a negative dependency graph for the events $\mathcal{E}_{i_1, j_1, i_2, j_2}$.

To prove the claim, fix $a = (i_1, j_1, i_2, j_2)$ and let $J \subseteq V \setminus \Gamma^+(a)$ be a set of non-neighboring 4-tuples. We want to prove that

$$\mathbb{P}\left(\mathcal{E}_a \mid \bigcap_{b \in J} \overline{\mathcal{E}_b}\right) \leq \mathbb{P}(\mathcal{E}_{i_1, j_1, i_2, j_2}) = \frac{1}{n(n-1)}.$$

We examine only the case $a = (1, 1, 2, 2)$, since all other cases can be similarly treated. Set

$$\mathcal{S}_{i,j} = \{\pi \in \mathbb{S}_n \mid \pi(1) = i, \pi(2) = j \text{ and } \pi \in \bigcap_{b \in J} \overline{\mathcal{E}_b}\}$$

for $i \neq j$. The $\mathcal{S}_{i,j}$'s partition $\bigcap_{b \in J} \overline{\mathcal{E}_b}$ into $n(n-1)$ blocks, determined by $\pi(1)$ and $\pi(2)$. Also, if we prove

$$|\mathcal{S}_{1,2}| \leq |\mathcal{S}_{i,j}| \tag{3.1}$$

for every $i \neq j$, this will finish the proof because then

$$\mathbb{P}(\mathcal{E}_a \mid \bigcap_{b \in J} \overline{\mathcal{E}_b}) = \frac{|\mathcal{S}_{1,2}|}{\sum_{k \neq l} |\mathcal{S}_{k,l}|} \leq \frac{1}{n(n-1)}.$$

To see why (3.1) holds, consider the following map, $T : \mathcal{S}_{1,2} \rightarrow \mathcal{S}_{i,j}$, defined as

$$T(\pi) = (1, i)(2, j)\pi$$

(apply π and then switch $1 \leftrightarrow i, 2 \leftrightarrow j$). This is well defined because

1. If $\pi \in \mathcal{S}_{1,2}$ then $\pi(1) = 1$ and $\pi(2) = 2$. Therefore, $T(\pi)(1) = i$ and $T(\pi)(2) = j$, i.e., $T(\pi) \in \mathcal{S}_{i,j}$.
2. If $\pi \in \bigcap_{b \in J} \overline{\mathcal{E}_b}$ then $T(\pi) \in \bigcap_{b \in J} \overline{\mathcal{E}_b}$ as well because the events \mathcal{E}_b do not involve 1 or 2, either in the domain or in the range.

It is clear that T is injective which proves equation 3.1.

Now let's focus on the 4-tuples (i_1, j_1, i_2, j_2) such that $A_{i_1, j_1} = A_{i_2, j_2}$. Let's call this set W . To apply the Lopsided Lovász local lemma, all we need to find is d , which the maximum number of events in W that a given event $\mathcal{E}_{i_1, j_1, i_2, j_2}$ is dependent on. As we argued above, such events correspond to 4-tuples $\mathcal{E}_{i'_1, j'_1, i'_2, j'_2}$ where either $\{i_1, i_2\} \cap \{i'_1, i'_2\} \neq \emptyset$, or the same for the j 's. Therefore there at most $(4n - 4)$ choices for one of the pairs and $k - 1$ for the other pair (since there are $4n - 4$ entries in the given pairs of rows/columns, for each of them at most $k - 1$ other pairs sharing the same entry in the matrix). Then indeed

$$(d + 1) < (4n - 4)k \leq (4n - 4) \frac{n}{4e} = \frac{n(n - 1)}{e}$$

which implies that

$$(d + 1)pe \leq 1.$$

Therefore there is a Latin transversal for A .

□

3.3 Rainbow trees

Definition 3.8 *A rainbow tree is a tree with colored edges such that each edge has a different color.*

Conjecture 2 (Brualdi-Hollingsworth) *Let $n \geq 6$ be an even integer. For any proper edge coloring (incident edges get different colors) of K_n using $n - 1$ colors, there is a decomposition of $E(K_n)$ into $n/2$ rainbow spanning trees.*

This conjecture above is still open. The following recent result is known.

Theorem 3.9 (Horn '17) *There exist $\Omega(n)$ edge-disjoint rainbow trees for any proper edge coloring of K_n .*

The result that the LLL gives is the following (from [Harvey-Vondrak '15]). Interestingly, we don't even need the lopsided version here.

Theorem 3.10 *For any edge coloring of K_n where each color appears at most $n/100$ times, there exist $n/100$ edge-disjoint rainbow spanning trees.*

The following lemmata will be needed for the proof of Theorem 3.10.

Lemma 3.11 (Lu-Mohr-Székely '12) *Let T be a uniformly random spanning tree of K_n . For any two vertex-disjoint disjoint sets of edges $A, B \subset E(K_n)$ the following holds:*

$$\mathbb{P}(A \subset T \ \& \ B \subset T) = \mathbb{P}(A \subset T) \mathbb{P}(B \subset T).$$

Lemma 3.12 For a forest $F \subset K_n$ with components of f_1, f_2, \dots, f_s vertices ($\sum f_i = n$), the number of spanning trees in K_n containing F is $n^{n-2} \prod_{i=1}^s \frac{f_i}{n^{f_i-1}}$.

The above lemmata are left as an exercise for the reader. (Lemma 3.11 follows from Lemma 3.12.) We note that it is an interesting fact that the appearances of disjoint edges in a random tree are independent. This happens because we are looking at random spanning trees of a complete graph. In a general graph, the appearances of edges in a random spanning tree are negatively correlated.

Proof:(of Theorem 3.10) Set $t = n/100$ and choose $\bar{T} = \{T_1, T_2, \dots, T_t\}$, a collection of t independently random spanning trees in K_n . There are two types of bad events:

- Type 1: a color appears twice in one of the trees. Formally, for all $i \in [t]$, $e, f \in E(K_n)$ such that e and f have the same color c , define

$$\mathcal{E}_{e,f}^i = \{\bar{T} \mid e, f \in T_i\}$$

- Type 2: two trees share an edge e . For all $i, j \in [t]$ and $e \in E(K_n)$, define

$$\mathcal{E}_e^{i,j} = \{\bar{T} \mid e \in T_i \cap T_j\}$$

Using Lemma 3.12, if e and f are two disjoint edges, we get that

$$\mathbb{P}(\mathcal{E}_{e,f}^i) = \frac{2}{n^{2-1}} \frac{2}{n^{2-1}} = \frac{4}{n^2}$$

while if e and f meet at a vertex,

$$\mathbb{P}(\mathcal{E}_{e,f}^i) = \frac{3}{n^2}.$$

Also because T_i is independent of T_j :

$$\mathbb{P}(\mathcal{E}_e^{i,j}) = \frac{2}{n^{2-1}} \frac{2}{n^{2-1}} = \frac{4}{n^2}.$$

So we can set $p = \frac{4}{n^2}$.

What are the dependencies? We define a dependency graph as follows:

$$(\mathcal{E}_A^I, \mathcal{E}_B^J) \in E(\mathcal{G})$$

iff

$$I \cap J \neq \emptyset \quad \& \quad V(A) \cap V(B) \neq \emptyset.$$

(Here, I, J could be single indices or pairs, and similarly A, B could be single edges or pairs of edges, as above.) It follows from Lemma 3.11 and the independent choice of different trees that this is a dependency graph (note — lopsidedness is not needed here). The degrees in the dependency graph are bounded as follows.

1. $\mathcal{E}_{e,f}^i$ is dependent on at most $4nt$ events of type 1. Notice that it is independent of all events of the form $\mathcal{E}_{e',f'}^{i'}$ where $i \neq i'$. So we should only count the events of the form $\mathcal{E}_{e',f'}^i$. Either e' or f' have to be incident to one of e and f and there are at most $4n$ such edges. The other edge will be chosen to have the same color as the first one picked, so there are t such options at most.

2. $\mathcal{E}_{e,f}^i$ is dependent on at most $4nt$ events of type 2. Now we are looking among events of the form $\mathcal{E}_{e'}^{i,j}$. There are at most $4n$ edges incident to e, f (in order to choose e') and then there are t trees to choose one from (for the choice of j).
3. $\mathcal{E}_e^{i,j}$ is dependent on at most $4nt$ events of type 1 and at most $4nt$ events of type 2, because of similar arguments.

Note that we counted the event $\mathcal{E}_{e,f}$ itself among the dependencies. Hence, we have

$$(d+1) \leq 8nt = \frac{n^2}{12.5}$$

and $p = \frac{4}{n^2}$, so

$$ep(d+1) < 1.$$

Therefore by the LLL, there are at least $n/100$ edge-disjoint rainbow spanning trees. □