

Lecture 5: Lovasz Local Lemma

Lecturer: Shayan Oveis Gharan

04/22/2025

Scribe:

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications.*

This lecture and the next are based on Alistair Sinclair's course on randomized algorithms

Recall from earlier lectures that the probabilistic method provides a useful non-constructive strategy for proving the existence (or non-existence) of an object that satisfies some prescribed property. Generally, the argument involves selecting an object randomly from a specific set and demonstrating that it has the desired property with strictly positive probability. This in turn proves the existence of at least one such object. In most of the examples we have seen, the desired property holds not just with positive probability but actually with quite large probability, even with probability approaching 1 as $n \rightarrow \infty$. This in turn often leads to an efficient randomized algorithm for constructing such an object: we just select an object at random and with high probability it has the desired property.

For some problems, it is natural to describe the selected object in terms of a set of "bad" events $\{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$, whose occurrences render the object undesirable, while the desired property is simply the avoidance of all these bad events in the set. In such scenarios, the existence of a non-trivial lower bound on

$$\mathbb{P} \left[\bigwedge_{i=1}^n \neg \mathcal{A}_i \right]$$

is of particular interest. One immediate approach is just to use the union bound. The union bound fails if the probability of each bad is way larger than $1/n$. Of particular interest to us are the case that each \mathcal{A}_i occurs with a small constant probability.

Clearly, if all "bad" events are independent, and if the probability of each of them satisfies $\mathbb{P}[\mathcal{A}_i] \leq p$, then the probability that none of the events $\{\mathcal{A}_i\}$ occur is simply the product

$$\mathbb{P} \left[\bigwedge_{i=1}^n \neg \mathcal{A}_i \right] = \prod_{i=1}^n \mathbb{P}[\neg \mathcal{A}_i] \geq (1-p)^n \quad (5.1)$$

which is strictly positive (provided only that the trivial condition $p < 1$ holds).

Informally, the Lovász Local Lemma can be viewed as extending the above result to a more general setting, in which we allow limited dependencies among the events in question. In light of (5.1), the resulting probability that no bad event occurs will typically be exponentially small. Thus the Local Lemma tends to apply in situations where we are looking for a "needle in a haystack," so does not immediately lead to an efficient randomized algorithm. (However, see below and the next lecture for more recent developments on constructive versions of the Lemma.)

Definition 5.1. *An event \mathcal{A} is said to be mutually independent of a set of events $\{\mathcal{A}_i\}$ if for any subset S of events or their complements contained in $\{\mathcal{A}_i\}$, we have $\mathbb{P}[\mathcal{A}|S] = \mathbb{P}[\mathcal{A}]$.*

Lemma 5.2 (Lovász Local Lemma). *Let $\mathcal{A}_1, \dots, \mathcal{A}_n$ be a set of "bad" events with $\mathbb{P}[\mathcal{A}_i] \leq p < 1$ such that each event \mathcal{A}_i is mutually independent of all but at most d of the other \mathcal{A}_j 's. If $e \cdot p(d+1) \leq 1$ then*

$$\mathbb{P} \left[\bigwedge_{i=1}^n \neg \mathcal{A}_i \right] > 0.$$

Often the Lovász Local Lemma is stated with the condition $e \cdot p(d+1) \leq 1$ replaced by $4pd \leq 1$, which is slightly weaker for $d \leq 2$ but asymptotically stronger. In fact, the constant e above is asymptotically optimal.

5.1 Application 1: Satisfying Solutions to k -SAT

Theorem 5.3. *Any instance ϕ of k -SAT in which no variable appears in more than $\frac{2^{k-2}}{k}$ clauses is satisfiable.*

As a quick example, the above claim implies that for $k = 10$, any formula in which no variable appears in more than 25 clauses is satisfiable. Note that there is no restriction at all on the total number of clauses!

Proof. Suppose we have an arbitrary instance ϕ of k -SAT consisting of n clauses. Let's pick a truth assignment to the variables of ϕ uniformly at random and let A_i denote the event "clause i is not satisfied". Noting that exactly one of the 2^k possible assignments fails to satisfy any particular clause, we have

$$\forall i \in \{1, 2, \dots, n\} : \mathbb{P}[A_i] = 2^{-k} \equiv p.$$

Furthermore, we observe that each event A_i is independent of all other events A_j except those corresponding to clauses j that share at least one variable with clause i . Let d denote the largest possible number of such clauses. Clearly, since each variable is assumed to appear in at most $\frac{2^{k-2}}{k}$ clauses, we have

$$d \leq k \cdot \frac{2^{k-2}}{k} = 2^{k-2}.$$

The condition $p \leq 1/4d$ in the local lemma now becomes

$$p = \frac{1}{2^k} = \frac{1}{4 \cdot 2^{k-2}} \leq \frac{1}{4d}$$

Hence, the lemma implies that

$$\mathbb{P}\left[\bigwedge_{i=1}^n \neg A_i\right] > 0.$$

Since the probability of picking an assignment that satisfies every clause in ϕ is non-zero, we can invoke the standard argument of the probabilistic method and infer the existence of a satisfying truth assignment. \square

In the above proof, we claimed that each \mathcal{A}_i is independent of all \mathcal{A}_j for which clauses i and j do not share any variables. This is an instance of the following general principle that is frequently useful in applications of the Local Lemma:

Proposition 5.4 (Mutual Independence Principle). *Suppose that Z_1, \dots, Z_m is an underlying sequence of independent variables, and suppose that each event \mathcal{A}_i is completely determined by some subset $S_i \subseteq \{Z_1, \dots, Z_m\}$. If $S_i \cap S_j = \emptyset$ for $j = j_1, \dots, j_k$, then \mathcal{A}_i is mutually independent of $\{\mathcal{A}_{j_1}, \dots, \mathcal{A}_{j_k}\}$.*

In our above application, the underlying independent events Z_i are the assignments to the variables.

5.2 Application 2: Min Congestion Routing with Dilation

Recall that in the last lecture we discussed the min congestion problem: We have k terminal pairs, $(s_1, t_1), \dots, (s_k, t_k)$ and let \mathcal{P}_i be the set of paths from s_i to t_i for all i and $\mathcal{P} = \cup_i \mathcal{P}_i$. Furthermore, let $y_{\mathcal{P}}$ be a feasible LP solution with fractional congestion 1 (everything we say will generalize for larger congestion, but here we choose 1 for simplicity of the arguments). So, $\sum_{P \in \mathcal{P}_i} y_P = 1$ for all i .

We say the LP solution has dilation D if

$$D = \max\{|P| : P \in \mathcal{P}, y_P > 0\},$$

i.e., D is the maximum length of all paths in the support of LP. We prove the following theorem:

Theorem 5.5 (Leighton-Rao-Srivisan [LRS98]). *Given a fractional solution with dilation D , it is possible to round this to integral flow with congestion at most $O(\log D / \log \log D)$.*

Note that $D \leq n$, so the above theorem only improves over what we proved in that last lecture when $D \ll n$. The proof uses Lovász Local Lemma, so the rounding procedure was not algorithmic at the time that the paper was published but it can be made algorithmic using new developments.

We assume that the $OPT = 1$ to simplify the argument, although it naturally extends to when $OPT > 1$.

Step 1: Discretization. It turns out that the analysis is simpler if y_P is invariant over all paths with $y_P > 0$. To do that, we "discretize" the LP solution, simply by choosing an $\epsilon \ll \frac{1}{kn}$ -small enough and repeat every path $\lfloor y_P / \epsilon \rfloor$ many times. This can incur an extra loss of $k\epsilon \ll 1/n$ on the congestion of every edge which we ignore for simplicity. So, assume that $y_P = \epsilon$ for every path P where $y_P > 0$. Furthermore, since the fractional congestion of every edge is ≤ 1 , there are at most $m := 1/\epsilon$ many paths going over an edge e .

Step 2: Bad Events The crux of the proof is to define the right set of bad events. Let C be the target congestion. Here the idea is as follows: For every edge e , and every set S of C paths going through e we add an event $\mathcal{A}_{e,S}$ which occurs if all of the paths in S are chosen. So, if none of the bad events happen the congestion of every edge is at most $C - 1$ and we are done. Next, we upper bound the probability of a bad event. Fix a bad event $\mathcal{A}_{e,S}$, where $S = \{(i_1, j_1), \dots, (i_C, j_C)\}$; We use the notation (i_1, j_1) to denote the j_1 'th path of the m many paths connecting s_{i_1} to t_{i_1} in S , and write $P_{i_1, j_1} \in \mathcal{P}_{i_1}$ to denote this path.

Now, notice if S contains two paths of the same source/destination pair, then $\mathbb{P}[\mathcal{A}_{e,S}] = 0$. So, we ignore these events from any further considerations. Otherwise, if all of the C paths of S come from distinct source/destination pair, we have

$$\mathbb{P}[\mathcal{A}_{e,S}] = \frac{1}{m^{|S|}} = \frac{1}{m^C}.$$

Step 3: Dependency Set Fix two bad event $\mathcal{A}_{e,S}, \mathcal{A}_{e',S'}$ where paths in S, S' are from distinct source/destination pairs. Say $S = \{(i_1, j_1), \dots, (i_C, j_C)\}$ and $S' = \{(i'_1, j'_1), \dots, (i'_C, j'_C)\}$. Observe that if $\{i_1, \dots, i_k\} \cap \{i'_1, \dots, i'_k\} = \emptyset$ then S, S' are independent. So assume $i'_1 \in \{i_1, \dots, i_k\}$, say $i'_1 = i_1$. For how many choices of i'_2, \dots, i'_k and j'_1, \dots, j'_k do we have a bad event? First, there are m choices for j'_1 . Fix such a choice. The path $P_{i'_1, j'_1}$ uses that D many edges and each edge has at most m paths going through it. Thus the number of bad events containing i_1 is at most

$$m \cdot D \cdot \binom{m}{C-1}.$$

Since there are C options for the common index between $\mathcal{A}_{e,S}, \mathcal{A}_{e',S'}$ we have

$$|\Gamma(\mathcal{A}_{e,S})| \leq C \cdot m \cdot D \cdot \binom{m}{C-1} \leq \frac{C \cdot D \cdot m^C}{(C-1)!}.$$

So, to apply LLL we need

$$e \cdot \frac{1}{m^C} \cdot \frac{C \cdot D \cdot m^C}{(C-1)!} \leq e \cdot \frac{CD}{(C-1)!} < 1$$

The latter holds as long $C > \frac{\log D}{\log \log D}$.

5.3 Proof of the Lovász Local Lemma

First, we prove the following lemma which is the main ingredient of the proof:

Lemma 5.6. *For any subset $S \subset \{1, \dots, n\}$, and any $i \in \{1, \dots, n\}$, such that $i \notin S$ we have*

$$\mathbb{P} \left[\mathcal{A}_i \mid \bigwedge_{j \in S} \neg \mathcal{A}_j \right] \leq \frac{1}{d+1}.$$

Proof. We prove by induction on $m = |S|$. The base case, $m = 0$, holds since $\mathbb{P}[\mathcal{A}_i] \leq p \leq \frac{1}{e(d+1)} \leq \frac{1}{d+1}$. For the inductive step ($m > 0$) we first partition S into the two sets $S_1 = S \cap \Gamma_i$ and $S_2 = S \setminus S_1$, where Γ_i is the "dependency set" of \mathcal{A}_i , i.e., the set of at most d indices j such that \mathcal{A}_i is independent of all \mathcal{A}_j except for those in this set. Note that by definition $|S_1| \leq d$. Then, by Bayes rule we may write

$$\mathbb{P} \left[\mathcal{A}_i \mid \bigwedge_{j \in S} \neg \mathcal{A}_j \right] = \frac{\mathbb{P} \left[\mathcal{A}_i \wedge \bigwedge_{j \in S_1} \neg \mathcal{A}_j \mid \bigwedge_{k \in S_2} \neg \mathcal{A}_k \right]}{\mathbb{P} \left[\bigwedge_{j \in S_1} \neg \mathcal{A}_j \mid \bigwedge_{k \in S_2} \neg \mathcal{A}_k \right]}$$

To bound the numerator we write,

$$\mathbb{P} \left[\mathcal{A}_i \wedge \bigwedge_{j \in S_1} \neg \mathcal{A}_j \mid \bigwedge_{k \in S_2} \neg \mathcal{A}_k \right] \leq \mathbb{P} \left[\mathcal{A}_i \mid \bigwedge_{k \in S_2} \neg \mathcal{A}_k \right] = \mathbb{P}[\mathcal{A}_i] \leq p.$$

So, it remains to lower bound the denominator. WLOG, perhaps after renaming, assume $S_1 = \{1, 2, \dots, r\}$. We use the Bayes rule together with IH:

$$\begin{aligned} \mathbb{P} \left[\bigwedge_{j \in S_1} \neg \mathcal{A}_j \mid \bigwedge_{k \in S_2} \neg \mathcal{A}_k \right] &= \prod_{j=1}^r \mathbb{P} \left[\neg \mathcal{A}_j \mid \left(\bigwedge_{\ell=j+1}^m \neg \mathcal{A}_\ell \right) \wedge \left(\bigwedge_{k \in S_2} \neg \mathcal{A}_k \right) \right] \\ &= \prod_{j=1}^r \left(1 - \mathbb{P} \left[\mathcal{A}_j \mid \left(\bigwedge_{\ell=j+1}^m \neg \mathcal{A}_\ell \right) \wedge \left(\bigwedge_{k \in S_2} \neg \mathcal{A}_k \right) \right] \right) \\ &\stackrel{\text{IH}}{\geq} \prod_{j=1}^r \left(1 - \frac{1}{d+1} \right) \\ &\stackrel{r=|S_1| \leq d}{\geq} \left(1 - \frac{1}{d+1} \right)^d > \frac{1}{e}. \end{aligned}$$

Putting this together with the bound on the numerator, we obtain that

$$\mathbb{P} \left[\mathcal{A}_i \mid \bigwedge_{j \in S} \neg \mathcal{A}_j \right] \leq \frac{p}{1/e} \leq \frac{1}{d+1}$$

as desired. □

To prove [Lemma 5.2](#), we applying the above lemma repeatedly,

$$\begin{aligned} \mathbb{P} \left[\bigwedge_{i=1}^n \neg \mathcal{A}_i \right] &= \prod_{i=1}^n \mathbb{P} \left[\neg \mathcal{A}_i \mid \bigwedge_{j=i+1}^n \neg \mathcal{A}_j \right] \\ &= \prod_{i=1}^n \left(1 - \mathbb{P} \left[\mathcal{A}_i \mid \bigwedge_{j=i+1}^n \neg \mathcal{A}_j \right] \right) \\ &\geq \prod_{i=1}^n \left(1 - \frac{1}{d+1} \right) = \left(1 - \frac{1}{d+1} \right)^n > 0. \end{aligned}$$

Note that the above proof is not algorithmic. After around 20 years of intense study Moser and Tardos manage to give an algorithmic proof of the lemma which we will discuss in the next lecture.

5.4 General Lovasz Local Lemma

In some settings it is useful to have a more flexible version of the Local Lemma, which allows large differences in the probabilities of the "bad" events. We state this next.

Lemma 5.7 (General Lovász Local Lemma). *Let $\mathcal{A}_1, \dots, \mathcal{A}_n$ be a set of "bad" events, and let $\Gamma_i \subseteq \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$ denote the "dependency set" of \mathcal{A}_i (i.e., \mathcal{A}_i is mutually independent of all events not in Γ_i). If there exists a set of real numbers $x_1, \dots, x_n \in [0, 1]$ such that*

$$\mathbb{P}[\mathcal{A}_i] \leq x_i \prod_{j \in \Gamma_i} (1 - x_j), \quad \forall 1 \leq i \leq n$$

then

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

Exercise 5.8. *It is left as a straightforward (and strongly recommended) exercise to prove this general version by mimicking the proof of [Lemma 5.2](#). Also, you should check that applying [Lemma 5.7](#) with $x_i = 1/(d+1)$ yields [Lemma 5.2](#) as a special case.*

Corollary 5.9 (Asymmetric Lovász Local Lemma). *In the same scenario as in [Lemma 5.7](#), if $\sum_{j \in \Gamma_i} \mathbb{P}[\mathcal{A}_j] \leq 1/4$ for all i then*

$$\mathbb{P} \left[\bigwedge_{i=1}^n \neg \mathcal{A}_i \right] \geq \prod_{i=1}^n (1 - 2\mathbb{P}[\mathcal{A}_i]) > 0.$$

Proof. The result follows easily by applying [Lemma 5.7](#) with $x_i = 2\mathbb{P}[\mathcal{A}_i]$. □