

Lecture 10

In which we go over a more powerful (but difficult to compute) alternative to the spectral norm, and discuss how to approximate it.

Today we'll discuss a solution to the issue of high-degree vertices distorting spectral norms, which will prepare us for next lecture's discussion on community detection in the stochastic block model using SDP. We'll discuss a new kind of norm, the *infinity-to-one* norm, and find an efficient way to approximate it using SDP.

1 Fun with Norms

In the past few lectures, we've been heavily relying on the spectral norm,

$$\|M\| = \max_{\|x\|=1} |x^T M x| = \max_{\|x\|=1, \|y\|=1} x^T M y$$

which is efficiently computable and turns out to be pretty handy in a lot of cases.

Unfortunately, high-degree vertices have a disproportionately large influence on the spectral norm's value, limiting its usefulness in graphs where such outliers exist. Often (as we did in lecture 9), people will try to modify the input so that there are no vertices of overly high degree or add some regularizing term to ameliorate this issue. Unfortunately, this can lead to less useful results - in lecture 9, for instance, we derived a bound that required that all high-degree vertices be excised from the input graph first.

In this lecture, we'll attack this problem in a different way by introducing a different norm, the *infinity-to-one* norm, defined as follows:

$$\|M\|_{\infty \rightarrow 1} := \max_{x, y \in \{-1, 1\}^n} x^T M y .$$

It can be shown that

$$\|M\|_{\infty \rightarrow 1} \leq n \|M\| .$$

So spectral norm always gives us a bound for the infinity-to-one norm. However, the infinity-to-one norm can come in even more handy than the spectral norm (if we can actually calculate it):

Theorem 1 Let P be a symmetric matrix with entries in $[0, 1]$. Pick a random graph G such that $\{i, j\}$ is in G w.p. P_{ij} . Then whp:

$$\|A - P\|_{\infty \rightarrow 1} \leq O\left(\sqrt{n \sum_{ij} P_{ij}}\right) = O(n\sqrt{d}).$$

For context, recall that we proved a similar bound in Lecture 9 for the spectral norm, but we required that all nodes with degree greater than twice the average degree be removed first. The zero-to-one norm allows us to bypass that restriction entirely.

PROOF: Fix $\mathbf{x}, \mathbf{y} \in \{+1, -1\}^n$. Examine the expression

$$\mathbf{P}[\mathbf{x}^T(A - P)\mathbf{y} \geq t] = \mathbf{P}\left[\sum_{ij} x_i y_j (A_{ij} - P_{ij}) \geq t\right]$$

We want to make sure that this probability exponentially decreases w.r.t. n .

Recall that Bernstein's inequality states that given N independent random variables, absolutely bounded by M , with expectation 0, we have $P[\sum_i X_i > t] \leq \exp\left(-\frac{t^2/2}{\sum_i EX_i^2 + 3Mt}\right)$.

$A_{ij} - P_{ij}$ is either $-P_{ij}$ or $1 - P_{ij}$ and therefore is bounded by $[-1, +1]$. Since $\mathbf{x}, \mathbf{y} \in \{+1, -1\}^n$, we can have the M from Bernstein's inequality take on value 1. Combining this with the fact that

$$\mathbb{E}(x_i y_i (A_{ij} - P_{ij}))^2 \leq P_{ij},$$

Bernstein's inequality gives us

$$\mathbf{P}\left[\sum_{ij} x_i y_j (A_{ij} - P_{ij}) \geq t\right] \leq \exp\left(-\frac{t^2/2}{\sum_{ij} P_{ij} + 3t}\right) \leq 2^{-3n}$$

provided that $t > \sqrt{n \sum_{i,j} P_{ij}}$, as desired. \square So this norm allows us to easily sidestep the issue of high-degree vertices. The problem is that it's NP-hard to compute.

2 Grothendieck's Inequality

However, it turns out that we can *approximate* the infinity-to-one norm to within a constant factor using SDP:

Theorem 2 (*Grothendieck's Inequality*) There exists some c (turns out to be around 1.7, but we won't worry about its exact value here) such that for all M ,

$$\begin{aligned} \max_{\mathbf{x}_1, \dots, \mathbf{x}_n} \sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle &\leq c \max_{\mathbf{x}, \mathbf{y} \in \{-1, 1\}^n} \sum_{i,j} M_{ij} x_i y_j = \|M\|_{\infty \rightarrow 1} \\ \mathbf{y}_1, \dots, \mathbf{y}_n & \\ \|\mathbf{x}_i\| &= 1 \\ \|\mathbf{y}_j\| &= 1 \end{aligned}$$

So instead of dealing the combinatorially huge problem of optimizing over all ± 1 vectors, we can just solve an SDP instead to get a good approximation. For convenience, let's denote the quantity on the left of the above expression as $\|M\|_{SDP}$.

Let's start with a warmup lemma. The proof techniques in this lemma - specifically, the trick of replacing a vector from a continuous distribution with a *random* vector from some discrete distribution, and then taking the expectation to relate the two quantities, will come in handy later on.

Lemma 3 $\max_{\mathbf{x}, \mathbf{y} \in \{-1, 1\}^n} \mathbf{x}^T M \mathbf{y} = \max_{\mathbf{x}, \mathbf{y} \in [-1, 1]^n} \mathbf{x}^T M \mathbf{y}$. *In other words, maximizing over the discrete space of ± 1 random vectors and maximizing over the continuous space of vectors in that range gives the same result.*

PROOF: It is obvious that the expression on the left is at least the right (since it's just a relaxation). In order to show that the right is at least the left: given some continuous vectors $\mathbf{x}_i, \mathbf{y}_i$, we can find discrete $-1, 1$ vectors $\mathbf{b}_i, \mathbf{c}_i$ such that their expectations are equal to $\mathbf{x}_i, \mathbf{y}_i$. We can do that by having \mathbf{b}_i take on value $+1$ w.p. $1/2 + x_i/2$ and -1 w.p. $1/2 - x_i/2$, likewise for \mathbf{c}_i .

That means that:

$$\begin{aligned} \mathbb{E} \sum_{i,j} M_{i,j} \mathbf{b}_i \mathbf{c}_j &= \sum_{i,j} M_{i,j} (\mathbb{E} \mathbf{b}_i) (\mathbb{E} \mathbf{c}_j) \\ &= \sum_{i,j} M_{i,j} \mathbf{x}_i \mathbf{y}_j \end{aligned}$$

which gives us the desired result. \square

Fact 4 $\|M\|_{SDP}$ is a norm.

PROOF: **Multiplicative scaling:** obvious.

Nonnegativity (except iff $M = 0$): It is obvious that the SDP norm is zero if M is zero.

Now suppose M is nonzero.

Notice that we can replace the constraints in the SDP norm requiring that $\|\mathbf{x}_i\| = 1$, $\|\mathbf{y}_i\| = 1$ with $\|\mathbf{x}_i\| \leq 1$, $\|\mathbf{y}_i\| \leq 1$. Why? We'll use the same trick as we did in the proof of Lemma 3:

Obviously maximizing over $\|\mathbf{x}_i\| \leq 1$, $\|\mathbf{y}_i\| \leq 1$ will give us at least as good a result as maximizing over $\|\mathbf{x}_i\| = 1$, $\|\mathbf{y}_i\| = 1$, since it's a relaxation, so it suffices to show that if we can obtain some value for $\sum_{i,j} M_{i,j} \langle \mathbf{x}_i, \mathbf{y}_j \rangle$ using $\|\mathbf{x}_i\| \leq 1$, $\|\mathbf{y}_i\| \leq 1$, we can do at least as well using $\|\mathbf{x}_i\| = 1$, $\|\mathbf{y}_i\| = 1$.

Let's suppose we have some vectors $\mathbf{x}_i, \mathbf{y}_i$ with length at most 1. Now let's replace them with *random* vectors $\mathbf{r}_i, \mathbf{s}_i$ of length exactly 1 whose expectation are $\mathbf{x}_i, \mathbf{y}_i$ respectively (just scale the $\mathbf{x}_i, \mathbf{y}_i$ up, and have $\mathbf{r}_i, \mathbf{s}_i$ be either the scaled value or its negative with probability

calibrated appropriately so their expectations work out to be $\mathbf{x}_i, \mathbf{y}_i$). Then we can just say:

$$\begin{aligned} \mathbb{E} \sum_{i,j} M_{ij} \langle \mathbf{r}_i, \mathbf{s}_j \rangle &= \sum_{i,j} M_{ij} \langle \mathbb{E} \mathbf{r}_i, \mathbb{E} \mathbf{s}_j \rangle \\ &= \sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle . \end{aligned}$$

So $\mathbf{r}_i, \mathbf{s}_j$ must take on some values (of length exactly 1) that make $\sum_{i,j} M_{ij} \langle \mathbf{r}_i, \mathbf{s}_j \rangle \geq \sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle$, as desired.

Since we assumed M is nonzero, let a, b be such that $M_{a,b} \neq 0$. If we set $\mathbf{x}_a = \pm \mathbf{y}_b$ to be some arbitrary vector of unit length - using a plus if $M_{a,b}$ is positive and a minus if it's negative - and all other $\mathbf{x}_i, \mathbf{y}_j$ to zero (which we can do without affecting the value of the max by the fact we just proved), we can immediately see that

$$\sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle = M_{ab} \langle \mathbf{x}_a, \mathbf{y}_b \rangle > 0,$$

giving a positive lower bound for the maximizer, as desired.

Triangle inequality: Just look at the x and y that maximize $\sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle$ for $M = A$ and $M = B$, and observe that

$$\begin{array}{ccc} \max_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_n \\ \mathbf{y}_1, \dots, \mathbf{y}_n \\ \|\mathbf{x}_i\| = 1 \\ \|\mathbf{y}_j\| = 1}} \sum_{i,j} (A+B)_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle \leq & \max_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_n \\ \mathbf{y}_1, \dots, \mathbf{y}_n \\ \|\mathbf{x}_i\| = 1 \\ \|\mathbf{y}_j\| = 1}} \sum_{i,j} A_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle + & \max_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_n \\ \mathbf{y}_1, \dots, \mathbf{y}_n \\ \|\mathbf{x}_i\| = 1 \\ \|\mathbf{y}_j\| = 1}} \sum_{i,j} B_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle \end{array}$$

since we can always match the quantity on the left hand side by choosing the same \mathbf{x}_i and \mathbf{y}_j for both terms on the right-hand-side. \square

2.1 Proof of Grothendieck's inequality

Now let's prove Grothendieck's inequality:

PROOF: Observe that, by Lemma 3, maximizing over the choice of $\mathbf{x}, \mathbf{y} \in \{-1, 1\}^n$ in the zero-to-one norm is equivalent to maximizing over the choice of $\mathbf{x}, \mathbf{y} \in [-1, 1]^n$, so we can rewrite our proof goal as

$$\begin{array}{l} \max_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_n \\ \mathbf{y}_1, \dots, \mathbf{y}_n \\ \|\mathbf{x}_i\| = 1 \\ \|\mathbf{y}_j\| = 1}} \sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle \leq c \max_{\mathbf{x}, \mathbf{y} \in [-1, 1]^n} \sum_{i,j} M_{ij} x_i y_j \end{array}$$

Let $\mathbf{x}_1, \dots, \mathbf{x}_n, \mathbf{y}_1, \dots, \mathbf{y}_n$ be of length 1 and be the optimal choices used in $\|M\|_{SDP}$, i.e. the maximizers of $\sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle$. It suffices to prove that there exist vectors $\mathbf{b}, \mathbf{c} \in [-B, B]^n$ (for some fixed constant B , since we can just scale the result by changing c) to plug into the right-hand side of the above expression such that $c \sum_{i,j} M_{ij} b_i c_j \geq \sum_{i,j} M_{ij} \langle \mathbf{x}_i, \mathbf{y}_j \rangle$.

Pick $\mathbf{g} = (g_1, \dots, g_m)$, with each coordinate being drawn from the normal distribution with mean 0 and variance 1, and let $b_i = \langle \mathbf{g}, \mathbf{x}_i \rangle$ and $c_i = \langle \mathbf{g}, \mathbf{y}_j \rangle$. Then

$$\begin{aligned} \mathbb{E} b_i c_j &= \mathbb{E} \mathbf{x}_i^T \mathbf{g} \mathbf{g}^T \mathbf{y}_j \\ &= \mathbf{x}_i^T (\mathbb{E} \mathbf{g} \mathbf{g}^T) \mathbf{y}_j \end{aligned}$$

But $\mathbb{E} \mathbf{g} \mathbf{g}^T$ is just the identity, since the diagonal elements are just the expectation of the square of the Gaussian (which is 1, its variance) and the off-diagonals are the expectation of the product of two independent Gaussians, which is zero (since the expectation of each individual Gaussian is zero).

So

$$\mathbb{E} b_i c_j = \mathbf{x}_i^T \mathbf{y}_j = \langle \mathbf{x}_i, \mathbf{y}_j \rangle .$$

At this point we've got something that looks a lot like what we want: if the expectation of $b_i c_j$ is equal to $\langle \mathbf{x}_i, \mathbf{y}_j \rangle$, then maximizing over them is definitely going to give us a quantity greater than or equal to $\langle \mathbf{x}_i, \mathbf{y}_j \rangle$. Unfortunately, there's an issue here: since Gaussian random variables are unbounded, b_i and c_i are unbounded; on the other hand, what we're trying to do is maximize over vectors whose elements are bounded by $[B, B]$.

Our approach will be to "clip" the b_i and c_i to a finite range, and then bound the error introduced by the clipping process. Formally, fix a constant B and pick a Gaussian random vector \mathbf{g} . Now define a 'clipped' inner product:

$$b_i := \begin{cases} -B & \text{if } \langle \mathbf{x}_i, \mathbf{g} \rangle < -B \\ \langle \mathbf{x}_i, \mathbf{g} \rangle & \text{if } -B \leq \langle \mathbf{x}_i, \mathbf{g} \rangle \leq B \\ B & \text{if } \langle \mathbf{x}_i, \mathbf{g} \rangle > B \end{cases}$$

and likewise for c_j . For convenience, let's define the *truncation error* as follows, to represent how far the clipped value differs from the actual value.

$$t(z) := \begin{cases} z + B & \text{if } z < -B \\ 0 & \text{if } -B \leq z \leq B \\ z - B & \text{if } z > B \end{cases}$$

So we can rewrite b as:

$$b_i := \langle \mathbf{x}_i, \mathbf{g} \rangle - t(\langle \mathbf{x}_i, \mathbf{g} \rangle)$$

and similarly we can define:

$$c_j := \langle \mathbf{y}_j, \mathbf{g} \rangle - t(\langle \mathbf{y}_j, \mathbf{g} \rangle) .$$

So now we have

$$\begin{aligned}
\mathbb{E} \sum_{ij} M_{ij} b_i c_j &= \sum_{ij} M_{ij} \mathbb{E} [(\langle \mathbf{x}_i, \mathbf{g} \rangle - t(\langle \mathbf{x}_i, \mathbf{g} \rangle)) (\langle \mathbf{y}_j, \mathbf{g} \rangle - t(\langle \mathbf{y}_j, \mathbf{g} \rangle))] \\
&= \sum_{ij} M_{ij} \mathbb{E} [\langle \mathbf{g}, \mathbf{x}_i \rangle \langle \mathbf{g}, \mathbf{y}_j \rangle] \\
&\quad - \sum_{ij} M_{ij} \mathbb{E} [\langle \mathbf{g}, \mathbf{x}_i \rangle t(\langle \mathbf{g}, \mathbf{y}_j \rangle)] \\
&\quad - \sum_{ij} M_{ij} \mathbb{E} [t(\langle \mathbf{g}, \mathbf{x}_i \rangle) \langle \mathbf{g}, \mathbf{y}_j \rangle] \\
&\quad + \sum_{ij} M_{ij} \mathbb{E} [t(\langle \mathbf{g}, \mathbf{x}_i \rangle) t(\langle \mathbf{g}, \mathbf{y}_j \rangle)]
\end{aligned}$$

Clearly $\sum_{ij} M_{ij} \mathbb{E} \langle \mathbf{g}, \mathbf{x}_i \rangle \langle \mathbf{g}, \mathbf{y}_j \rangle$ is just the SDP norm, since $\langle \mathbf{g}, \mathbf{x}_i \rangle$ and $\langle \mathbf{g}, \mathbf{y}_j \rangle$ were the 'original' b_i and c_j .

What we will show now is that the remaining three terms are bounded by constant factors of SDP norm, so the entire sum of all four terms getting a constant factor approximation of it. The analysis is the same for all three terms so, for brevity, we'll just look at the last one:

$$\sum_{ij} M_{ij} \mathbb{E} t(\langle \mathbf{g}, \mathbf{x}_i \rangle) t(\langle \mathbf{g}, \mathbf{y}_j \rangle) .$$

For convenience, let's define $f_i, h_j : \mathbb{R}^m \rightarrow \mathbb{R}$ as:

$$f_i(\mathbf{g}) := t(\langle \mathbf{x}_i, \mathbf{g} \rangle)$$

$$h_j(\mathbf{g}) = t(\langle \mathbf{y}_j, \mathbf{g} \rangle) .$$

Now, it's convenient to think of the f_i and h_i as vectors of infinite dimension indexed by input g (we'll bring the dimensionality down to a finite value at the end of the proof). Let's define the following inner product and norm in this space:

$$\langle f_1, f_2 \rangle = \mathbb{E}_g f_1(\mathbf{g}) f_2(\mathbf{g})$$

$$\|f\|^2 = \mathbb{E}_g f^2(\mathbf{g})$$

Now, since the Gaussian distribution is rotation-independent (we can just rotate a Gaussian random variable around without changing its distribution), the squared norms of the f_i and the h_j are all the same (since all the $\mathbf{x}_i, \mathbf{y}_j$ have length 1, and dotting with them can be thought of a rotation). That means that the above norm takes on the same value for all f, h , so all we need to do is figure out a constant bound on it.

Fortunately, this value is pretty easy to bound. Notice that the function t is zero if the dot product's absolute value is smaller than B . If we have $w \sim N(0, 1)$,

$$|t(w)|^2 \leq |w|^2 \mathbf{1}_{\{|w| \geq B\}} \leq e^{-\Omega(B^2)} \leq 1/(10B)$$

for sufficiently large B , which is a constant.

So if the squared norms are bounded above by $1/(10B)$, which means we can substitute this and the norm bound into the above expression to get:

Now, armed with this, we can conclude (with similar results for the other two error terms)

$$\sum_{ij} M_{ij} \mathbb{E} t(\langle g, \mathbf{x}_i \rangle) t(\langle g, \mathbf{y}_j \rangle) = \sum_{ij} M_{ij} \langle f_i, h_j \rangle \leq \|M\|_{SDP} (1/100B^2).$$

which tells us that $\mathbb{E} \sum_{ij} M_{ij} b_i c_j$, is within a constant bound of the SDP norm $\sum_{ij} M_{ij} \mathbb{E} \langle g, \mathbf{x}_i \rangle \langle g, \mathbf{y}_j \rangle$ as desired.

There's only one slightly fishy bit in the proof we used above, though, and that's the treatment of functions of infinite-dimensional vectors indexed by a Gaussian vector. Let's conclude by constructing a (finite-dimensional) solution to the SDP from the functions:

Claim: If $X_{ij} = \langle f_i, f_j \rangle$, then $X_{ij} \succeq 0$.

Proof:

$$\begin{aligned} a^T X a &= \sum_{ij} a_i \langle f_i, f_j \rangle a_j \\ &= \sum_{ij} \langle a_i f_i, a_j f_j \rangle \\ &= \langle \sum_i a_i f_i, \sum_j a_j f_j \rangle \\ &= \left\| \sum_i a_i f_i \right\|^2 \\ &\geq 0 \end{aligned}$$

as desired. \square

So the matrix X_{ij} comprises a nice finite-dimensional solution to the SDP, and we're done with the proof.

Also, noticing that we have four terms in the expansion of $\mathbb{E} \sum_{ij} M_{ij} b_i c_j$, each one of which is a feasible value for the SDP, we can figure out how much we deviate from the actual optimum, i.e. $\|M\|_{SDP}$. Since each term can't exceed the value of the SDP optimum at all, $\mathbb{E} \sum_{ij} M_{ij} b_i c_j$ is separated from $\sum_{ij} M_{ij} \mathbb{E} [\langle g, \mathbf{x}_i \rangle \langle g, \mathbf{y}_j \rangle]$ by a factor of four, giving us a bound on the constant c in the inequality. \square

To recap, notice that there were two key "tricks" here:

- 1) Assuming that we were rounding an optimal solution to our SDP (i.e. starting with $\mathbf{x}_i, \mathbf{y}_j$ as optimizers). We don't get any bounds otherwise!
- 2) Treating the rounding error itself as a feasible solution of the SDP.

This proof was communicated to us by James Lee.