Applications of the Hoffman/Lovász bound

Yuval Filmus

February 21, 2019

1 The bounds

Suppose that $G = (V, E)$ is a graph. We will give two bounds on the size of independent sets in G.

The first bound is due to Lovász:

Lemma 1. *Suppose that M is a symmetric $V \times V$ matrix such that $M(x, y) = 1$ whenever $(x, y) \notin E$. If F is an independent set then*

$$|F| \le \lambda_{\max}(M).$$

Furthermore, if $|F| = \lambda_{\max}(M)$ then 1_F lies in the eigenspace of $\lambda_{\max}$.

Proof. On the one hand, $1_F' M 1_F = |F|^2$. On the other hand, $1_F' M 1_F \le \lambda_{\max}(M) 1_F' 1_F = \lambda_{\max}(M)|F|$. Therefore $|F| \le \lambda_{\max}$, with equality if and only if 1_F lies in the eigenspace of $\lambda_{\max}$. □

The minimum value of $\lambda_{\max}(M)$ over all such matrices M is known as $\theta(G)$. This is the famous Lovász theta function.

The second bound is due to Hoffman:

Lemma 2. *Suppose that M is a symmetric $V \times V$ matrix such that $M(x, y) = 0$ whenever $(x, y) \notin E$, and furthermore $M\mathbf{1} = \mathbf{1}$. If F is an independent set then*

$$\frac{|F|}{|V|} \le \frac{-\lambda_{\min}(M)}{1 - \lambda_{\min}(M)} = 1 - \frac{1}{1 - \lambda_{\min}(M)}.$$

Furthermore, in case of equality, $1_F - \frac{|F|}{|V|}\mathbf{1}$ lies in the eigenspace of $\lambda_{\min}(M)$.

Proof. Let $g = 1_F - \frac{|F|}{|V|}\mathbf{1}$, so that $g'\mathbf{1} = 0$. This implies that

$$g'g = 1_F' g = |F| - \frac{|F|^2}{|V|}.$$

Also, we have

$$0 = 1_F' M 1_F = (g + \tfrac{|F|}{|V|}\mathbf{1})' M (g + \tfrac{|F|}{|V|}\mathbf{1}) = g' M g + \frac{|F|^2}{|V|},$$

since $g' M \mathbf{1} = g'\mathbf{1} = 0$. Therefore

$$\frac{|F|^2}{|V|} = -g' M g \le -\lambda_{\min}(M) g' g = -\lambda_{\min}(M) \left(|F| - \frac{|F|^2}{|V|} \right).$$

Rearranging,

$$(1 - \lambda_{\min}(M)) \frac{|F|}{|V|} \le -\lambda_{\min}(M),$$

implying the upper bound. Moreover, equality holds if and only if g lies in the eigenspace of $\lambda_{\min}(M)$. □

This gives rises to the "Hoffman theta function", which we will denote by $\theta_H(G)$.

It is not hard to show that the Lovász bound is actually stronger.

Lemma 3. *For all graphs G, $\theta(G) \leq \theta_H(G)$.*

Proof. Suppose that M_H is a matrix witnessing $\theta_H(G)$. Define

$$M = \mathbf{1}\mathbf{1}' - \frac{|V|}{1 - \lambda_{\min}(M_H)} M_H.$$

On the one hand,

$$M\mathbf{1} = \left(|V| - \frac{|V|}{1 - \lambda_{\min}(M_H)} \right) \mathbf{1} = \theta_H(G)\mathbf{1}.$$

On the other hand, if v is any eigenvector orthogonal to $\mathbf{1}$ then its eigenvalue is at most

$$-\frac{|V|}{1 - \lambda_{\min}(M_H)} \cdot -\lambda_{\min}(M_H) = \theta_H(G). \qquad \square$$

This gives an alternative proof of Lemma 2.

Proof. Let M_H be a matrix satisfying the conditions of Lemma 2. Then the matrix M constructed in Lemma 3 satisfies the constraints of Lemma 1, hence the upper bound. When the upper bound is achieved, 1_F must lie in the eigenspace of $\lambda_{\max}(M)$, which is spanned by $\mathbf{1}$ and the eigenspace of $\lambda_{\min}(M_H)$. The orthogonal projection of 1_F to $\mathbf{1}$ is $\frac{|F|}{|V|}\mathbf{1}$, and so $1_F - \frac{|F|}{|V|}\mathbf{1}$ must lie in the eigenspace of $\lambda_{\min}(M_H)$. (Here we used the fact that the eigenspaces are orthogonal.) $\square$

Both arguments can be extended in various ways. Two prominent extensions are *stability* and *cross-independent sets*.

In the setting of the Lovász bound, suppose F is an independent family of size close to $\lambda_{\max}(M)$. Then 1_F must be close to the eigenspace to $\lambda_{\max}(M)$. Quantitatively, the closeness would depend on the gap between $\lambda_{\max}(M)$ and the next largest eigenvalue.

Another extension handles cross-independent families, which are families F, G such that no $x \in F$ and $y \in G$ form an edge. An application of Cauchy–Schwarz allows us to obtain an upper bound on $\sqrt{|F||G|}$. In fact, we get a stronger upper bound on $\sqrt{\frac{|F|}{|V|-|F|}\frac{|G|}{|V|-|G|}}$.

2 Traffic light puzzle

A traffic light is controlled by n three-way switches. We are given that if the state of all switches is changed, then the light changes. We will show that in this case, the traffic light is actually controlled by a single switch!

Let F denote the set of states leading to red light. Then F is an independent set in the graph on $\mathbb{Z}_3^n$ in which two vertices are connected if they differ on all coordinates. The adjacency matrix of this graph is

$$M_n = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}^{\otimes n}.$$

Let ω be a primitive third root of unity. Since $1 + \omega + \omega^2 = \frac{1-\omega^3}{1-\omega}$, we can check that

$$M_1 \begin{pmatrix} 1 & 1 & 1 \\ 1 & \omega & \omega^2 \\ 1 & \omega^2 & \omega \end{pmatrix} = \begin{pmatrix} 2 & -1 & -1 \\ 2 & -\omega & -\omega^2 \\ 2 & -\omega^2 & -\omega \end{pmatrix}.$$

Therefore the columns of this matrix are the eigenvectors of M_1, with eigenvalues $2, -1, -1$. In particular, the matrix $M := M_n/2^n$ satisfies all constraints of Hoffman's bound (Lemma 2), with $\lambda_{\min}(M) = -1/2$. Applying the bound, we get that

$$\frac{|F|}{3^n} \leq \frac{1/2}{1+1/2} = \frac{1}{3}.$$

This is true also for the sets of states leading to yellow and green lights, and since together these must cover all states, we conclude that $|F| = 3^{n-1}$.

According to Hoffman's bound, this means that $1_F - 1/3$ lies in the eigenspace of $\lambda_{\min}(M)$, which is spanned by all functions depending on a single coordinate. That is, we can write

$$1_F(x_1, \ldots, x_n) = \frac{1}{3} + \sum_{i=1}^{n} \phi_i(x_i).$$

At most one of the functions ϕ_i can be non-constant, since otherwise 1_F would not be Boolean. Hence F depends on a single coordinate i. Since F has measure $1/3$, it is of the form $\{x : x_i = j\}$.

The families corresponding to the other two colors have a similar shape. Since they must form a partition of the domain, it is not hard to check that they have to depend on the same coordinate.

3 Triangle-intersecting families

The Hoffman bound can be used to prove various intersection theorems. Here we will concentrate on one example: triangle-intersecting families of graphs.

Let $I \subseteq 2^{[n]}$. A family $F \subseteq 2^{[n]}$ is *I-intersecting* if the intersection of any two sets in F contains an element of I. It is *I-agreeing* if the agreement $\overline{\cdot \triangle \cdot}$ of any two sets in F contains an element of I.

Lemma 4. *The maximum size of an I-intersecting family is the same as the maximum size of an I-agreeing family.*

Proof. For the proof, it suffices to show that if F is I-agreeing then there is an I-intersecting family of the same size. We prove this by shifting.

An *upshift along the i'th coordinate* is the following operation. Consider all pairs of sets $\{(S, S{+}i) : i \notin S\}$. We modify F to a new family F' in the following way: if $S \in F$ but $S + i \notin F$, then we replace S with $S + i$. We claim that F' is still I-agreeing. Indeed, suppose $A, B \in F'$. For every set $A \in F$, either $A \in F'$ or $i \in A$ and $A - i \in F$. If $A, B \in F$ or $i \in A, B$ and $A - i, B - i \in F$, then this is clear, so suppose $i \in A$ and $A - i, B \in F$. If $i \in B$ then $\overline{A \triangle B} = \overline{(A - i) \triangle B} + i$. If $i \notin B$ then by construction $B + i \in F$, and then $\overline{A \triangle B} = \overline{(A - i) \triangle (B + i)}$.

Each upshift increases the total size of sets in the family, so if we perform upshifts, eventually we will reach an upshifted family G which is I-agreeing. For every $S \in G$ and $i \notin S$, we have $S + i \in G$, and so G is monotone. For every two sets $A, B \in G$, we also have $B \cup \overline{A} \in G$. The agreement of A and $B \cup \overline{A}$ is the same as $A \cap B$, and so the latter contains an element of I. In other words, G is I-intersecting. ☐

This motivates studying I-agreeing families. Let G be the graph on $[n]$ in which two sets are connected if their agreement doesn't contain any set in I, and let M be any matrix satisfying the conditions of Hoffman's bound. For set z, we can form a new matrix $M_z(x, y) = M(x \triangle z, y \triangle z)$. Since $M_z \mathbf{1} = \mathbf{1}$ and $(x \triangle z) \triangle (y \triangle z) = x \triangle y$, the matrix M_z also satisfies the conditions of Hoffman's bound. Furthermore, M_z is similar to M (apply the involution corresponding to XORing with z), and so has the same minimal eigenvalue. It follows that $M' = \mathbb{E}_z[M_z]$ also satisfies the conditions of Hoffman's bound, with a possibly large minimal eigenvalue (this can be seen by the von Neumann formula for the minimal eigenvalue, as minimizing $x'Mx$ over $x'x = 1$). So if we are interested in applying Hoffman's bound to the problem, we might as well consider the matrix M' instead. By construction, $M'_z = M'$, and so $M'(x, y) = \phi(x \triangle y)$. Later we will learn that M' belongs to the *Bose–Mesner algebra of the Hamming association scheme.*

Let us use the notation M instead of M'. The fact that M is symmetric implies that the Fourier characters are its eigenvectors:

$$(M\chi_S)(x) = \sum_y \phi(x \triangle y)\chi_S(y) = \sum_y \phi(y)\chi_S(x \triangle y) = \chi_S(x) \sum_y \phi(y)\chi_S(y).$$

In fact, we get $M\chi_S = 2^n \hat{\phi}(S)\chi_S$. This reduces computing the Hoffman bound from a semidefinite program to a linear program.

Concretely, we can consider the matrix X_z defined by $(X_z v)(x) = v(x+z)$. This matrix satisfies $X_z(x,y) = [x \triangle y = z]$ and has eigenvalues $X_z \chi_S = \chi_S(z)$.

The matrix X_z satisfies the condition of Hoffman's bound for I-intersecting families if and only if $\bar{z}$ doesn't contain any element of I. This shows that the matrices satisfying Hoffman's bound are of the form

$$M = \sum_{z \in Z} \alpha_z X_{\bar{z}},$$

where Z consists of all sets not containing any element of I. Furthermore, $\sum_{z \in Z} \alpha_z = 1$. The eigenvalue corresponding to χ_S is

$$\sum_{z \in Z} \alpha_z \chi_S(\bar{z}) = \sum_{z \in Z} \alpha_z (-1)^{|\bar{z} \cap S|} = (-1)^{|S|} \sum_{z \in Z} \alpha_z (-1)^{|z \cap S|}.$$

We call such matrices *admissible*, and the corresponding spectrum an *admissible spectrum.*

It is well-known that $\{(-1)^{|w \cap S|} : w \subseteq z\}$ spans the set of all functions on $z \cap S$. Since the set Z is downward-closed, the following spectrum is admissible for all $z \in Z$, as long as $\sum_{z \in Z} \alpha_z \phi_z(\emptyset) = 1$:

$$(-1)^{|S|} \sum_{z \in Z} \phi_z(S \cap z).$$

In fact, these are all admissible spectra.

3.1 Triangle-intersecting families of graphs

We now apply the forgoing to triangle-intersecting families of graphs. These are families of subsets of K_n in which the intersection of any two graphs contains a triangle. One way to construct such a family is to take a fixed triangle and all the graphs containing it. This family contains 1/8 of all graphs. Can we do better?

We will show that we can't do better even if we consider the class of non-bipartite-agreeing families of graphs. In this case, the set Z consists of all bipartite graphs. Let $q_k(G)$ be the probability that a random partition of G cuts exactly k edges. Since the number of edges cut by a specific partition depends only on the intersection with the corresponding complete bipartite graph, we see that the following is an admissible spectrum, assuming $\sum_k c_k = 1$:

$$(-1)^{|G|} \sum_k c_k q_k(G).$$

Since the upper bound in Hoffman's bound gets better as $-\lambda_{\min}$ gets smaller, we would like to minimize $-\lambda_{\min}$. We can do so by considering the following graphs: empty graph, one edge, two connected edges, triangle, 4-matching, K_4^-. We get that the best $-\lambda_{\min}$ is $1/7$, and this is achieved by

$$(-1)^{|G|} \left(q_0(G) - \frac{5}{7} q_1(G) - \frac{1}{7} q_2(G) + \frac{3}{28} q_3(G) \right).$$

If this quantity were always at least $-1/7$, it would follow that any non-bipartite-agreeing family of graphs contains at most a fraction of $\frac{1/7}{1+1/7} = 1/8$ of graphs.

Here is what happens for the small graphs mentioned above:

G	$q_0(G)$	$q_1(G)$	$q_2(G)$	$q_3(G)$	$q_4(G)$
$\emptyset$	1	0	0	0	0
$\vert$	1/2	1/2	0	0	0
$\vert\vert$	1/4	1/2	1/4	0	0
$\vert\vert\vert$	1/8	3/8	3/8	1/8	0
$\vert\vert\vert\vert$	1/16	1/4	3/8	1/4	1/16
$\triangle$	1/4	0	3/4	0	0
$\boxslash$	1/8	0	1/4	1/2	1/8

Actually, the uniqueness condition in Hoffman's bound shows that for subgraphs of the triangle, the eigenvalue must be exactly $-1/7$. Together with the condition for the empty graph, we can deduce the coefficients c_0, c_1, c_2 of q_0, q_1, q_2. We also get the value of $c_3 + c_4/4$ by considering $||||$ and $\boxslash$. For simplicity, we choose $c_4 = 0$.

To show that all eigenvalues are at least $-1/7$, note that if G contains many edges, then it is unlikely that a random cut will cut few edges. Therefore as $|G| \to \infty$, the eigenvalue will tend to 0. Quantitatively, we can find some edge threshold m beyond which the eigenvalue is at least $-1/7$. We can then in principle check all graphs with at most m edges. All eigenvalues will be at least $-1/7$.

With a bit more work, we can design a function (which also includes $q_H(G)$, functions which measure the probability that a random cut results exactly in a graph H) in which the eigenspace of $\lambda_{\min}$ consists of all graphs with one, two or three edges. It follows that a non-bipartite-agreeing family of measure 1/8 must be a degree three function, and so a 3-junta, and so a triangle-junta. This function is obtained from the preceding one by adding a small enough positive multiple of the following:

$$(-1)^{|G|} \left(q_{F_4}(G) - q_{\square}(G)\right),$$

where $q_{F_4}(G)$ is the probability that a random partition of G cuts exactly 4 edges arranged as a forest, and $q_{\square}(G)$ is the probability that it cuts exactly 4 edges arranged as a C_4.

4 Planted clique

- The problem.
- Cliques of size $\sqrt{n \log n}$ using a degree argument.
- Boosting by guessing vertices.
- Feige–Krauthgamer.