Eigenvalues of graphs

Lecture 8+9+10 - Advanced Discrete Mathematics 2013 - EPFL - Frank de Zeeuw

---

1. Definition and examples • 2. Eigenvalues and properties of graphs • 3. Proofs with eigenvalues

---

1 Definition and examples

Definitions: An *eigenvalue* of a graph G is an eigenvalue of its adjacency matrix A_G; ie a $\lambda \in \mathbb{R}$ for which there is an *eigenvector* $v \in \mathbb{R}^{|V(G)|}$, $v \neq 0$, such that

$$A_G v = \lambda v.$$

The *multiplicity* $m(\lambda)$ of λ is the dimension of the subspace of $\mathbb{R}^{|V(G)|}$ spanned by all eigenvectors for λ (its *eigenspace*). The *spectrum* of G is the multiset of eigenvalues λ with their multiplicities $m(\lambda)$, which we denote as follows:

$$\text{Spec}(G) = \text{Spec}(A_G) = (\lambda)^{m(\lambda)}(\mu)^{m(\mu)} \cdots (\omega)^{m(\omega)}.$$

Usually we will name the eigenvalues $\lambda_1, \lambda_2, ..., \lambda_n$, where $n = |V(G)|$, so that

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n.$$

Note that these two notations clash a little, since in the first the $\lambda, \mu, ...$ are distinct and there can be $< n$ of them, while in the second they need not be distinct, but there are always exactly n.

Remarks:
• There is no easy way to visualize the eigenvalues of a graph, but here is one attempt. You could think of an eigenvector as a function f on the vertices. Then if at any vertex v you sum up the values of f on its neighboring vertices, you should get λ times the values of f at v. Formally:

$$\sum_{u \in N(v)} f(u) = \lambda f(v).$$

• As you may know, eigenvalues can be computed by finding the roots of the characteristic polynomial $f(x) = \det(A - xI)$. Here we will avoid that as much as possible, because the typical graphs that we will deal with (complete graphs, cycles, etc) have some kind of symmetry that allows the eigenvalues (and eigenvectors) to be found in smarter and less computational ways.
• There are other kinds of eigenvalues of graphs, for instance the Laplacian eigenvalues, which are often used in applications. To avoid confusion, we'll ignore these completely. The same goes for eigenvalues of directed graphs or multigraphs.

We will frequently use the following theorem from linear algebra, but we won't prove it here. I'll refer to it as the Spectral Theorem in these lectures, although that theorem has many different forms.
It is based on the fact that an adjacency matrix is real and symmetric. This implies that the eigenvalues are real, though note that the eigenvectors need not be.

Theorem 1.1 (Spectral Theorem)**.**

a) The eigenvalues of a graph G are always real.

b) The adjacency matrix A_G is diagonalizable.

c) There is an orthonormal basis of eigenvectors.

We will now derive the spectra of several graphs. We avoid determinant calculations completely.
By J we mean the matrix with all entries 1, and by j the vector with all entries 1.

Example: $\text{Spec}(K_n) = (n-1)^1(-1)^{n-1}$

Proof. The adjacency matrix is $A = J - I$. It's easy to guess that j is an eigenvector: $Jj = nj$, so $Aj = Jj - Ij = (n-1)j$. That means $(n-1)$ is an eigenvalue, but we don't know its multiplicity yet.
By Theorem 1.1, when looking for other eigenvectors v we can assume that they are orthogonal to j. That implies that $Jv = 0$, so $Av = Jv - Iv = -v$ for each such orthogonal eigenvector. Therefore -1 is the only other eigenvalue and its multiplicity is $n-1$ (the dimension of the orthogonal complement of j). It follows that $n-1$ has multiplicity 1. □

Example: $\text{Spec}(K_{mn}) = (\sqrt{mn})^1(0)^{m+n-2}(-\sqrt{mn})^1$

Proof. We have

$$A = \begin{bmatrix} 0_{mm} & J_{mn} \\ J_{nm} & 0_{nn} \end{bmatrix},$$

so an eigenvalue and its eigenvector (split into vectors v, w of size m, n) would satisfy

$$A \begin{bmatrix} v \\ w \end{bmatrix} = \begin{bmatrix} J_{mn}w \\ J_{nm}v \end{bmatrix} = \begin{bmatrix} \lambda v \\ \lambda w \end{bmatrix}.$$

Now j_{m+n} is not an eigenvector (unless $m = n$), but v is a vector of size m that is orthogonal to j_m, then $J_{nm}v = 0$ so

$$A \begin{bmatrix} v \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} = 0 \cdot \begin{bmatrix} v \\ 0 \end{bmatrix}.$$

The same works for a vector w of size n orthogonal to j_n. The dimension of such orthogonal vectors is $m-1$ for j_m and $n-1$ for j_n, so the multiplicity of 0 is at least $m+n-2$.
Any other eigenvector would have to be of the form

$$a \begin{bmatrix} j_m \\ 0 \end{bmatrix} + b \begin{bmatrix} 0 \\ j_n \end{bmatrix} = \begin{bmatrix} a \cdot j_m \\ b \cdot j_n \end{bmatrix}.$$

So the corresponding eigenvalue would satisfy

$$\lambda \begin{bmatrix} a \cdot j_m \\ b \cdot j_n \end{bmatrix} = A \begin{bmatrix} a \cdot j_m \\ b \cdot j_n \end{bmatrix} = \begin{bmatrix} b \cdot J_{mn} j_n \\ a \cdot J_{nm} j_m \end{bmatrix} = \begin{bmatrix} bn \cdot j_m \\ am \cdot j_n \end{bmatrix}.$$

We get that $\lambda a = bn$ and $\lambda b = am$. If $ab \neq 0$, multiplying these together gives $\lambda^2 = mn$, so $\lambda = \pm\sqrt{mn}$ could be eigenvalues. If $ab = 0$, we don't get an eigenvector.
Each of these potential eigenvalues does indeed have an eigenvector: For $\sqrt{mn}$ we can take $a = 1/\sqrt{m}$, $b = 1/\sqrt{n}$, and for $-\sqrt{mn}$ we can take $a = -1/\sqrt{m}$, $b = 1/\sqrt{n}$.
It follows that $\pm\sqrt{mn}$ both have multiplicity 1 and 0 has multiplicity $m + n - 2$. □

The *Petersen graph* Pet can be constructed by taking all 2-element subsets of $\{1, 2, 3, 4, 5\}$ as vertices, and connecting two by an edge if they are disjoint.
It is 3-regular since a 2-element sets is disjoint from 3 other 2-element sets.

Example: $\text{Spec}(Pet) = (3)^1(1)^5(-2)^4$

Proof. Let $A = A_{Pet}$. It is easy to see that $Aj = 3j$, because the graph is 3-regular, so 3 is an eigenvalue.
We consider A^2. Its ST-entry A^2_{ST} counts walks of length 2 from S to T. If S and T are neighbors (ie disjoint sets), then there is no such walk (it would mean there is a set disjoint from both, not possible with 5 elements), so $A_{ST} = 0$. If S and T are not neighbors, then they share exactly one element, their union has 3 elements, and so there is exactly one set disjoint from both, which means there is one walk of length 2 from S to T, so $A_{ST} = 1$. When $S = T$ we have $A_{ST} = 3$.
This tells us exactly that

$$A^2 + A = J + 2I.$$

Let v be an eigenvector orthogonal to j with eigenvalue λ. Then

$$0 = (A^2 + A - J - 2I)v = A^2 v + Av - 2Iv = (\lambda^2 + \lambda - 2)v.$$

So $\lambda = 1$ or $\lambda = -2$. This implies that $\lambda = 3$ has multiplicity 1, since otherwise it would show up here.
Write a for the multiplicity of 1 and b for that of -2. Then

$$0 = \text{Tr}(A) = \sum_{i=1}^{n} \lambda_i = 3 + a \cdot 1 + b \cdot (-2),$$

so $2b - a = 3$. We also have $a + b = 9$, and by solving we get $a = 5$, $b = 4$. □

To conclude we will compute the eigenvalues of cycles C_n. This is a bit harder, and there are several ways of doing it, but this one seems the most natural to me.
Note that for some n, some of the eigenvalues listed in the spectrum will coincide, so technically we should be putting those together as a single eigenvalue with multiplicity 2. But writing it like this is more convenient.

Example: $\mathrm{Spec}(C_n) = (2)^1(2\cos(2\pi/n))^1(2\cos(2\cdot 2\pi/n))^1\cdots(2\cos((n-1)\cdot 2\pi/n))^1$

Proof. Let W be the $n\times n$ matrix that has first row $(0,1,0,\ldots,0)$, and each subsequent row equals the one above it, but shifted to the right by one position. So the second-to-last row is $(0,\ldots,0,1)$, and the last row is $(1,0,\ldots,0)$. In other words, W has ones only right above the diagonal, and in the bottom-left corner.
Note that W^k is then the permutation matrix whose first row has a single 1, in position $k+1$, and whose subsequent rows are shifted to the right as for W. This makes sense for any k by just taking it modulo k.
The crucial thing is that

$$A_{C_n} = W + W^{-1},$$

which will let us determine the eigenvalues of C_n from those of W and W^{-1}. We can determine the eigenvectors of W as follows. It acts on a vector by shifting each entry up by one position, with the first entry becoming the last. So if $v = (v_1,\ldots,v_n)^T$ is an eigenvector with eigenvalue λ, then we have

$$v_1 = \lambda v_n = \lambda^2 v_{n-1} = \cdots = \lambda^n v_1,$$

so $\lambda^n = 1$ (no $v_i = 0$ because then this equation would make all of them 0). This means that the eigenvalues are among the nth roots of unity, which we can write as ω^l for $\omega = e^{\frac{2\pi i}{n}}$. Indeed, each of these is an eigenvalue $\lambda_l = \omega^l$ with multiplicity 1, because if we choose $v_1 = 1$, the equations above give the vector

$$u_l = (1, \omega^l, \omega^{2l}, \ldots, \omega^{(n-1)l})^T.$$

So

$$\mathrm{Spec}(W) = (1)^1(\omega)^1(\omega^2)^1\cdots(\omega^{n-1})^1, \quad \text{and} \quad \mathrm{Spec}(W^k) = (1)^1(\omega^k)^1(\omega^{2k})^1\cdots(\omega^{(n-1)k})^1.$$

Because each W^k has the same eigenvector for corresponding eigenvalues, we easily get the eigenvalues of C_n, because

$$A_{C_n}u_l = Wu_l + W^{-1}u_l = (\omega^l + \omega^{-l})u_l.$$

Therefore the eigenvalues of C_n are

$$\omega^l + \omega^{-l} = 2\cos(2\pi l/n), \quad \text{for} \quad l = 0,1,\ldots,n-1.$$

□

2 Eigenvalues and properties of graphs

Now we will see several lemmas that illustrate the connection between the eigenvalues of a graph and its graph-theoretical properties, like degrees, bipartiteness, or connectedness. These are just a sample and there are many more such connections.

Lemma 2.1. *The largest eigenvalue λ_1 of a graph G lies between the average and maximum degrees:*

$$d_{avg} \leq \lambda_1 \leq d_{max}$$

In particular, if G is d-regular, then $\lambda_1 = d$.

Proof. To prove $\lambda_1 \leq d_{max}$, let $x = (x_v)_{v \in V(G)}$ be an eigenvector corresponding to λ_1, and let x_u be the entry of x with maximum absolute value. Then we have (with $N(u) = \{v \in V(G) : uv \in E(G)\}$ the neighborhood of u)

$$\lambda_1 x_u = \sum_{v \in N(u)} x_v,$$

so (with $d_u = |N(u)|$ the degree of u)

$$|\lambda_1| \cdot |x_u| \leq \sum_{v \in N(u)} |x_v| \leq \sum_{v \in N(u)} |x_u| \leq d_u \cdot |x_u| \leq d_{max} \cdot |x_u|.$$

Since $x \neq 0$, we have $|x_u| \neq 0$, so we get $\lambda_1 = |\lambda_1| \leq d_{max}$ (we have $\lambda_1 > 0$ since $0 = \mathrm{Tr}(A) = \sum \lambda_i$).

To prove $d_{avg} \leq \lambda_1$, we consider $j^T A j$. On the one hand

$$j^T A j = \sum_{v \in V(G)} d_v = 2|E(G)|.$$

On the other hand, take an orthonormal basis $\{v_1, \ldots, v_n\}$ of eigenvectors of A, and let $j = \sum c_i v_i$ the representation of j in this basis. So we have $Av_i = \lambda_i v_i$, $j^T v_i = c_i$, and $\sum c_i^2 = \|j\|^2 = n$. Then

$$j^T A j = \sum c_i j^T (Av_i) = \sum c_i j^T (\lambda_i v_i) = \sum \lambda_i c_i (j^T v_i) = \sum \lambda_i c_i^2 \leq \lambda_1 \sum c_i^2 = \lambda_1 n.$$

So we get $\lambda_1 \geq 2|E(G)|/n = d_{avg}$. □

A *walk* in a graph is a sequence of adjacent vertices (ie each vertex in the sequence is adjacent to the next one); the walk is *closed* if the first and last vertex are the same. Note that different walks may correspond to the same set of edges; in the following lemma such walks are really counted separately.

Lemma 2.2. *The number of closed walks of length k in G equals*

$$\sum_{i=1}^{n} \lambda_i^k.$$

Proof. The matrix A^k has in its uv-entry the number of walks from u to v. So a diagonal uu-entry gives the number of closed walks starting and ending at u, and the sum of the diagonal entries, $\text{Tr}(A^k)$, gives the number of closed walks of length k.
On the other hand, if λ is an eigenvalue of A then λ^k is an eigenvalue of A^k, with the same multiplicity. Diagonalizing A^k as $S^{-1}A^kS = D$, where D is the diagonal matrix with the eigenvalues of A^k on the diagonal, gives

$$\text{Tr}(A^k) = \text{Tr}(SDS^{-1}) = \text{Tr}(D) = \sum \lambda_i^k.$$

□

Lemma 2.3 (Bipartiteness)**.** *A graph is bipartite if and only if its spectrum is symmetric (ie if λ is an eigenvalue, then so is $-\lambda$, and with the same multiplicity).*

Proof. First suppose G is bipartite, with parts S and T of sizes s and t. This means that for some $s \times t$ matrix B we have

$$A = \begin{bmatrix} 0_{ss} & B \\ B^T & 0_{tt} \end{bmatrix}.$$

If λ is an eigenvalue, then

$$\begin{bmatrix} \lambda v \\ \lambda w \end{bmatrix} = \lambda \begin{bmatrix} v \\ w \end{bmatrix} = A \begin{bmatrix} v \\ w \end{bmatrix} = \begin{bmatrix} Bw \\ B^T v \end{bmatrix}.$$

So $Bw = \lambda v$ and $B^T v = \lambda w$. Then $-\lambda$ is also an eigenvalue:

$$A \begin{bmatrix} v \\ -w \end{bmatrix} = \begin{bmatrix} -Bw \\ B^T v \end{bmatrix} = \begin{bmatrix} -\lambda v \\ \lambda w \end{bmatrix} = -\lambda \begin{bmatrix} v \\ \lambda w \end{bmatrix}.$$

If λ has multiplicity m, then so does $-\lambda$, since the subspace spanned by the $[v\ w]^T$will have the same dimension as that spanned by the corresponding $[v\ \ -w]^T$. This means that the spectrum is symmetric.

Conversely, suppose that the spectrum is symmetric. Then for any odd k we have

$$\sum \lambda_i^k = 0,$$

since any term λ_i^k will be cancelled by a term $(-\lambda_i)^k = -\lambda_i^k$. By Lemma 2.2, it follows that there are no closed walks of odd length in G, so in particular there are no odd cycles. That means that G is bipartite. □

For K_{mn} this clearly checks out.
For even cycles C_{2m}, which are bipartite, it also holds, as $\cos(l \cdot \pi/m) = \cos((2m - l) \cdot \pi/m)$.

We'll just mention two similar theorems that relate familiar graph properties to eigenvalues. They will be on a problem set, and we might use them later on.

Lemma 2.4 (Connectedness)**.** *If G is d-regular, then the multiplicity of the eigenvalue λ_1 equals the number of connected components of G.*

Lemma 2.5 (Diameter)**.** *If G is connected, then the diameter of G is strictly less than its number of distinct eigenvalues.*

The following lemma is a bit harder, but next time we will see a cool application of it to extremal set theory.
Recall that $\alpha(G)$ is the *independence number* of G, the size of the largest independent set, ie a set of vertices without any edges between them.
By the *least* eigenvalue we mean the last one, λ_n, when we order them $\lambda_1 \geq \cdots \geq \lambda_n$; we don't mean the one with least absolute value.

Lemma 2.6 (Independence number)**.** *Suppose G is d-regular and λ_n is its least eigenvalue, then*

$$\alpha(G) \leq \frac{n}{1 - \frac{d}{\lambda_n}}.$$

Proof. Let S be an independent set, and let $x_S \in \mathbb{R}^{|V(G)|}$ be its characteristic vector. The fact that it is independent implies that $x_S^T A x_S = 0$.
Write $\lambda = \lambda_n$ for now. The fact that it is the least eigenvalue implies that $A - \lambda I$ has nonnegative eigenvalues. Also, since $Aj = dj$ by regularity,

$$(A - \lambda I)j = dj - \lambda j = (d - \lambda)j = \frac{d - \lambda}{n} Jj,$$

where in the last step we used $Jj = nj$.
We define a new matrix

$$M = (A - \lambda I) - \frac{d - \lambda}{n} J.$$

We just saw that $Mj = 0$, so j is an eigenvector of M with eigenvalue 0. Let v be an eigenvector of M that is orthogonal to j, with eigenvalue μ. So $Jv = 0$. Then

$$\mu v = Mv = (A - \lambda I)v,$$

hence μ is an eigenvalue of $A - \lambda I$, so it is nonnegative.
Since we can choose an orthonormal basis of eigenvectors of M, this implies that all the eigenvalues of M are nonnegative. You can check that this implies that $x_S^T M x_S \geq 0$. On the other hand,

$$0 \leq x_S^T M x_S = x_S^T A x_S - \lambda x_S^T x_S - \frac{d - \lambda}{n} x_S^T J x_S = -\lambda |S| - \frac{d - \lambda}{n} |S|^2,$$

which gives

$$|S| \leq -\lambda \frac{n}{d - \lambda} = \frac{n}{1 - \frac{d}{\lambda_n}}.$$

□

This bound is in general not tight. But bounding $\alpha(G)$ is quite hard (determining it is $\mathcal{NP}$-hard), so we're happy with anything we can get. Let's see what it does for some examples:

- For K_n, it gives $\alpha \leq 1$, which is tight but not very exciting.
- For K_{mm}, with $m = n/2$, we have $\lambda_n = -m$ and $d = m$, so $\alpha \leq \frac{n}{1 - \frac{m}{-m}} = n/2$, again tight.
- For the Petersen graph, we get $\alpha \leq 4$, which is tight, since for instance $\{12, 13, 14, 15\}$ is an independent set. It would probably be tedious to prove $\alpha = 4$ directly.
- For an even cycle C_{2m}, we have $d = 2$ and $\lambda_n = -2$, so the bound is $\alpha \leq n/2$, tight.
- But for an odd cycle C_{2m+1}, where $\lambda_n = 2\cos\left(2\pi \cdot \frac{m}{2m+1}\right)$, the bound is not tight, since you can check that it's $\geq m$, whereas $\alpha = m - 1$. Still, pretty close.

3 Proofs with eigenvalues

Theorem 3.1. *The complete graph K_{10} cannot be decomposed into 3 Petersen graphs.*

Proof. We know that $\text{Spec}(K_{10}) = (9)^1(-1)^9$ and $\text{Spec}(Pet) = (3)^1(1)^5(-2)^4$. Suppose that there is a decomposition like in the statement, and let A, B, C be the adjacency matrices of the 3 copies of Pet. Since the matrix of K_{10} is $J - I$, we have

$$J - I = A + B + C.$$

Note that we can "add up" the 3 eigenvalues 3 of the Petersen graphs, because they all have the same eigenvector j, and this gives the eigenvalue 9 of K_{10}. If we could add up the other eigenvalues like this, we could get a contradiction. Of course we can't always add eigenvalues like that, but this is the idea for the proof.
Let V_A and V_B be the eigenspaces of A and B corresponding to the eigenvalue 1. They have dimension 5 and are both orthogonal to j (since there is an orthonormal basis), so

$$\dim(V_A \cap V_B) \geq 5 + 5 - 9 = 1.$$

Hence there is an eigenvector $x \in V_A \cap V_B$ that is an eigenvector of 1 for both A and B. Since x is also orthogonal to j, we have $Jx = 0$, so

$$Cx = (J - I - A - B)x = -3x.$$

That means that -3 is an eigenvalue of C. Wait, no, it isn't. Contradiction. □

The Windmill Theorem
The following theorem is a classic from 1966 by Paul Erdős, Alfréd Rényi, and Vera Sós. There are many proofs, but the following is probably the nicest. Some of the other proofs use less linear algebra, but are considerably longer.
It is usually referred to as the *friendship theorem*, because it says that in a group of people where every two have exactly one common friend, there is a politician who is friends with everyone. Being Dutch, I'm going to phrase it differently. A *windmill graph* on $2m + 1$ vertices consists of m triangles that all share a vertex. Equivalently, there is one vertex of degree $2m$ and $2m$ vertices of degree 2.

Theorem 3.2. *If a graph G has the property that every two vertices have exactly one common neighbor, then it is a windmill graph.*

Proof. We will show that there must be a vertex c that is a neighbor to every other vertex, ie $N(c) = V(G)\backslash\{c\}$. Then it will follow that G is a windmill graph, because for any other vertex u, its unique common neighbor with c will be the other corner of the arm of the windmill.
We will write $f(x, y)$ for the unique common neighbor of vertices x and y. Note that a subgraph $C_4 = K_{22}$ would violate the condition.
Suppose there is no such c and $|V(G)| \geq 4$ (the smaller cases are easily checked). We will get a contradiction in 3 steps.

• G **is k-regular:** Take two non-adjacent vertices u, v and their unique common neighbor $w = f(u, v)$. Then because $uv \notin E$, $f(u, w)$ is not v, and $f(w, v)$ is not u. Also $f(u, w) \neq f(w, v)$ because that would give a C_4. So $uwf(u, w)$ and $wvf(w, v)$ form two triangles sharing only the vertex w.
For any other neighbor t of u, $f(t, v)$ is not one of the 5 previous vertices, because in each case we would get a C_4. So we can pair up t with a neighbor of v. We can repeat this for any other $t' \in N(u)$, and then also $f(t', v) \neq f(t, v)$ and $f(t', v) \neq t$, because that would give a C_4. Repeating this we can pair off the remaining neighbors of u and v, hence $d(u) - 2 = d(v) - 2$, showing that $d(u) = d(v)$. So any 2 non-adjacent vertices have the same degree. But any vertex other than w is non-adjacent to at least one of u, v, so these all have the same degree k. And by the assumption at the start, w is not a neighbor of all vertices, so it will also have degree k.

• $|V(G)| = k^2 - k + 1$**:** Pick any vertex u. The defining property implies that every vertex is a neighbor of a neighbor of u (including u itself), and in only one way. So

$$|V(G)| = 1 + \sum_{v \in N(u)} (d(v) - 1) = 1 + \sum (k - 1) = k^2 - k + 1.$$

• **Contradiction:** If A is the adjacency matrix of G, then we have

$$A^2 = (k - 1)I + J.$$

Indeed, the condition says that between every 2 vertices there is exactly 1 path of length 2, which means off the diagonal A^2 has everywhere 1. And because G is k-regular, there are k paths of length 2 from a vertex to itself, so the diagonal entries of A^2 are k.
Then from $\mathrm{Spec}(J) = (k^2 - k + 1)^1 (0)^{n-1}$ and $\mathrm{Spec}((k - 1)I) = (k - 1)^n$ it follows that

$$\mathrm{Spec}(A^2) = (k^2)^1 (k - 1)^{n-1},$$

so

$$\mathrm{Spec}(A) = (k)^1 (\sqrt{k - 1})^a (-\sqrt{k - 1})^b,$$

for some integers $a, b \geq 0$ with $a + b = n - 1$. Therefore

$$0 = \mathrm{Tr}(A) = k + a\sqrt{k - 1} - b\sqrt{k - 1},$$

which gives $(b - a)\sqrt{k - 1} = k$. But a, b, and k are integers, so $k - 1 = l^2$ for an integer l. Then l divides both $k - 1$ and k, hence $l = 1$, so $k = l^2 + 1 = 2$ and $|V(G)| = k^2 - k + 1 = 3$, which we excluded above. Contradiction. □

Intersecting Set Systems

The following theorem is a classic from extremal set theory, first proved in 1961 by Paul Erdős, Chao Ko, and Richard Rado. It is very similar to the theorems about set systems with certain intersection properties that we proved using the dimension bound, but it seems not to have a proof like that. There are many different proofs, some of them quite short, but the following one, using eigenvalues, is particularly straightforward once you know two general lemmas. It can also be used for extensions and generalizations for which the other proofs are no use.

We say that a set system $(X, \mathcal{S})$ is *k-uniform* if all its sets have size k, and it is *intersecting* if $|S \cap T| \geq 1$ for all $S, T \in \mathcal{S}$. How large can a set system be if it has both properties? Clearly, if $k > |X|/2$, then $\mathcal{S} = [n]^{(k)}$ is the largest. For $k \leq |X|/2$, we can pick some fixed $x \in X$ and take

$$\mathcal{S} = \{S \in [n]^{(k)} : S \ni x\}.$$

This system has size $\binom{|X|-1}{k-1}$, and the theorem says that this is the best possible.

Theorem 3.3 (Erdős-Ko-Rado)**.** *Let $|X| = n$ and $k \leq n/2$. If $(X, \mathcal{S})$ is k-uniform and intersecting, then*

$$|\mathcal{S}| \leq \binom{n-1}{k-1}.$$

In the proof we will use the *Kneser graph* $K(n, k)$ defined by

$$V(K(n,k)) = [n]^{(k)}, \quad E(K(n,k)) = \{ST : S \cap T = \}.$$

It is clearly $\binom{n-k}{k}$- regular. We have already seen $K(5, 2)$, which is the Petersen graph. An intersecting uniform set system is the same as an independent set in this graph, so the theorem says exactly that, if $k \leq n/2$, then

$$\alpha(K(n,k)) \leq \binom{n-1}{k-1}.$$

So Lemma 2.6 will come in handy here. The only other thing we need is the least eigenvalue of the Kneser graph. We will simply state here without proof what it is, because determining it is somewhat lengthy and unpleasant.

Lemma 3.4. *The Kneser graph $K(n, k)$ has least eigenvalue*

$$\lambda_{least} = -\binom{n-k-1}{k-1}.$$

Proof of Erdős-Ko-Rado. Because the graph is $\binom{n-k}{k}$-regular, we have $\lambda_1 = \binom{n-k}{k}$. Now we just calculate

$$\alpha(K(n,k)) \leq \frac{|V(K(n,k))|}{1 - \frac{\lambda_1}{\lambda_{\text{least}}}} = \frac{\binom{n}{k}}{1 + \frac{\binom{n-k}{k}}{\binom{n-k-1}{k-1}}} = \frac{\binom{n}{k}}{1 + \frac{n-k}{k}} = \frac{k\binom{n}{k}}{k+n-k} = \frac{k}{n}\binom{n}{k} = \binom{n-1}{k-1}.$$

□

Equiangular Lines

For the next theorem we return to equiangular sets of lines – sets of concurrent lines such that every two have the same (smaller) angle between them. We proved earlier with the dimension bound that the size of such a set in $\mathbb{R}^d$ is at most $\binom{d+1}{2}$, and we saw that this is tight for $d = 2$ (lines with 60-degree angles) and $d = 3$ (lines through the vertices of a regular isocahedron). It is not tight for $d = 4, 5, 6$, although we won't prove that here. But weirdly it is tight again for $d = 7$. In a bonus problem about spherical two-distance sets, we saw such a tight example, but there the vectors where just given without any motivation or explanation of how one might come up with them. Here we will we use eigenvalues to construct this example, in a way that (hopefully) you can imagine coming up with (given the right ingredients).

Theorem 3.5. *There exists an equiangular set of 28 lines in* $\mathbb{R}^7$.

Proof/Explanation. From basic linear algebra we have a correspondence between n linearly independent vectors $x_i \in \mathbb{R}^d$, which we will view as columns in a $d \times n$ matrix $X = [x_1 \cdots x_n]$, and a positive semidefinite $n \times n$ matrix M such that $M = X^T X$, so $M_{ij} = x_i^T x_j$ and $\text{rk}(M) = \text{rk}(X)$. Now suppose the x_i are unit vectors and represent equiangular lines, so we have $x_i^T x_j = \pm\alpha$ for some fixed α (the smaller angle is unique, but the inner product can be $\pm$ the cosine of the smaller angle). Then M has the particular form

$$M = I + \alpha S,$$

with S having 0s on the diagonal, and 1 or -1 everywhere else. This matrix S should remind us of an adjacency matrix of a graph: If we set

$$A = \frac{1}{2}(J - I - S),$$

then A is a 0/1-matrix with 0s on the diagonal, so corresponds to a graph.

This we can now do in reverse: Given a graph G, let A be its adjacency matrix, set $S = J - I - 2A$, choose α so that $M = I + \alpha S$ is positive semidefinite, find the corresponding X such that $X^T X = M$, and then finally the columns of X should represent equiangular lines. The number of such columns will be n, and the dimension that the columns lie in will equal the rank of M. So we need to choose the graph so that $d = \text{rk}(M)$ is small compared to n, and we need to choose α so that $I + \alpha S$ is positive semidefinite.

To see how to choose α, let μ be the least eigenvalue of S. Then if we choose

$$\alpha = -1/\mu,$$

$I + \alpha S$ will have all eigenvalues nonnegative, which implies that it is positive semidefinite.

To see how to choose the graph, let its eigenvalues be $\lambda_1, \ldots, \lambda_n$. The eigenvalues of $S = J - I - 2A$ are then $n - 1 - 2\lambda_1$ and $-1 - 2\lambda_i$ for $i \geq 2$, by a calculation that we have seen several times before. In particular, the least eigenvalue of S will probably be $-1-2\lambda_i$. Then we want $\text{rk}(M)$ to be small, and

$$\text{rk}(M) = n - \text{m}_M(0) = n - \text{m}_S(\mu) \approx n - m_G(\lambda_2).$$

Here $\text{m}_M(0)$ is the multiplicity of the eigenvalue 0 of M, which equals the dimension of its kernel; μ is the least eigenvalue of S; and the last equality need not be exactly true, but we can use it to guide our choice of G.

So we know what we want: A graph G with the multiplicity of λ_2 especially large. We can now look through the examples that we've seen, and we'll find that complements of line graphs (which we saw on a problem set) have this property. In particular, $\overline{L(K_8)}$ stands out, because

$$\mathrm{Spec}(\overline{L(K_8)}) = (15)^1(1)^{20}(-5)^7 \quad \Rightarrow \quad \mathrm{Spec}(S) = (9)^7(-3)^{21}.$$

Here by lucky coincidence we have $n-1-2\lambda_1 = -1-\lambda_2$, making the multiplicity of $\mu = -3$ a little larger.

So choosing $G = \overline{L(K_8)}$ we get

$$\mathrm{rk}(M) = \mathrm{rk}(I + \frac{1}{3}S) = n - m_S(-3) = 28 - 21 = 7.$$

Now we can find the 7×28 matrix X such that $X^T X = M$, and its 28 columns will give equiangular lines in $\mathbb{R}^7$. We won't actually do the tedious but straightforward calculation needed to determine them, but we have proven that they exist. □

Turán's Theorem

We will end with a reproof of one of the main theorems of extremal graph theory, first proved by Paul Turán in 1941. We proved it earlier using probability, and there are several purely combinatorial proofs.

Recall that the *clique number* of a graph is the number of vertices in a largest clique (complete subgraph) of the graph.

Theorem 3.6 (Turán). *If a graph with n vertices and e edges has clique number ω, then*

$$e \le \frac{1}{2}\left(\frac{\omega - 1}{\omega}\right) n^2.$$

Proof. Let λ_1 be the largest eigenvalue of the graph G. We proved above that $\lambda_1 \ge d_{\text{avg}}$, and combining that with Lemma 3.7 below we get

$$\frac{2e}{n} = d_{\text{avg}} \le \lambda_1 \le \frac{\omega - 1}{\omega} n.$$

□

So, like for Erdős-Ko-Rado, the proof comes down to proving an eigenvalue bound for the clique number. The Lemma says that $\omega \ge \frac{n}{n-\lambda_1}$, but unfortunately we haven't really seen examples with interesting clique numbers to test it on.

Lemma 3.7. *The largest eigenvalue λ_1 and the clique number satisfy*

$$\lambda_1 \le \frac{\omega - 1}{\omega} n.$$

Proof. Let A_G be the adjacency matrix of the graph. We will consider the following function of vectors $x \in \mathbb{R}^n$:

$$x^T A_G x = \sum_{uv \in E} x_u x_v.$$

If the graph is complete we can bound this function in terms of $j^T x = \sum x_u$; we will prove this claim at the end.

Claim 1: *If $G = K_t$ is complete, then for all $x \in \mathbb{R}^t$ we have*

$$x^T A_{K_t} x \leq \frac{t-1}{t} \cdot (j^T x)^2.$$

We will use the following notation for the *support* of a vector $x \in \mathbb{R}^{|V(G)|}$:

$$s(x) = \text{the graph induced by } u \in V(G) \text{ with } x_u \neq 0.$$

Now suppose that G has a complete subgraph $H = K_T$. Then for any vector x with $s(x) = H$ we have

$$\frac{x^T A_G x}{(j^T x)^2} = \frac{x^T A_H x}{(j^T x)^2} \leq \frac{t-1}{t}.$$

We claim that this is actually true for all vectors.

Claim 2: *The maximum of $\frac{x^T A_G x}{(j^T x)^2}$ over all $x \in \mathbb{R}^n$ is attained on some vector y with $s(y)$ a complete graph.*

With this claim we can prove the Lemma. Take a unit eigenvector v for λ_1. Then

$$\frac{v^T A_G v}{(j^T v)^2} = \frac{v^T(\lambda_1 v)}{(j^T v)^2} = \lambda_1 \frac{v^T v}{(j^T v)} \geq \frac{\lambda_1}{n},$$

using Cauchy-Schwarz again. On the other hand, if the clique number of the graph is ω, then it has a subgraph $H = K_\omega$, so for some x with $s(x) = H$ we have

$$\frac{v^T A_G v}{(j^T v)^2} \leq \frac{x^T A_H x}{(j^T x)^2} \leq \frac{\omega - 1}{\omega}.$$

This proves the Lemma. $\square$

It remains to prove the two claims.

Proof of Claim 1.

$$x^T A_{K_t} x = \sum_{(u,v): u \neq v} x_u x_v = \sum_u x_u \left(\sum_{v \neq u} x_v \right) = \sum_u x_u (j^T x - x_u)$$

$$= j^T x \left(\sum x_u \right) - \left(\sum x_u^2 \right) = (j^T x)^2 - x^T x \leq \frac{t-1}{t} (j^T x)^2,$$

where we used the Cauchy-Schwarz inequality to get $(j^t x)^2 \leq j^T j \cdot x^T x = t \cdot x^T x$. $\square$

Proof of Claim 2. Let y be a vector that maximizes $\frac{x^T A_G x}{(j^T x)^2}$, scaled so that $j^T y = 1$. We will show that if $y_u, y_v \neq 0$ for some $uv \notin E$, then we can make one of these y_u, y_v zero without changing $y^T A y$. If we keep repeating this, we must end up with a maximizing y such that $uv \in E$ whenever $y_u, y_v \neq 0$, which means that $s(y)$ is complete.
So suppose $y_u, y_v \neq 0$ but $a_{uv} = 0$. Then

$$y^T A y = \sum a_{uv} y_u y_v = y_u \sum_{w \neq v} a_{uw} y_w + y_v \sum_{w \neq u} a_{vw} y_w,$$

and without loss of generality we can assume $\sum_{w \neq v} a_{uw} y_w \geq \sum_{w \neq u} a_{vw} y_w$, which gives

$$y^T A y \leq (y_u + y_v) \cdot \left(\sum_{w \neq v} a_{uw} y_w \right) = z^T A z,$$

if we define z by $z_u = y_u + y_v$, $z_v = 0$, and $z_w = y_w$ for $w \neq u, v$. Therefore z is also maximizing (and in fact the inequality is an equality), and it has one more entry zero, as we wanted. $\square$