

Spectral graph theory

Uri Feige

January 2010

1 Background

With every graph (or digraph) one can associate several different matrices. We have already seen the vertex-edge incidence matrix, the Laplacian and the adjacency matrix of a graph. Here we shall concentrate mainly on the adjacency matrix of (undirected) graphs, and also discuss briefly the Laplacian. We shall show that spectral properties (the eigenvalues and eigenvectors) of these matrices provide useful information about the structure of the graph. It turns out that for regular graphs, the information one can deduce from one matrix representation (e.g., the adjacency matrix) is similar to the information one can deduce from other representations (such as the Laplacian). We remark that for nonregular graphs, this is not the case, and the choice of matrix representation may make a significant difference. We shall not elaborate on this issue further, as our main concern here will be either with regular or nearly regular graph.

The adjacency matrix of a connected undirected graph is nonnegative, symmetric and irreducible (namely, it cannot be decomposed into two diagonal blocks and two off-diagonal blocks, one of which is all-0). As such, standard results in linear algebra, including the Perron-Frobenius theorem, imply that:

1. All its eigenvalues are real. Let us denote them by $\lambda_1 \geq \lambda_2 \dots \geq \lambda_n$. (Equality between eigenvalues corresponds to eigenvalue multiplicity.)
2. Eigenvectors that correspond to different eigenvalues are orthogonal to each other.
3. The eigenvector that corresponds λ_1 is all positive.
4. $\lambda_1 > \lambda_2$ and $\lambda_1 \geq |\lambda_n|$.

An eigenvector with eigenvalue λ can be interpreted as associating values (the coordinate entries of the eigenvector) with the vertices of the graph, such that each vertex has a value that is the sum of its neighbors, scaled by $1/\lambda$.

Observe that the trace of the adjacency matrix is 0 and hence $\sum \lambda_i = 0$, implying in particular that $\lambda_n < 0$.

There is useful characterization of eigenvalues by Rayleigh quotients. Let v_1, \dots, v_n be an orthonormal basis of eigenvectors. For a nonzero vector x , let $a_i = \langle x, v_i \rangle$ and hence $x = \sum a_i v_i$. Observe that:

$$\frac{x^t Ax}{x^t x} = \frac{\sum \lambda_i (a_i)^2}{\sum (a_i)^2}$$

This implies that $\lambda_n \leq \frac{x^t Ax}{x^t x} \leq \lambda_1$. Moreover, if x is orthogonal to v_1 then $\frac{x^t Ax}{x^t x} \leq \lambda_2$.

For a d -regular graph, $\lambda_1 = d$, and the corresponding eigenvector is the all 1 vector. If the graph is disconnected, then its adjacency matrix decomposes into adjacency matrices of its connected components. In this case (and again, we assume here d -regularity), the multiplicity of the eigenvalue d is equal to the number of connected components. This indicates that a small gap between λ_1 and λ_2 corresponds to the graph having small cuts. There is also a converse, saying that a large gap between λ_1 and λ_2 implies that the graph has no small cuts in some exact sense. Namely, it is an expander.

To understand intuitively this expansion property (in fact, a related property as we shall also consider λ_n here), consider a random walk starting at an arbitrary distribution x . After one step, the distribution is Ax (if G is regular) normalized to 1. By representing x in the basis of eigenvectors, and assuming that λ_1 is much larger (say, twice as large) than $\max[\lambda_2, -\lambda_n]$ it is clear that in $O(\log n)$ steps, the walk mixes. Hence there cannot be any “bottlenecks” (sets of vertices with only few outgoing edges), as then the walk would not mix.

If the graph is not regular, the largest eigenvalue is at least as large as the average degree (by taking Rayleigh quotient for all 1 vector) and not larger than the maximum value c for which the graph has a c -core (by sorting according to largest eigenvalue - no vertex has more than λ_1 higher neighbors). For a star, the largest eigenvalue is $\sqrt{n-1}$.

If the graph is bipartite, nonzero eigenvalues come in pairs (flipping the entries of vertices of one side in the eigenvectors). Hence $\lambda_n = -\lambda_1$. If $|\lambda_n|$ is much smaller than λ_1 , this implies that the graph does not have very large cuts (that cut almost all the edges). The converse is not true, in the sense that $|\lambda_n|$ might be close to λ_1 , and still the graph has no large cuts. This may happen for example if a d -regular graph G contains a small subgraph that is nearly a complete bipartite graph on $2d$ vertices. This small subgraph affects λ_n (as can be seen by taking Rayleigh quotients) but is insignificant in terms of the existence of large cuts in G . Extensions of the eigenvalue based technique to semidefinite programming (a topic beyond the scope of this course) give an approximation algorithm for max-cut with a very good approximation ratio (roughly 0.87856).

The following theorem is due to Hoffman, and illustrates the connection between eigenvalues and graph properties.

Theorem 1 *Let G be a d -regular n -vertex graph and let λ_n be the most negative eigenvalue of its adjacency matrix A . Then the size of its largest independent is at most*

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

Proof: Let S be an independent set in G . Consider the vector x with value $n - |S|$ on S and value $-|S|$ on the remaining vertices. Then we use the Rayleigh quotient to bound λ_n from above.

$$\frac{x^t Ax}{x^t x} = \frac{|S|^2(nd - 2|S|d) - 2d|S|(n - |S|)|S|}{|S|(n - |S|)^2 + (n - |S|)|S|^2} = \frac{-n|S|^2d}{n|S|(n - |S|)} = \frac{-|S|d}{n - |S|}$$

The proof follows by rearranging. \square

The eigenvalues of A^2 are $(\lambda_i)^2$. The trace of A^2 implies that $\sum(\lambda_i)^2 = \sum d_i$. For d -regular graphs, this implies that the average absolute value of eigenvalues is at most \sqrt{d} . For random regular graphs, or random graphs that are nearly regular, it is in fact the case that with high probability over the choice of graph, all but the largest eigenvalue have absolute value $O(\sqrt{d})$. This can be proved by the trace method (considering higher even powers of A). An alternative proof can be based on considering the form $x^t A x$ for all possible unit vectors x orthogonal to v_1 . (There are infinitely many such vectors so a certain discretization needs to be used.)

If G is regular or nearly regular, the adjacency matrix is most helpful. If the graph is irregular, the Laplacian may be more informative. Note that for regular graphs, the spectra of these two matrices are easily related. For general graphs, the Laplacian gives the quadratic form $\sum_{(i,j) \in E} (x_i - x_j)^2$. Hence it is positive semidefinite, with the all 1 vector as a unique eigenvector of eigenvalue 0 (if the graph is connected). The next eigenvalue is known as the Fiedler value and the associated eigenvector as the Fiedler vector. It gives a useful heuristic for spectral graph partitioning. The Fiedler vector and the next eigenvector form a convenient coordinate system for graph drawing (see [3]).

Much of what we said about the spectrum of regular graphs applies to graphs that are nearly regular. In establishing that a random graph is nearly regular, it is useful to have some bounds of the probability that a degree of a vertex deviates from its expectation. These bounds are quantitative versions of the law of large numbers and of the central limit theorem (that says that sums of independent random variables converge to the normal distribution), and we present one representative such bound.

Theorem 2 *Let $p_1, \dots, p_n \in [0, 1]$ and set $p = \frac{p_1 + \dots + p_n}{n}$. Assume that X_1, \dots, X_n are mutually independent random variables with $\Pr[X_i = 1 - p_i] = p_i$ and $\Pr[X_i = -p_i] = 1 - p_i$. Let $X = X_1 + \dots + X_n$. (Observe that $E[X_i] = 0$, and $E[X] = 0$.) Then*

$$\Pr[X > a] < e^{-a^2/2pn + a^3/2(pn)^2}$$

and

$$\Pr[X < -a] < e^{-a^2/2pn}$$

The theorem above is typically used when $p \leq 1/2$. Note that the bounds for $X < -a$ seem stronger than those for $X > a$. They can be used for $X > a$ after replacing p by $1 - p$. Namely, $\Pr[X > a] < e^{-a^2/2(1-p)n}$. Note also that the bounds can be tightened in certain cases. See [1] for more details.

2 Refuting random 4SAT

We illustrate here an application (due to Goerdt and Krivelevich [?]) of spectral techniques to a seemingly unrelated problem.

A 4CNF formula contains clauses, where each clause contains four literals, and each literal is either a Boolean variable or its negation. An assignment satisfies the formula if in each clause at least one literal is set to true. Checking whether a given 4CNF formula

is satisfiable is NP-complete. Namely, it is in NP (if the answer is yes there is a satisfying assignment certifying this), and it is NP-hard. Checking whether a given 4CNF formula is not satisfiable is coNP-complete. Hence there are no witnesses for non-satisfiability unless $\text{coNP}=\text{NP}$. This applies for worst case instances. Here we study the average case (for some natural distribution).

Consider the distribution $F_{n,m}$ of random 4CNF formulas with n variables and m clauses. The formula is random in the sense that every literal in the formula is chosen independently at random. (We may further require that clauses are distinct and that variables in a clause are distinct, but this has negligible effect on the results that we present.) Observe that given any assignment to the variables, a random clause is satisfied by this assignment with probability exactly $15/16$ (every literal has probability $1/2$ of getting polarity opposite to the assignment). Hence the probability that it satisfies the random formula is $(15/16)^m$, and the expected number of satisfying assignments is exactly $2^n(15/16)^m$. Hence when m is sufficiently large (say, $m = 16n$) it is very likely that the formula is not satisfiable. We shall be interested in developing refutation heuristics that run in polynomial time and combine the following two properties:

1. *Perfect Soundness.* If the heuristic outputs *not satisfiable* then indeed the 4CNF formula is not satisfiable.
2. *Average completeness.* For most non-satisfiable formulas from a given distribution of formulas (this distribution will be specified shortly), the heuristic answers *not satisfiable*.

The distribution that we shall consider is $F_{n,m}$ with m very large compared to n . For fixed m (above the threshold for satisfiability) the heuristic refutation problem becomes easier as m grows (because refuting a subformula implies refuting the whole formula). Hence we shall consider first very large values of m , and then address smaller values of m .

When $m = \Omega(n^4)$ refutation becomes very easy. With high probability, there is a set of four variables that provide the support to a super constant number of clauses. With high probability, all possible polarity combinations (there are 16 of them) appear, and hence any assignment to these four variables leaves at least one clause unsatisfied. It is easy to detect this subformula of 16 clauses in polynomial time.

When $m = n^3$, we have the following refutation algorithm. Pick two variables (say x_1 and x_2). Try all possible assignments to these two variables (four combinations altogether). For each such assignment, simplify the resulting 4CNF formula by removing satisfied clauses, and removing unsatisfied literals. The resulting formula is expected to have roughly $3n$ 2CNF clauses. (Each 4CNF clause has 6 pairs of variables, there are $\binom{n}{2}$ distinct pairs, and there are 4 assignments.) The probability that they are satisfiable is at most $2^n(3/4)^{3n} \ll 1/8$ (for sufficiently large n). (In fact, it is known that a random 2CNF formula with $(1 + \epsilon)n$ clauses is already unlikely to be satisfiable, but we are not interested in optimizing the parameters for the approach described here, so we will not use this fact.) It follows that with high probability, for each of the four assignments simultaneously, the corresponding 2CNF formula is not satisfiable. As satisfiability for 2CNF can be tested in polynomial time, we get a refutation heuristic for 4SAT (when $m = n^3$).

Now let us consider the case $m = cn^2$ for large enough c . To simplify the presentation (though with tighter analysis this simplification is not needed), let us give ourselves more slackness and view c not as a constant, but rather as n^ϵ (where $\epsilon > 0$ and n is sufficiently large).

Given the 4CNF formula ϕ , partition it into three subformulas. ϕ^+ contains only those clauses in which all literals are positive, ϕ^- contains only those clauses in which all literals are negative, and ϕ' contains the remaining clauses. We completely ignore ϕ' , and construct two graphs, G^+ based on ϕ^+ , and G^- based on ϕ^- . We describe the construction of G^+ , and the construction of G^- is similar.

The vertex set V contains $\binom{n}{2}$ vertices, where each vertex is labelled by a distinct pair of distinct variables. For every clause in ϕ^+ (that we assume contains 4 distinct variables), put an edge in G^+ between the vertex labelled by the first two variables in the clause and the vertex labelled by the last two variables in the clause.

Lemma 3 *If ϕ is satisfiable, then at least one of the two graphs G^+ and G^- has an independent set of size at least $\binom{n/2}{2} \simeq |V|/4$.*

Proof: Consider an arbitrary satisfying assignment for ϕ , let S^+ be the set of variables assigned to true, and let S^- be the set of variables assigned to false. Consider the set of $\binom{|S^-|}{2}$ vertices in G^+ labelled by pairs of vertices from S^- . They must form an independent set because ϕ cannot have a clause containing only variables from S^- in which all literals are positive. Likewise, G^- has an independent set of size at least $\binom{|S^+|}{2}$. As $\min[|S^+|, |S^-|] \geq n/2$, the proof follows. \square

Observe that for $F_{n,m}$ with $m = cn^2$, both G^+ and G^- are random graphs with roughly $m/16 \simeq c|V|/8$ edges, and hence average degree $c/4$. For large enough c , they are nearly regular (the standard deviation of degrees is $O(\sqrt{c})$ which becomes negligible as c grows). Hence it will be the case that for their adjacency matrices, the eigenvalues satisfy $\lambda_1 \simeq c/4$ and $|\lambda_{|V|}| = O(\sqrt{c})$. Using Theorem 1 (adapted to nearly regular graphs), this can be used in order to certify that neither G^+ nor G^- have independent sets larger than $O(|V|/\sqrt{c})$. For sufficiently large c , this establishes that ϕ cannot have a satisfying assignment.

It is a major open question to design refutation heuristics for $F_{n,m}$ when $m = n^{2-\epsilon}$.

References

- [1] Noga Alon and Joel Spencer. The Probabilistic Method. Wiley-Interscience Series in Discrete Mathematics and Optimization.
- [2] Andreas Goerdt, Michael Krivelevich: Efficient Recognition of Random Unsatisfiable k-SAT Instances by Spectral Methods. STACS 2001: 294-304
- [3] Dan Spielman. Lecture notes in <http://cs-www.cs.yale.edu/homes/spielman/>.