CS 860 Spectral graph theory, Spring 2019, Waterloo.

Lecture 13: Interlacing polynomials

We go through some background on interlacing polynomials and real stable polynomials, and use them to develop a new probabilistic method, which will be used to solve some important problem in next lectures.

Spectral sparsification and Kadison-Singer problem

It was observed that the linear-size spectral sparsification result looks similar to the following conjecture by Weaver, which is known to be equivalent to the Kadison-Singer problem, whose positive resolution would have implications in several areas of mathematics.

Weaver's conjecture There exist positive constants α and ε so that for every m and n and every set of vectors $v_1, \ldots, v_m \in \mathbb{R}^n$ such that $\|v_i\|^2 \leq \alpha$ for all i and $\sum_i v_i v_i^T = I$, there exists a partition of $\{1, \ldots, m\}$ into two sets S_1 and S_2 so that $\|\sum_{i \in S_j} v_i v_i^T\| \leq 1 - \varepsilon$ for $j \in \{1,2\}$.

Note that since $\sum_{i \in S_1} v_i v_i^T + \sum_{i \in S_2} v_i v_i^T = I$, the conclusion $\|\sum_{i \in S_1} v_i v_i^T\| \leq 1-\varepsilon$ is equivalent to

$\varepsilon I \preccurlyeq \sum_{i \in S_1} v_i v_i^T \preccurlyeq (1-\varepsilon) I$, and so the vectors in S_1 is a spectral approximation of I.

In the BSS theorem in L12, the task was to find scalars with few non-zeros so that

$(1-\varepsilon) I \preccurlyeq \sum_i w_i v_i v_i^T \preccurlyeq (1+\varepsilon) I$, or equivalently $\frac{(1-\varepsilon)}{2} I \preccurlyeq \sum_i \frac{w_i}{2} v_i v_i^T \preccurlyeq \frac{(1+\varepsilon)}{2} I$.

If all the $\frac{w_i}{2}$ are either zero or one, then it would have given a positive resolution to Weaver's conj.

This is not always possible, however, since if there is a long vector v_i (say $\|v_i\| > 1-\varepsilon$), then setting w_i to be zero or one would violate the minimum eigenvalue or maximum eigenvalue bound.

This is why there is an additional condition $\|v_i\|^2 \leq \alpha$ in Weaver's conjecture, and with this we want to set the scalars to be zero or one (but not arbitrary real values), so that we get the stronger conclusion that the vectors can be partitioned into two groups.

For graph sparsification, the question in Weaver's conjecture corresponds to finding an unweighted sparsifier.

One known result is by Karger, who showed that if the min-cut size is $\Omega(\log n)$, then uniform sampling works to produce an unweighted cut sparsifier.

Recall the reduction from spectral sparsification to finding spectral approximation of the identity matrix. the length $\|v_i\|^2$ is equal to the effective resistance of the i-th edge in the original graph.

So, Weaver's conjecture in the spectral sparsification setting is asking if the maximum effective resistance of an edge is at most α, then there is a partitioning of the edges into two groups so that the subgraph formed by each group is a (somewhat) good spectral approximation of the original graph.

Some examples of graphs with maximum effective resistance small are expander graphs or edge-transitive graphs (since effective resistance will be the same for every edge. e.g. hypercubes, Cayley graphs, etc).

One can apply matrix Chernoff bounds to this problem, and it will work for $\alpha = O(1/\log^2 n)$ with high probability, but this is not enough for Weaver's conjecture.

The approach by Batson-Spielman-Srivastava heavily depends on a careful choice of scalars and also does not seem applicable here.

Probabilistic formulation

To resolve Weaver's conjecture, Marcus, Spielman and Srivastava proved a probabilistic statement that implies Weaver's conjecture.

Theorem Let $v_1, v_2, \ldots, v_m \in \mathbb{R}^n$ be independent random vectors with finite support such that

$$E\left[\sum_{i=1}^{m} v_i v_i^T\right] = I_n \quad \text{and} \quad E\left[\|v_i\|^2\right] \le \varepsilon \quad \text{for } 1 \le i \le m.$$

Then $$\Pr\left[\left\|\sum_{i=1}^{m} v_i v_i^T\right\| \le (1+\sqrt{\varepsilon})^2\right] > 0.$$

Reduction

Weaver's conjecture is about partitioning and the above statement is about sum of random matrices. There is a simple reduction from the former to the latter.

For each vector $u_i \in \mathbb{R}^n$ in Weaver's problem, we create a random vector $v_i \in \mathbb{R}^{2n}$ with two choices, so that $v_i = \sqrt{2}\begin{pmatrix} u_i \\ 0 \end{pmatrix}$ with probabity $\frac{1}{2}$ and $v_i = \sqrt{2}\begin{pmatrix} 0 \\ u_i \end{pmatrix}$ with probability $\frac{1}{2}$,

such that the first choice corresponds to putting u_i into the first group and the second choice corresponds to putting u_i into the second group.

Then, $$E[v_i v_i^T] = \frac{1}{2}\left(2\begin{pmatrix} u_i \\ 0 \end{pmatrix}\begin{pmatrix} u_i^T & 0 \end{pmatrix} + 2\begin{pmatrix} 0 \\ u_i \end{pmatrix}\begin{pmatrix} 0 & u_i^T \end{pmatrix}\right) = \begin{pmatrix} u_i u_i^T & 0 \\ 0 & u_i u_i^T \end{pmatrix},$$

and thus $$\sum_{i=1}^{m} E[v_i v_i^T] = \sum_{i=1}^{m}\begin{pmatrix} u_i u_i^T & 0 \\ 0 & u_i u_i^T \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^{m} u_i u_i^T & 0 \\ 0 & \sum_{i=1}^{m} u_i u_i^T \end{pmatrix} = \begin{pmatrix} I_n & 0 \\ 0 & I_n \end{pmatrix} = I_{2n}.$$

Similarly, $$E[\|v_i\|^2] = E[v_i^T v_i] = \frac{1}{2}\left(2\begin{pmatrix} u_i^T & 0 \end{pmatrix}\begin{pmatrix} u_i \\ 0 \end{pmatrix} + 2\begin{pmatrix} 0 & u_i^T \end{pmatrix}\begin{pmatrix} 0 \\ u_i^T \end{pmatrix}\right) = 2\|u_i\|^2 \le 2\alpha.$$

By the above statement, there exists a choice of v_i such that $\|\sum_{i=1}^m v_i v_i^T\| \le (1+\sqrt{2\alpha})^2$

We put i into S_1 if we select the first choice for i, and put i into S_2 otherwise.

Then the conclusion implies that $\left\| \begin{pmatrix} 2\sum_{i\in S_1} u_i u_i^T & 0 \\ 0 & 2\sum_{i\in S_2} u_i u_i^T \end{pmatrix} \right\| \le (1+\sqrt{2\alpha})^2$

and thus $\|\sum_{i\in S_j} u_i u_i^T\| \le \frac{1}{2}(1+\sqrt{2\alpha})^2$.

So, definitely when α is small enough (say $\alpha \le \frac{1}{8}$), then $\frac{1}{2}(1+\sqrt{2\alpha})^2 < 1$, and Weaver's conjecture follows.

It is quantatively stronger than Weaver's conjecture as when α is small enough, we can bound how far it is from $\frac{1}{2}I$, the ideal partitioning, and it will be useful in applications.

New probabilistic method

Marcus, Spielman and Srivastava proved the theorem using a probabilistic method - but in an unusual way.

In standard probabilistic method, we compute the expectation of a random variable $E[X]$, and then conclude that there is an outcome in the sample space with value at most/least $E[X]$.

In this problem, the natural random variable would be $\|\sum_i v_i v_i^T\|$ and so we should compute $E[\|\sum_i v_i v_i^T\|]$.

If the quantity $E[\|\sum_i v_i v_i^T\|]$ is difficult to estimate, then we could consider related quantities such as $\left(tr\left(\sum_i v_i v_i^T\right)^k\right)^{\frac{1}{k}}$.

Instead of working with the random matrix directly, MSS take an unusual route and consider the characteristic polynomial of the random matrix.

Note that $\|\sum_i v_i v_i^T\| = \lambda_{max}\left(\sum_i v_i v_i^T\right) = \text{max-root}\left(\det\left(\lambda I - \sum_i v_i v_i^T\right)\right)$.

The standard probabilistic method says that $\|\sum_i v_i v_i^T\| \le E\left[\text{max-root}\left(\det\left(\lambda I - \sum_i v_i v_i^T\right)\right)\right]$ with positive prob.

MSS proved a surprising variant of the probabilistic method.

Theorem: Let $v_1, v_2, \dots, v_m \in \mathbb{R}^n$ be independent random vectors with finite support.

Then $\|\sum_i v_i v_i^T\| \le \text{max-root}\left[E\left(\det\left(\lambda I - \sum_i v_i v_i^T\right)\right)\right]$ with positive probability.

In general, $E\left[\text{max-root}\left(\det\left(\lambda I - \sum_i v_i v_i^T\right)\right)\right] \ne \text{max-root}\left[E\left(\det\left(\lambda I - \sum_i v_i v_i^T\right)\right)\right]$ and in fact the latter term could be smaller than the former term, so the theorem is not trivial at all.

Characteristic polynomials have not played an important role in much of spectral graph theory.

For example, one disadvantage is that we lost the information about the eigenvectors.

But the proof of MSS uses their algebraic and analytical properties in crucial ways.

- The characteristic polynomials satisfy a number of algebraic identities which make calculating their averages tractable. It will become very convenient when we study its application in Ramanujan graphs.
- The proof of the above probabilistic method uses that the random matrix is the sum of rank-one PSD matrices, whose characteristic polynomials satisfy some very nice interlacing properties that we are going to study next.

In this lecture, we will study some background on interlacing polynomials and real-stable polynomials.

In the next lectures, we will see some amazing applications such as constructing bipartite Ramanujan graphs, proving Weaver's conjecture, constructing thin trees, etc.

Interlacing polynomials

Recall that we would like to study under what conditions it holds that $\|\sum_i A_i\| \le \text{max-root}\left(\mathbb{E}\left(\det(\lambda I - \sum_i A_i)\right)\right)$ with positive probability when A_i are random matrices.

Let's consider the more general question when $\min_i \text{max-root}(f_i) \le \text{max-root}\left(\sum_i f_i\right)$ when f_i are polynomials.

In general, it is usually not true.

Consider $f_1 = (x-1)(x-2)$ and $f_2 = (x-3)(x-4)$.

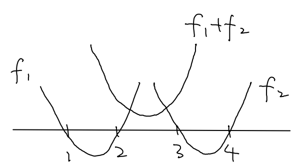

The polynomial f_1+f_2 is not even real-rooted.

Even if f_1+f_2 is real-rooted, the relation does not necessarily hold.

For example, consider $f_1 = (x+5)(x-9)(x-10)$ and $f_2 = (x+6)(x-1)(x-8)$, f_1+f_2 has roots $\approx -5.3, 6.4, 7.4$.

f_1 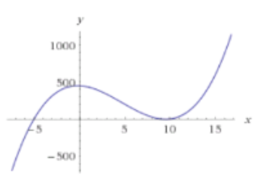 f_2 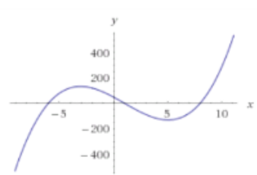 f_1+f_2 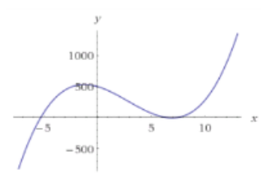

pictures from Wolfram alpha

There are, however, some very nice properties in the polynomials in Weaver's setting, when the random matrix is a sum of rank-one PSD matrices.

Definition (Interlacing) Let f be a degree n polynomial with real roots $\alpha_1 \ge \alpha_2 \ge \dots \ge \alpha_n$ and let g be a degree n polynomial with real roots $\beta_1 \ge \beta_2 \ge \dots \ge \beta_n$ (or g a degree $n-1$ polynomial). We say that g interlaces f if $\alpha_1 \ge \beta_1 \ge \alpha_2 \ge \beta_2 \ge \dots \ge \beta_{n-1} \ge \alpha_n \ge \beta_n$ (just drop β_n if $\deg(g) = n-1$).

Example This is closely related to Weaver's setting although we won't use it directly.

Let $g = \det(\lambda I - A)$ and $f = \det(\lambda I - A - vv^T)$, a rank one update of A.

Then g interlaces f.

The proof follows from the Courant-Fischer theorem in L02 and is left as an exercise.

Common interlacing and comparing max roots

Suppose two real-rooted polynomials with positive leading coefficients are interlacing.

Then it is easy to see that $\min_{i \in \{1,2\}} \text{max-root}(f_i) \le \text{max-root}(f_1 + f_2)$.

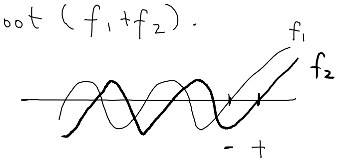

Say $\text{max-root}(f_1) \le \text{max-root}(f_2)$.

Since both f_1 and f_2 have positive leading coefficients,

both are positive in the range $(\text{maxroot}(f_2), \infty)$, and $(f_1 + f_2)(\text{maxroot}(f_2)) > 0$.

We can assume that f_1 and f_2 have no common roots, as otherwise we can divide both polynomials by the common factors, prove that the resulting polynomials are interlacing and this would imply that the original polynomials are interlacing as well

Since f_1 interlaces f_2, the second largest root of f_2 is smaller than the largest root of f_1, and thus $f_2(\text{maxroot}(f_1))$ must be negative, and hence $(f_1 + f_2)(\text{maxroot}(f_1)) < 0$.

By the intermediate value theorem, there must be a root of $f_1 + f_2$ in the range $(\text{maxroot}(f_1), \text{maxroot}(f_2))$, and this proves the claim that $\text{max-root}(f_1) \le \text{max-root}(f_1 + f_2) \le \text{max-root}(f_2)$.

This can be generalized to a set of polynomials in a natural way.

Definition (common interlacing) We say a set of polynomials $f_1, \ldots, f_m$ have a common interlacing if there is a polynomial g which interlaces each f_i.

Equivalently, $f_1, \ldots, f_m$ have a common interlacing if there are disjoint intervals $I_1 \geq I_2 \geq \ldots \geq I_n$ so that the k-th largest root of each f_i is contained in I_k, i.e. $\beta_5 \; I_5 \; \beta_4 \; I_4 \; \beta_3 \; I_3 \; \beta_2 \; I_2 \; \beta_1 \; I_1$

The following lemma follows by applying the intermediate value theorem on each I_k.

Lemma Suppose $f_1, \ldots, f_m$ have a common interlacing where each has a positive leading coefficient.

Let $\lambda_k(f_j)$ be the k-th largest root of f_j, and $\mu_1, \mu_2, \ldots, \mu_m$ be non-negative numbers with $\sum_{i=1}^m \mu_i = 1$.

Then $\min_j \lambda_k(f_j) \le \lambda_k\left(\sum_{j=1}^m \mu_j f_j\right) \le \max_j \lambda_k(f_j)$.

So, if we could show that a set of polynomials have a common interlacing, then we can apply the new probabilistic method to show that one polynomial has small max-root by showing that the "average" polynomial has small max-root.

We will show how to use this approach for the Weaver's conjecture next week.

Today we focus on some general techniques to prove that a set of polynomials has common interlacing.

Common interlacing and real-rootedness

From the lemma in the previous section, if $f_1, \dots, f_m$ are real-rooted and have a common interlacing, then any convex combination of $f_1, \dots, f_m$ is also real-rooted.

It turns out that the converse is also true.

We need the following simple fact for the proof.

Fact $f_1, \dots, f_m$ have a common interlacing if and only if f_i, f_j have a common interlacing $\forall i \neq j$.

proof $(\Rightarrow)$ is trivial.

$(\Leftarrow)$ if every pair has a common interlacing, then in any interval $[c, \infty]$, no polynomial can have two more roots than any polynomial, and so there is a common interlacing. $\square$

Also, we use the following result from complex analysis without proof.

Theorem The roots of a polynomial are continuous functions of its coefficients on the complex plane.

Now we are ready to prove the converse.

Lemma Given $f_1, f_2, \dots, f_m$, if all convex combinations $\sum_{i=1}^{m} \mu_i f_i$ are real-rooted, then $f_1, f_2, \dots, f_m$ have a common interlacing.

Proof By the fact, we only need to prove this for two polynomials.

Again, we will assume without loss of generality that f_1 and f_2 have no common roots.

Let $f_t = (1-t) f_1 + t f_2$ for $t \in [0,1]$.

If we keep track of the roots of f_t from $t=0$ to $t=1$ as a continuous function of t, then each root of f_t is a continuous curve on the complex plane as t varies from 0 to 1 by the theorem.

Since by assumption each f_t is real-rooted, the curve of each root is an interval on the real line.

If two intervals overlap at a point r which is a root of f_1,

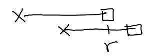

then $0 = f_t(r) = (1-t) f_1(r) + t f_2(r) = t f_2(r)$, and $f_2(r)$ must be zero if $t \in (0,1)$,

contradicting that f_1 and f_2 have no common roots.

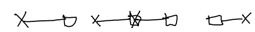

So, the intervals must be disjoint except at endpoints, and so f_1, f_2 have a common interlacing. $\square$

By the lemma, to prove a set of polynomials have a common interlacing (to apply the probabilistic method), it is equivalent to prove that all convex combinations (of two polynomials) are real-rooted.

Henceforth, we will focus on methods to prove that a polynomial is real-rooted.

Real-rooted polynomials

What are some examples of real-rooted polynomials? The following example is familiar.

Example The characteristic polynomial of a real symmetric matrix (more generally Hermitian matrix).

The roots are the eigenvalues of the matrix, and we know from L01 that the eigenvalues are real.

There is a characterization of when a polynomial is real-rooted, but it is not easy to use.

Theorem (Hermite-Sylvester theorem) A polynomial $p(x) = \prod_{i=1}^{n} (x - \lambda_i)$ is real-rooted if and only if the $n \times n$ matrix H with $H_{ij} = \sum_{\ell=1}^{n} \lambda_\ell^{i+j-2}$ (i.e. the $(i+j-2)$-th moment of the roots) is PSD.

Note that the entries of H can be computed from the coefficients in $p(x)$ efficiently, and this gives a polynomial time algorithm to check whether a polynomial is real-rooted.

We will not use this theorem. The proof is left as a (harder) homework problem.

Another way to show that a polynomial $p(x)$ is real-rooted is to start with a known real-rooted polynomial $q(x)$ (e.g. characteristic polynomial of a real symmetric matrix) and show that $p(x)$ can be obtained from $q(x)$ by some real-rootedness preserving operations.

The following are some examples of real-rootedness preserving operations:

- If $p(x)$ is real-rooted, then $p(cx)$ is real-rooted for any $c \in \mathbb{R}$.
- If $p(x)$ is a degree n real-rooted polynomial, then so is $x^n p(\frac{1}{x})$.

- If $p(x)$ is real-rooted, then so is $p'(x)$, the derivative of $p(x)$.

The proofs are left as exercises.

Real stable polynomials

The approach of MSS to prove a polynomial is real-rooted is a generalization of the above idea.

They consider a multivariate generalization of real-rooted univariate polynomial, called real stable polynomials.

To show that a univariate polynomial is real-rooted, they will start from some real stable polynomial and show that the univariate polynomial can be obtained from the multivariate real stable polynomial through a sequence of real stability preserving operations.

Definition (real stable polynomials) A multivariate polynomial $f \in \mathbb{R}[x_1, \ldots, x_m]$ is real stable if there are no roots $(y_1, y_2, \ldots, y_m)$ with $\operatorname{Im}(y_j) > 0$ for all $1 \le j \le m$.

Examples : $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$, $f(x_1, \ldots, x_n) = a_1 x_1 + \ldots + a_n x_n$ where $a_i > 0$ for $1 \le i \le n$.

Non-examples : $f(x_1, x_2) = x_1 - x_2$, $f(x_1, x_2, x_3, x_4) = x_1 x_2 - x_3 x_4$

Note that it is a generalization of real-rootedness for univariate polynomials with real coefficients.

Fact A univariate polynomial $f \in \mathbb{R}[x]$ is real stable if and only if it is real-rooted.

Proof Let $f(x) = \sum_{k=0}^{d} c_k x^k$. The proof follows from the observation that complex roots come in conjugate pair.

Suppose $f(a+ib) = \sum_{k=0}^{d} c_k (a+ib)^k = 0$.

Then $0 = \sum_{k=0}^{d} \overline{c_k (a+ib)^k} = \sum_{k=0}^{d} c_k \overline{(a+ib)^k} = \sum_{k=0}^{d} c_k (\overline{a+ib})^k = \sum_{k=0}^{d} c_k (a-ib)^k = f(a-ib)$.

One of these must have positive imaginary part, contradicting to real stability of f. $\square$

We can check whether a multivariate polynomial is real-stable by checking whether certain derived univariate polynomial is real-rooted. The proof follows from the definition and is left as an exercise.

Lemma A polynomial $f \in \mathbb{R}[x_1, \ldots, x_n]$ is real stable if and only if for any $b \in \mathbb{R}^n_{>0}$ with positive coordinates and $a \in \mathbb{R}^n$, it holds that $f(a + xb)$ is real-rooted (with x as the only variable).

We can use this lemma to check whether some polynomials are real stable.

Example

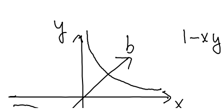

the restriction of the polynomial to $b > 0$ has two real roots, Example

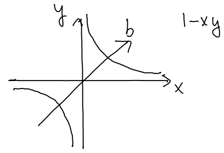

the restriction of the polynomial to $b>0$ has two real roots, so $f(x,y) = 1-xy$ is real stable.

Non-example

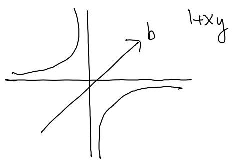

the restriction of the polynomial to $b>0$ has no real roots, so $f(x,y) = 1+xy$ is not real stable.

It is proved that the support of any homogenous multiaffine real stable polynomial form the bases of a matroid, e.g. $x_1x_2+x_3x_4$ is not real stable. This shows some limitation of this class of polynomials.

At least for these lectures, the source of all real stable polynomials come from determinants.

Lemma If $A_1, \dots, A_m$ are positive semidefinite matrices,

then $f(x_0, x_1, \dots, x_m) := \det\left(x_0 I + \sum_{i=1}^{m} x_i A_i\right)$ is a real stable polynomial.

Proof We show that if $\operatorname{Im}(x_i) > 0$ for all $0 \le i \le m$, then the matrix $x_0 I + \sum_{i=1}^{m} x_i A_i$ is of full rank, and hence $\det\left(x_0 I + \sum_{i=1}^{m} x_i A_i\right) \neq 0$, implying real stability.

Let $\vec{v} = \vec{c} + i\vec{d}$ where $\vec{c}$ is the real part and $\vec{d}$ is the imaginary part of $\vec{v}$.

Let $X = x_0 I + \sum_{i=1}^{m} x_i A_i$. Write X as $\operatorname{Re}(X) + i \operatorname{Im}(X)$.

When $\operatorname{Im}(x_j) > 0$ for all $0 \le j \le m$, this implies that $\operatorname{Im}(X) \succ 0$, as $A_i \succcurlyeq 0$ and $I \succ 0$.

We will show that $X\vec{v} = 0$ only if $\vec{c} = \vec{d} = 0$, and hence X is of full rank.

To show this, we show that $(\vec{c} - i\vec{d})^T (\operatorname{Re}(X) + i\operatorname{Im}(X))(\vec{c} + i\vec{d}) = 0$ only if $\vec{c} = \vec{d} = 0$.

Note that $\operatorname{Im}\left[(\vec{c} - i\vec{d})^T(\operatorname{Re}(X) + i\operatorname{Im}(X))(\vec{c} + i\vec{d})\right] = \vec{c}^T \operatorname{Im}(X)\vec{c} + \vec{d}^T \operatorname{Im}(X) \vec{d} = 0$ only if $\vec{c} = \vec{d} = 0$, as $\operatorname{Im}(X) \succ 0$ when all $\operatorname{Im}(x_j) > 0$. This completes the proof. $\square$

Observe that the polynomial in this lemma is similar to the characteristic polynomial of the sum of matrices A_i, those that appear in Weaver's setting, except that it is a multivariate polynomial.

Later, we will start from this multivariate real-stable polynomial to show interlacing properties of those characteristic polynomials in Weaver's setting.

The only missing pieces are the stability preserving operations in the next section.

Real stability preserving operations

There are several real-stability preserving operations, and there are some deep results about these. We just present the proofs of two operations that we need for the Weaver's conjecture, and state others without proofs.

Specialization

This operation will be useful in reducing the number of variables of the multivariate polynomial.

Lemma Let $f(x_1, x_2, \ldots, x_m)$ be a non-zero real-stable polynomial.

Then, for $c \in \mathbb{R}$, $f(c, x_2, \ldots, x_m)$ is a real-stable polynomial.

proof As all coefficients of f are real and c is real, all coefficients of $f(c, x_2, \ldots, x_m)$ are real.

Suppose by contradiction that $f(c, x_2, \ldots, x_m)$ is not real stable.

This means that there exist $y_2, \ldots, y_m$ such that $\mathrm{Im}(y_j) > 0$ for $2 \le j \le m$,

and $f(c, y_2, \ldots, y_m) = 0$.

Consider the polynomial $f(c + i\delta, x_2, \ldots, x_m)$ for some small enough $\delta > 0$ to be chosen.

By the theorem that the roots of polynomials are continuous functions of the coefficients,

for every $\varepsilon > 0$, there exists $\delta > 0$ such that there exists a root $y_2', y_3', \ldots, y_m'$ with

$|y_j - y_j'| < \varepsilon$ for $2 \le j \le m$ and $f(c + i\delta, y_2', \ldots, y_m') = 0$.

We can choose ε small enough so that $\mathrm{Im}(y_j') > 0$ for all $2 \le j \le m$, but this contradicts

the real stability of f, since all coordinates of this root $(c + i\delta, y_2', \ldots, y_m')$ have positive

imaginary part. □

Partial derivative

Lemma For any real t, the polynomial $(1 + t\frac{\partial}{\partial x_1}) f(x_1, x_2, \ldots, x_m)$ is real stable if f is.

proof We substitute $x_2 = y_2, \ldots, x_m = y_m$ with $\mathrm{Im}(y_j) > 0$ into f.

Since f is stable, the resulting univariate polynomial $g(x)$ is stable (note that

it may not be real since y_j are complex).

If we could prove that $g(x) + t g'(x)$ is always stable (assuming g is), then we prove

that $(1 + t\frac{\partial}{\partial x_1}) f(x_1, x_2, \ldots, x_m)$ is stable, because if $(1 + t\frac{\partial}{\partial x_1}) f(x_1, x_2, \ldots, x_m)$ is

not stable, then there exist $y_2, \ldots, y_m$ with $\mathrm{Im}(y_j) > 0$ such that $g(x) + t g'(x)$ not stable.

Since $g(x)$ is stable, it can be written as $c\prod_{j=1}^{n}(x-w_j)$ with $\mathrm{Im}(w_j)\le 0$ for all $1\le j\le m$.

Then $g(x)+tg'(x) = g(x)\left(1+\sum_{j=1}^{n}\frac{t}{x-w_j}\right)$.

For z with $\mathrm{Im}(z)>0$, $g(z)\neq 0$ as g is stable, and furthermore since $\mathrm{Im}(z)>0$,

we have $\mathrm{Im}\left(\frac{t}{z-w_j}\right)<0$ for all j, and thus $1+\sum_{j=1}^{n}\frac{t}{z-w_j}\neq 0$,

proving that $g(x)+tg'(x)$ is stable. □

The following are some other operations that preserve stability.

- Symmetrization: If $p(x_1,x_2,\dots,x_n)$ is real stable, then so is $p(x_1,x_1,x_3,\dots,x_n)$.
- external field: If $p(x_1,x_2,\dots,x_n)$ is real stable, then so is $p(w_1x_1,\dots,w_nx_n)$ for any $w\in\mathbb{R}^n_{>0}$.
- inversion: If $p(x_1,\dots,x_n)$ is real stable and degree of x_i is d_i, then $p(1/x_1,\dots,1/x_n)\prod_{i=1}^{n}x_i^{d_i}$ is real stable.
- differentiation: If $p(x_1,\dots,x_n)$ is real stable, then so is $\partial p/\partial x_1$.

Borcea and Brändén proved a general theorem characterizing what linear differential operators with real coefficients are stability preserving.

Theorem (Borcea, Brändén) For vectors $\alpha,\beta\in\mathbb{N}^n$, we write $x^\alpha = x_1^{\alpha(1)}x_2^{\alpha(2)}\cdots x_n^{\alpha(n)}$ and $\partial^\beta = \partial_{x_1}^{\beta(1)}\partial_{x_2}^{\beta(2)}\cdots\partial_{x_n}^{\beta(n)}$,

and let $D=\sum_{\alpha,\beta\in\mathbb{N}^n} a_{\alpha\beta}\,x^\alpha\partial^\beta$ with $a_{\alpha,\beta}\in\mathbb{R}$ for all α,β.

Then D is stability preserving if and only if $\sum_{\alpha,\beta\in\mathbb{N}^n} a_{\alpha,\beta}\,x^\alpha(-w)^\beta\in\mathbb{R}[x_1,\dots,x_n,w_1,\dots,w_n]$ is real stable.

For example, $1-\partial x_1\partial x_2$ is stability preserving because $1-(-w_1)(-w_2)=1-w_1w_2$ is real stable,

but $1+\partial x_1\partial x_2$ is not stability preserving because $1+(-w_1)(-w_2)=1+w_1w_2$ is not real stable.

Homework: ① Prove that for $1\le k\le n$, the k-th elementary symmetric polynomial $\sum_{S\subseteq\binom{[n]}{k}} x^S$ is real stable.

② Let MAP be the operator that only retains the multiaffine monomials of a given polynomial,

e.g. $\mathrm{MAP}(1+x+3x^3y+2xy)=1+x+2xy$. Prove that MAP is stability preserving.

Interlacing family

We are ready to prove the theorem of the new probabilistic method, which is restated below.

Theorem Let $v_1,v_2,\dots,v_m\in\mathbb{R}^n$ be independent random vectors with finite support.

Then $\|\sum_i v_iv_i^T\|\le$ max-root $\left[\,\mathbb{E}\left(\det\left(\lambda I-\sum_i v_iv_i^T\right)\right)\right]$ with positive probability.

To illustrate the idea, let's consider the Weaver's setting, where each vector v_i has only two possibilities.

We represent the decision on the i-th vector by a binary variable s_i, which is $+1$ if $v_i = \sqrt{2}\begin{pmatrix} u_i \\ 0 \end{pmatrix}$ and

-1 if $v_i = \sqrt{2}\begin{pmatrix} 0 \\ u_i \end{pmatrix}$.

For each fixed choice of all m vectors, we write $A_s = \sum_{i=1}^{m} v_i v_i^T$ as the resulting sum where $s \in \{\pm 1\}^n$.

Using this notation, $\mathbb{E}_{v_1, v_2, \ldots, v_m}\left[\det(\lambda I - \sum_i v_i v_i^T)\right]$ is written as $\mathbb{E}_{s \in \{\pm 1\}^n}\left[\det(\lambda I - A_s)\right]$, which is the average

characteristic polynomial of the 2^m possible sum of the m random outer product of the vectors.

By the result about common interlacing, if we can show that these 2^m characteristic polynomials

have a common interlacing, then the theorem would follow.

But it is easy to see that this is not true in general (exercise: find an example).

This does not mean that theorem is not true, just that we need to work more carefully.

The idea is to consider the "conditional" expectation polynomials and use a tree structure to prove it.

A conditional expectation polynomial is of the form $\mathbb{E}_{\substack{s \in \{\pm 1\}^n \\ \text{given } s_1 = b_1, s_2 = b_2, \ldots, s_k = b_k}} \det(\lambda I - A_s)$, i.e. the conditional expectation

given the first k bits are fixed.

Informally, say m=3, we will prove the following interlacing tree structure:

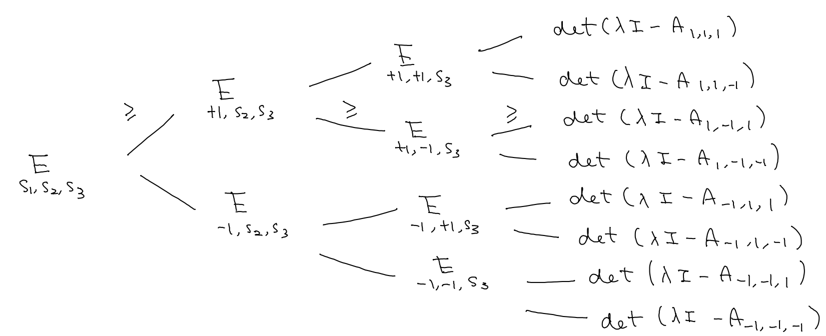

We will show that for each node in the tree, all the polynomials in its children have a common

interlacing (in Weaver's setting, all internal nodes of the tree have two children).

This is enough to establish the theorem, since starting from the root, we know that there is a child with

maxroot at most its parent, and we can repeat this until we reach a leaf, and this is

a polynomial that corresponds to an actual outcome, and we will be done.

Mixed characteristic polynomials

Let A_i be a random symmetric rank-one matrix (e.g. $A_i = \begin{cases} aa^T & \text{with prob } \frac{1}{4} \\ bb^T & \text{with prob } \frac{1}{3} \\ cc^T & \text{with prob } \frac{5}{12} \end{cases}$)

Let $A = \sum_{i=1}^{m} A_i$ be a sum of random rank-one matrices.

We are interested in showing that $\det(\lambda I - \sum_{i=1}^{m} A_i)$ form an interlacing family.

The following identity is at the heart of this approach.

<u>Lemma</u> If $A_1, A_2, \ldots, A_m$ are independent random symmetric rank-one matrices, then

$$\mathop{\mathbb{E}}_{A_1,\ldots,A_m} \det\left(\lambda I - \sum_{i=1}^{m} A_i\right) = \left(\prod_{i=1}^{m}\left(1 - \frac{\partial}{\partial x_i}\right)\right)\det\left(\lambda I + \sum_{i=1}^{m} x_i\, \mathbb{E}[A_i]\right)\Bigg|_{x_1=x_2=\cdots=x_m=0}$$

We will prove this in the next section.

We will first see how this implies that these polynomials form an interlacing family.

<u>Corollary</u> The expected characteristic polynomial $\mathop{\mathbb{E}}_{A_1,\ldots,A_m} \det(\lambda I - \sum_{i=1}^{m} A_i)$ is real-rooted for any independent random symmetric rank-one matrices.

<u>Proof</u> We start from the RHS of the theorem.

Since A_i is a random symmetric rank-one matrices, $\mathbb{E}[A_i] = \sum_j p_j v_j v_j^T \succcurlyeq 0$ is PSD.

By the result in real stable polynomials, we see that $\det(\lambda I + \sum_{i=1}^{m} x_i\, \mathbb{E}[A_i])$ is real stable.

By the result in stability preserving operations, applying the differential operator $1 - \frac{\partial}{\partial x_i}$ and substitution of real numbers preserve stability.

Hence, the LHS of the theorem is a real-stable univariate polynomial, and thus real-rooted. $\square$

<u>Theorem</u> The set of all possible polynomials in $\{\det(\lambda I - \sum_{i=1}^{m} A_i)\}$ form an interlacing family, when $A_1, \ldots, A_m$ are independent random symmetric rank one matrices.

<u>Proof</u> Following the approach in the previous section, we just need to prove that the children of each internal node of the tree have a common interlacing.

More precisely, suppose we fix the first k variables to be $A_1 = v_1 v_1^T, \ldots, A_k = v_k v_k^T$, and let A_{k+1} has l random choices $u_1 u_1^T, \ldots, u_l u_l^T$.

Then, we need to prove that the l conditional polynomials $\mathop{\mathbb{E}}_{A_{k+2},\ldots,A_m} \det(\lambda I - \sum_{i=1}^{k} v_i v_i^T - u_j u_j^T - \sum_{i=k+2}^{m} A_i)$ for $1 \le j \le l$ have a common interlacing.

By the result in common interlacing and real-rootedness, it is equivalent to proving that for any convex combination $\mu \in \mathbb{R}^l$, the polynomial $\sum_{j=1}^{l} \mu_j \mathop{\mathbb{E}}_{A_{k+2},\ldots,A_m} \det(\lambda I - \sum_{i=1}^{k} v_i v_i^T - u_j u_j^T - \sum_{i=k+2}^{m} A_i)$ is real-rooted.

By the result in common interlacing and real-rootedness, it is equivalent to proving that for any convex combination $\mu \in \mathbb{R}^{\ell}$, the polynomial $\sum_{j=1}^{\ell} \mu_j \underset{A_{k+2},\dots,A_m}{\mathbb{E}} \det(\lambda I - \sum_{i=1}^{k} v_i v_i^T - u_j u_j^T - \sum_{i=k+2}^{m} A_i)$ is real-rooted.

Note that this is just the expected characteristic polynomial $\underset{B_1,\dots,B_m}{\mathbb{E}} \det(\lambda I - \sum_{i=1}^{m} B_i)$ for a related set of independent random symmetric rank-one matrices, where B_1 to B_k are just the (deterministic) random variables with $B_i = v_i v_i^T$ with probability one, B_{k+1} is the random variable with $B_{k+1} = u_j u_j^T$ with probability μ_j for $1 \le j \le \ell$, and B_{k+2} and B_m are just the same as the random variables A_{k+2} to A_m.

So, by the previous corollary, this convex combination is real-rooted, and hence the children have a common interlacing, and hence the polynomials form an interlacing family. □

This proves the theorem in the new probabilistic method, assuming the identity, which we will prove in the next section.

Multilinear formula

It remains to prove the identity.

Lemma If $A_1, A_2, \dots, A_m$ are independent random symmetric rank-one matrices, then

$$\underset{A_1,\dots,A_m}{\mathbb{E}} \det(\lambda I - \sum_{i=1}^{m} A_i) = \left(\prod_{i=1}^{m} \left(1 - \frac{\partial}{\partial x_i}\right)\right) \det(\lambda I + \sum_{i=1}^{m} x_i \mathbb{E}[A_i]) \Big|_{x_1 = x_2 = \dots = x_m = 0}$$

We start with the one variable case, using the matrix determinantal formula.

Lemma For a non-singular M, $\det(M + vv^T) = \det(M) \cdot (1 + v^T M^{-1} v) = \det(M) \cdot (1 + tr(M^{-1} vv^T))$.

proof $\det(M + vv^T) = \det(M(I + M^{-1}vv^T)) = \det(M)\det(I + M^{-1}vv^T)$ since $\det(AB) = \det(A)\det(B)$.

Recall that the determinant of a matrix is equal to the product of its eigenvalues.

For the matrix, its eigenvalues are the eigenvalues of $M^{-1}vv^T$ plus one.

Since $M^{-1}vv^T$ is a rank one matrix, its only eigenvalue is $Tr(M^{-1}vv^T) = v^T M^{-1} v$.

Therefore, the spectrum of $M^{-1}vv^T$ is $(v^T M^{-1} v, 1, 1, \dots, 1)$, and so $\det(I + M^{-1}vv^T) = v^T M^{-1} v$. □

Note that the matrix determinantal formula implies that $\det(\lambda I - \sum_{i=1}^{m} A_i)$ is multilinear in terms of A_i, and this uses crucially that A_i is rank one.

Lemma If $A_1, A_2, \dots, A_m$ are independent random symmetric rank-one matrices, then

$$\mathop{\mathbb{E}}_{A_1,\dots,A_m} \det\left(\lambda I - \sum_{i=1}^m A_i\right) = \det\left(\lambda I - \sum_{i=1}^m \mathbb{E}[A_i]\right).$$

Proof. As it is multilinear, $\det(\lambda I - \sum_{i=1}^m A_i)$ can be written as a sum of "monomials" $\sum_{S\subseteq[m]} \prod_{i\in S} f_{i,S}(A_i)$, where each $f_{i,S}$ is a linear function of A_i depending on i and S.

Since $A_1,\dots,A_m$ are independent random variables,

$$\begin{aligned}
\mathop{\mathbb{E}}_{A_1,\dots,A_m} \det\left(\lambda I - \sum_{i=1}^m A_i\right) &= \mathop{\mathbb{E}}_{A_1\dots A_m} \sum_{S\subseteq[m]} \prod_{i\in S} f_{i,S}(A_i) = \sum_{S\subseteq[m]} \prod_{i\in S} \mathop{\mathbb{E}}_{A_i} f_{i,S}(A_i) \quad \text{by independence} \\
&= \sum_{S\subseteq[m]} \prod_{i\in S} f_{i,S}(\mathbb{E}[A_i]) \quad \text{by linearity of } f_{i,S} \\
&= \det\left(\lambda I - \sum_{i=1}^m \mathbb{E}[A_i]\right). \quad \square
\end{aligned}$$

The above lemma proves the first part of the identity, where we moved the expectation inside.

The second part of the identity is also by the multi-linearity of $\det(\lambda I - \sum_{i=1}^m A_i)$.

Lemma. $\det(B + t_1A_1 + \dots + t_mA_m) = \prod_{i=1}^m \left(1 + t_i\frac{\partial}{\partial x_i}\right) \det(B + x_1A_1 + \dots + x_mA_m)\Big|_{x_1=x_2=\dots=x_m=0}$.

Proof. Let $f(t_1,t_2,\dots,t_m)$ be a multilinear polynomial in $t_1,\dots,t_m$.

We write $f(t_1,\dots,t_m) = \sum_{S\subseteq[m]} c_S \prod_{i\in S} t_i$, where c_S is the coefficient of the monomial $\prod_{i\in S} t_i$.

Note that $c_S = \prod_{i\in S} \frac{\partial}{\partial x_i} f(x_1,\dots,x_m)\Big|_{x_1=x_2=\dots=x_m=0}$, as the differentiation and substitution kill all the terms except c_S.

So, $f(t_1,\dots,t_m) = \sum_{S\subseteq[m]} \left(\prod_{i\in S} t_i\right)\left(\prod_{i\in S} \frac{\partial}{\partial x_i} f(x_1,\dots,x_m)\Big|_{x_1=x_2=\dots=x_m=0}\right) = \prod_{i=1}^m \left(1 + t_i\frac{\partial}{\partial x_i}\right) f(x_1,\dots,x_m)\Big|_{x_1=\dots=x_m=0}$.

Now, put it $f(t_1,\dots,t_m) = \det(B + t_1A_1 + \dots + t_mA_m)$ gives the lemma. $\square$

Finally, putting $B = \lambda I$ and $t_i = -1$ for $1 \le i \le m$ proves the identity.

Remark: There are various proofs of this identity, some are shorter and more elegant, but this proof (presented by Tao) is more insightful about why this is true and how it is come up.

References

- Ramanujan graphs and the solution of the Kadison-Singer problem, by Marcus, Spielman, Srivastava.
- Real stable polynomials and the Kadison-Singer problem, blogpost by Terence Tao.
- Real stable and hyperbolic polynomials, course notes by Oveis Gharan.
- Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem, by MSS.