Math 270: Geometry of Polynomials **Fall 2015**

Lecture 1: Introduction, the Matching Polynomial

Lecturer: Nikhil Srivastava *September 1*

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications.*

The Geometry of Polynomials, also known as the analytic theory of polynomials, refers the study of the zero loci of polynomials with complex coefficients (and their dynamics under various transformations of the polynomials) using methods of real and complex analysis. The course will focus on the fragment of this subject which deals with *real-rooted* polynomials and their multivariate generalizations, *real stable* and *hyperbolic* polynomials. We will explore this area via its interactions with questions in combinatorics, probability, and linear algebra, some of which will be algorithmically motivated. Specifically, we will be interested in the following kind of question: how are the properties of a graph/matrix/probability distribution reflected in the zeros of various generating polynomials associated with it? We begin by presenting two of the simplest examples of this interplay.

1.1 Poission Binomial Distributions

The distribution of a sum of independent (not necessarily identically distributed) Bernoulli random variables is called a Poisson Binomial Distribution. A simple question that one might ask about such a distribution is: is it unimodal? That is, letting

$$X = \sum_{i=1}^{n} X_i$$

where X_i are independent Bernoullis with $\mathbb{E}X_i = b_i \in (0,1)$, and taking $p_k = \mathbb{P}[X = k]$, is there some m such that $p_0 \leq p_1 \leq \ldots \leq p_m \geq \ldots \geq p_n$?

This question is quickly answered by studying the generating function

$$q(x) := \sum_{k=0}^{n} p_k x^k = \prod_{i=1}^{n} (b_i x + (1 - b_i)),$$

of the distribution, where the important point is that the independence of the X_i yields a factorization of $q(x)$ into linear terms. In particular, this factorization immediately implies that $q(x)$ is real-rooted with strictly negative roots $\lambda_i := -\frac{1-b_i}{b_i} < 0$. We now appeal to the Newton Inequalities:

Theorem 1.1 (Newton Inequalities). *If $\sum_{k=0}^{n} a_k x^k$ is real-rooted, then*

$$\left(\frac{a_k}{\binom{n}{k}} \right)^2 \geq \frac{a_{k-1}}{\binom{n}{k-1}} \frac{a_{k+1}}{\binom{n}{k+1}}, \tag{1.1}$$

for $k = 1, \ldots, n-1$.

This condition is also known as *ultra log-concavity* (ULC), as after cancelling factorials it reduces to

$$a_k^2 \geq \left(1 + \frac{1}{k}\right)\left(1 + \frac{1}{n-k}\right) a_{k-1}a_{k+1},$$

which is strictly stronger than

$$a_k^2 \geq a_{k-1}a_{k+1} \qquad \text{(log-concavity)}.$$

It is easy to see that log-concavity implies unimodality, whence the probabilities p_k must be unimodal.

The Newton Inequalities are a consequence of two simple *closure properties* of real-rooted polynomials:

- **Differentiation.** If $q(x)$ is real-rooted so is $q'(x)$. The proof is by Rolle's theorem.
- **Inversion.** If $q(x)$ has degree n and is real-rooted, so is $r(x) = x^n q(1/x)$, which has the same coefficients in reverse order. The reason is that the roots of $r(x)$ are the reciprocals of the nonzero roots of $q(x)$.

Proof of Theorem 1.1. Differentiate the polynomial of interest $k-1$ times, reverse the coefficients, and differentiate $n-k-1$ more times to obtain a quadratic polynomial with coefficients equal to a_{k-1}, a_k, and a_{k+1} times some binomial coefficients. This quadratic must be real-rooted by the above closure properties, so its discriminant must be nonnegative, which implies the inequalities. The reader is encouraged to fill in the details as an exercise. □

Thus, we have used facts about polynomials to deduce properties of a probability distribution. However, our proof was entirely reversible, so the implication also goes the other way.

Proposition 1.2. *Suppose* $p(x) = \sum_{k=0}^n a_k x^k$ *is a real-rooted polynomial with nonnegative coefficients and* $a_0 \neq 0$, $p(1) = 1$. *Then there are independent Bernoulli random variables* $X_1, \ldots, X_n$ *such that*

$$a_k = \mathbb{P}\left[\sum_{i=1}^n X_i = k\right].$$

Proof. Factor $p(x)$ as $C\prod_{i=1}^n (x + \lambda_i)$ for some $\lambda_i > 0$. Since $p(1) = 1$ we must have

$$C = \frac{1}{\prod_{i=1}^n (1 + \lambda_i)}.$$

Thus, we have

$$p(x) = \prod_{i=1}^n (b_i x + (1 - b_i))$$

for $b_i = \frac{1}{1+\lambda_i} \in (0,1)$. Taking X_i with $\mathbb{E}X_i = b_i$ proves the claim. □

This allows us to deduce, for instance, that the coefficients of appropriately normalized real-rooted polynomials must must decay exponentially. Letting $S = \sum_{i=1}^n \mathbb{E}X_i$ and $\mu = \mathbb{E}S$, we have by a version of the Chernoff bound:

$$\sum_{k>(1+\epsilon)\mu} a_k = \mathbb{P}\left[S > (1+\epsilon)\mu\right] < \exp(-\epsilon^2\mu/3), \quad 0 < \epsilon < 1,$$

$$\sum_{k<(1-\epsilon)\mu} a_k = \mathbb{P}[[S < (1-\epsilon)\mu] < \exp(-\epsilon^2\mu/2),$$

(as well as the many other more sophisticated variations of such bounds that exist.)

Remark 1.3 (Mixed Volumes)**.** Newton was trying to characterize the set of real-rooted polynomials via inequalities satisfied by their coefficients. It turns out the ultra log-concavity does not characterize real-rootedness (i.e., it is necessary but not sufficient). However, it does characterize another natural class of polynomials, namely: a nonnegative sequence $a_0, \ldots, a_n$ is ULC iff there are convex compact sets A and B in $\mathbb{R}^n$ such that

$$\sum_{k=0}^n a_k t^k = \mathrm{Vol}(tA + B), t \geq 0$$

where Vol is the Lebesgue measure in $\mathbb{R}^n$ and $tA + B$ is the Minkowski sum. This was shown by Shephard in 1960; the interested reader is directed to [She60, Gur09].

Remark 1.4 (Edrei's Equivalence Theorem)**.** A complete set of inequalities characterizing real-rooted polynomials with nonnegative coefficients was discovered by Edrei [Edr53], who showed that $\sum_{k=0}^n a_k x^k$ is real-rooted if and only if the infinite matrix

$$A_{ij} = a_{i-j} \quad i, j \in \mathbb{N}$$

(where we take $a_k = 0$ if $k < 0$) is totally positive, i.e., all of its minors are nonegative. Such sequences are called *Polya Frequency Sequences*; see e.g. [Pit97] for more details.

Note that a real-rooted polynomial has nonnegative coefficients if and only if all its roots are nonpositive.

Remark 1.5 (Hankel Matrices)**.** There is a fully general characterization of real-rootedness in terms of positive semidefiniteness of an associated Hankel matrix[1], due to Hermite and Sylvester (see [BPR11][Theorem 4.37]). In particular, a polynomial

$$p(x) = (x - \lambda_1) \ldots (x - \lambda_n) = \sum_{k=0}^n x^{n-k}(-1)^k e_k$$

is real-rooted if and only if the $n \times n$ matrix defined by

$$H_{ij} = m_{i+j-2},$$

[1]Thanks to Steven Karp for pointing this out.

where $m_k = \sum_{i=1}^n \lambda_i^k$ is the k^th power sum, is positive semidefinite. Since the power sums can be written as polynomials in the coefficients due to Newton's identities:

$$ke_k = \sum_{i=1}^{k}(-1)^{i-1}e_{k-i}m_i, \tag{1.2}$$

this yields a set of inequalities on the coefficients. Positive semidefiniteness can be checked in strongly polynomial time, so this also gives a strongly polynomial time test for real-rootedness.

1.2 Counting Matchings

We now consider a more combinatorially intricate situation: given a graph G on n vertices, let m_k be the number of matchings[2] with k edges in G. For instance, m_1 is just the number of edges in G and $m_{n/2}$ is the number of perfect matchings. We may ask the same question: the sequence m_k unimodal? The difficulty is that the edges in a random matching are not independently distributed, so it we cannot simply appeal to the results of the previous section.

This question is answered (along with many others) in an important statistical mechanics paper from 1972, *Theory of Monomer-Dimer Systems*, by Heilmann and Lieb. The result is most natural in the more general context of weighted graphs, so given a graph with nonnegative edge weights $w_e \geq 0, e \in E$, we define

$$m_k := \sum_{\text{matching } M, |M|=k} \prod_{e \in M} w_e.$$

The relevant generating function is the *matching polynomial*:

$$\mu_G(x) := \sum_{k=0}^{n/2} x^{n-2k}(-1)^k m_k,$$

and the main theorem of that paper is the following:

Theorem 1.6 (Heilmann-Lieb [HL72]). *For every weighted graph G with nonnegative edge weights, $\mu_G(x)$ is real-rooted.*

Observe that we can write $\mu_G(x) = p(-x^2)$ when n is even (since it is a polynomial in x^2 with alternating signs) and $\mu_G(x) = xp(-x^2)$ when n is odd (since in this case there are no perfect matchings), for some polynomial p with nonnegative coefficients equal to the m_k. Theorem 1.6 tells us that p is real-rooted, implying that the m_k must be unimodal (in fact, a Poisson Binomial Distribution by Proposition 1.2).

Before proving the theorem, we cursorily review the statistical physics motivation behind studying such polynomials.

[2] A *matching* is a subset of edges such that no two share a common vertex.

1.2.1 Motivation: Phase Transitions

In Heilmann and Lieb's model, a a "monomer-dimer configuration" consists of a number of non-overlapping "dimers" (a matching) and "monomers" (corresponding to unmatched vertices) in a graph. Such a configuration is entirely determined by the matching, which we will denote by the variable M. There is an energy functional $E(M)$ (the "Hamiltonian") which associates each configuration with a nonnegative real energy, and the state of the system is described by a probability distribution over configurations:

$$p_G(M) = \frac{\exp(-\beta E(M))}{Z},$$

where $\beta > 0$ is a parameter (the "inverse temperature") and Z is a normalization constant (the "partition function").

When E is a sum of local terms involving vertices and pairs of vertices, the above density factors in certain cases of interest into a product over edges and vertices, and may be rewritten as:

$$p_G(M) = \frac{\prod_{uv\in M} w_{uv}\lambda^{n-2|M|}}{Z_G(\lambda)},$$

for weights $w_{uv} > 0$ (depending on β), a parameter $\lambda > 0$, and

$$Z_G(\lambda) = \sum_M \prod_{uv\in M} w_{uv}\lambda^{n-2|M|},$$

where $|M|$ is the number of edges in a matching.

The physicists are interested in whether certain macroscopic properties of this distribution, which correspond to physical observables, vary analytically with the parameter λ. One such observable is the "free entropy":

$$\log Z_G(\lambda).$$

For finite n this quantity is analytic whenever $\lambda > 0$ since Z_G is a polynomial with no positive zeros. However if one takes a "scaling limit"

$$z_G(\lambda) = \lim_{n\to\infty} \frac{1}{n} \log Z_{G_n}(\lambda)$$

for a sequence of graphs G_n converging in a certain appropriate sense, then this is not necessarily the case since the *complex* zeros of the Z_{G_n} may have a limit point on the positive real axis. The main result of [HL72] is that the zeros of Z_G all lie on the imaginary axis, so this does not happen. After performing the change of variable $\lambda = ix$, this is just Theorem 1.6. The physical consequence is that the monomer-dimer model does not exhibit a "phase transition", which corresponds to non-analyticity of z_G. We refer the interested reader to [HL72, Pem12] for more details.

1.2.2 Proof of the Theorem

We now return to the combinatorial setting. The proof of Theorem 1.6 is based on an important recurrence satisfied by $\mu_G(x)$, which may actually be seen as a definition of $\mu_G(x)$. Namely, for any vertex $v \in G$:

$$\mu_G(x) = x\mu_{G\setminus v}(x) - \sum_{u\sim v} w_{uv}\mu_{G\setminus uv}(x), \tag{1.3}$$

where $G \setminus v, G \setminus uv$ refer to vertex-deleted subgraphs and $u \sim v$ denotes vertices adjacent to v in G. The recurrence is easily established by considering matchings which do not contain v and those that contain exactly one edge incident to v. The base case is $\mu_\varnothing(x) = 1$ for the empty graph.

The key structure exploited by the proof is that of *interlacing polynomials*.

Definition 1.7. Let $p(x) = C_1 \prod_{i=1}^n (x - \lambda_i)$ and $q(x) = C_2 \prod_{i=1}^m (x - \nu_i)$ be real-rooted polynomials of degrees differing by at most 1, with $n = \deg(p) \geq \deg(q) = m$. We say that q *interlaces* p if

$$\nu_n \leq \lambda_n \leq \nu_{n-1} \leq \ldots \nu_1 \leq \lambda_1 \quad \text{when } m = n, \text{ or}$$

$$\lambda_n \leq \nu_{n-1} \leq \ldots \nu_1 \leq \lambda_1 \quad \text{when } m = n - 1.$$

If all the inequalities are strict, we say q *strictly interlaces* p.

Proof of Theorem 1.6. Assume that G is a complete graph on n vertices with $w_{uv} > 0$ for all pairs $u, v \in V$ (we will remove this assumption later by a limiting argument). Assume inductively that for every such graph H with at most $n - 1$ vertices:

1. $\mu_H(x)$ is real-rooted with all roots distinct.
2. For every $w \in H$, $\mu_{H\setminus w}(x)$ strictly interlaces $\mu_H(x)$.

We will show that (1) and (2) must be satisfied by G. Fix a vertex $v \in G$ and let $\lambda_{n-1} < \ldots < \lambda_1$ be the roots of $\mu_{G\setminus v}$. We know by induction that each $\mu_{G\setminus uv}$ strictly interlaces $\mu_{G\setminus v}$. Since each of these polynomials is monic, this implies in particular that $\mu_{G\setminus uv}(\lambda_1) > 0$, and since each interval $(\lambda_i, \lambda_{i+1})$ contains exactly one root of each $\mu_{G\setminus uv}$, we deduce that

$$\operatorname{sign}(\mu_{G\setminus uv}(\lambda_i)) = (-1)^{i+1}, \quad i = 1, \ldots, n-1,$$

for all $u \sim v$. Since the weights w_{uv} are positive, the sum

$$r(x) = \sum_{u\sim v} w_{uv}\mu_{G\setminus uv}(x)$$

must also alternate sign at the λ_i. Considering the recurrence (1.3), we now have

$$\mu_G(\lambda_i) = \lambda_i\mu_{G\setminus v}(\lambda_i) - r(\lambda_i) = -r(\lambda_i),$$

so

$$\operatorname{sign}(\mu_G(\lambda_i)) = (-1)^i.$$

By the intermediate value theorem, this means that μ_G has at least one root in each interval $(\lambda_i, \lambda_{i+1})$, yielding $n - 2$ distinct roots. Since $\mu_G(\lambda_1) < 0$ and $\mu_G(x) \to \infty$ as $x \to \infty$, we must also have $\mu_G(\lambda_0) = 0$ for some $\lambda_0 > \lambda_1$. A similar argument yields another root $\lambda_n < \lambda_{n-1}$, for a total of n distinct real roots which are strictly interlaced by the roots of $\mu_{G\setminus v}$.

The base case corresponds to a single edge uv, for which $\mu_G(x) = x^2 - w_{uv}$ and $\mu_{G\setminus v}(x) = x$, for which the claim is true since $w_{uv} > 0$.

To handle the case of general nonnegative weights, consider a sequence of graphs $G^{(k)}$ with weights

$$w_{uv}^{(k)} = \begin{cases} w_{uv} & w_{uv} > 0 \\ 1/k & w_{uv} = 0 \end{cases}$$

converging to the weights in G. Then the polynomials $\mu_{G^{(k)}}(x)$ converge to $\mu_G(x)$ coefficient-wise. Since a limit of real-rooted polynomials is either zero or real-rooted (see the next section), we conclude that $\mu_G(x)$ is real-rooted.

□

1.3 Continuity of Roots

As in the previous section, we will frequently use the fact that a limit of real-rooted polynomials is real-rooted. This is a consequence of the more general fact that the roots of a polynomial are continuous functions of its coefficients (the converse is trivially true). However, some care is required in formalizing what we mean by this statement, since a sequence of polynomials may converge to a polynomial of lower degree, which thereby has strictly fewer roots[3]; for instance, consider the sequence of polynomials

$$f_n(x) := \frac{1}{n}x^2 + x + 1,$$

with roots $\frac{n}{2}(-1 \pm \sqrt{1 - 4/n})$. Perhaps the most elementary formulation is the following.

Theorem 1.8. *Suppose $f_1, f_2, \ldots \in \mathbb{C}[z]$ is a sequence of polynomials of bounded degree with no zeros in an open set $\Omega \subset \mathbb{C}$. If $f_n \to f$ coefficient-wise, then either f is identically 0 or f has no zeros in Ω.*

Proof. Suppose f is not identically zero and $f(w) = 0$ for some $w \in \Omega$. Choose a $\rho > 0$ so that the open disk $D = \{|z - w| < \rho\}$ is contained in Ω, contains no other zeros of f, and $f(z) \neq 0$ on ∂D. Since $f_n \to f$ uniformly on ∂D (because the degrees are bounded and ∂D is compact), we may assume by passing to a subsequence that

$$\min_{z \in \partial D} f_n(z) \geq c > 0 \quad \text{for} \quad c := \frac{1}{2} \min_{z \in \partial D} f(z).$$

Thus, we can conclude that

$$\frac{f_n'(z)}{f_n(z)} \to \frac{f'(z)}{f(z)} \tag{1.4}$$

uniformly on ∂D.

[3]This subtlety goes away if we restrict attention to monic polynomials.

Observe that we can write the number of zeros $m(n)$ of each f_n inside D (counting multiplicity) as an integral of this rational function [4]:

$$\frac{1}{2\pi i}\oint_{\partial D}\frac{f_n'(z)}{f_n(z)}dz = \frac{1}{2\pi i}\oint_{\partial D}\sum_{i=1}^{\deg(f_n)}\frac{1}{z-\lambda_i(f_n)}dz = m(n).$$

By (1.4) we have

$$\oint_{\partial D}\frac{f_n'(z)}{f_n(z)}dz \to \oint_{\partial D}\frac{f'(z)}{f(z)}dz,$$

whence the $m(n)$ must converge to some positive integer, a contradiction. □

Remark 1.9. The above is a special case of a more general theorem about holomorphic functions called Hurwitz's theorem, but we have chosen to present the version above to keep the presentation as self-contained as possible.

Remark 1.10. For polynomials of bounded degree, coefficient-wise convergence is equivalent to uniform convergence on compact subsets. This is not so when the degree is unbounded, and we will use the latter notion of convergence for sequences of unbounded degree.

References

[BPR11] Saugata Basu, Richard Pollack, and Marie-Francoise Roy. Algorithms in real algebraic geometry. *AMC*, 10:12, 2011.

[Edr53] Albert Edrei. Proof of a conjecture of schoenberg on the generating function of a totally positive sequence. *Canad. J. Math*, 5:86–94, 1953.

[Gur09] Leonid Gurvits. A short proof, based on mixed volumes, of liggetts theorem on the convolution of ultra-logconcave sequences. *the electronic journal of combinatorics*, 16(1):N5, 2009.

[HL72] Ole J Heilmann and Elliott H Lieb. Theory of monomer-dimer systems. *Communications in Mathematical Physics*, 25(3):190–232, 1972.

[Pem12] Robin Pemantle. Hyperbolicity and stable polynomials in combinatorics and probability. *arXiv preprint arXiv:1210.3231*, 2012.

[Pit97] Jim Pitman. Probabilistic bounds on the coefficients of polynomials with only real zeros. *Journal of Combinatorial Theory, Series A*, 77(2):279–303, 1997.

[She60] GC Shephard. Inequalities between mixed volumes of convex sets. *Mathematika*, 7(02):125–138, 1960.

[4] The logarithmic derivative f'/f goes by many additional names, including the Cauchy Transform, Stieltjes Transform, and barrier function, and will play a recurring important role in controlling the roots of polynomials in this course.