CS294 Partition Functions **Fall 2020**

Lecture 23: November 17

Instructor: Alistair Sinclair

Disclaimer: *These notes have not been subjected to the usual scrutiny accorded to formal publications. They may be distributed outside this class only with the permission of the Instructor.*

23.1 Interpolation method (continuation)

In this lecture, we continue our discussion of the "polynomial interpolation" approach to computing partition functions, as applied to the classical ferromagnetic Ising model. Recall that the partition function of this model on a graph $G = (V, E)$ is

$$Z_G(\lambda, \mu) := \sum_{S \subseteq V} \lambda^{|E(S,\bar{S})|} \mu^{|S|},$$

where $\lambda \leq 1$ is the neighbor interaction and μ is the external field. We fix λ to an arbitrary value and view Z_G as a univariate polynomial in μ:

$$Z(\mu) = \sum_{k=0}^{n} a_k \mu^k \quad \text{where } a_k := \sum_{|S|=k} \lambda^{|E(S,\bar{S})|}. \tag{23.1}$$

Our approach is to compute the first m terms of the Taylor expansion of $\log Z(\mu)$ around $\mu = 0$, evaluated at any desired μ lying inside a disk $B(0, \ell)$ within which Z has no (complex) zeros. In the last lecture, we showed that taking $m = O(\log(n/\varepsilon)$ is sufficient to obtain a $1 \pm \varepsilon$ multiplicative approximation of $Z(\mu)$.

In this lecture, we complete the analysis of this algorithm, as in [LSS19b], by showing how to compute the Taylor expansion efficiently, and (most importantly) how to identify a zero-free disk $B(0, \ell)$.

23.2 Computing the Taylor expansion

Recall that our goal is to compute the Taylor expansion of $f(\mu) := \log Z(\mu)$, truncated to m terms:

$$f_m(\mu) := \sum_{i=0}^{m} \frac{f^{(i)}(0)\mu^i}{i!}.$$

We show first how to obtain the derivatives $f^{(i)}(0)$ from the coefficients of Z itself. Note that

$$f'(\mu) = \frac{1}{Z(\mu)} \frac{dZ(\mu)}{d\mu} \qquad \Rightarrow \qquad \frac{dZ(\mu)}{d\mu} = Z(\mu) f'(\mu).$$

Taking iterative derivatives of this expression, and plugging in $\mu = 0$, we get

$$Z^{(m)}(0) = \sum_{j=0}^{m-1} \binom{m-1}{j} Z^{(j)}(0) f^{(m-j)}(0).$$

This is a triangular system of linear equations in the variables $\{f^{(j)}(0)\}_{j=1}^{m}$ with coefficients involving the $\{Z^{(i)}(0)\}_{i=0}^{m-1}$, and is non-degenerate since $Z(0) = 1$. Hence we can solve the system in polynomial time, assuming we know $\{Z^{(i)}(0)\}_{i=0}^{m-1}$. But these values are just the first m coefficients of Z!

How do we compute the first $m = O(\log(n/\varepsilon))$ coefficients of Z? Naively, we just enumerate all subsets $S \subseteq V$ with at most m elements, which takes time $n^{O(m)}$, thus leading to a quasi-polynomial time algorithm. If G has bounded degree, however, we can use a combinatorial trick due to Patel and Regts [PR17] to reduce the running time to polynomial. The main idea is that we can restrict the enumeration to sets S that are *connected* in G; and if G has maximum degree Δ then the number of such subsets of size m is only $O(\Delta^m)$, which is $\mathrm{poly}(n, \varepsilon^{-1})$ for $m = O(\log(n/\varepsilon))$.

To sketch how to do this, note from (23.1) that each coefficient a_k of Z is a *weighted subgraph sum*, i.e., a sum of the form $\sum_{S \sqsubseteq G} \alpha_S$, where the sum is over subgraphs induced by subsets of vertices $S \subseteq V(G)$[1] of G and the α_S are real coefficients that depend only on S. (In this case the only non-zero coefficients α_S are those for $|S| = k$, and their values are the weighted cut values as in (23.1).)

It will be useful to rewrite the coefficients in terms of the roots of Z (in similar fashion to the last lecture). Namely, if $r_1, \ldots, r_n$ are the complex roots, then we can write

$$Z(\lambda) = \prod_{i=1}^{n} \left(1 - \frac{\mu}{r_i}\right) = \sum_k (-1)^k e_k \mu^k,$$

where e_k is the kth elementary symmetric polynomial evaluated at the point $(\frac{1}{r_1}, \ldots, \frac{1}{r_n})$. Therefore, since they are (up to signs) equivalent to the a_k, these e_k must also be subgraph sums.

Now we may appeal to *Newton's identities* to relate the symmetric polynomials e_k to the *power sums* $p_k := \sum_{i=1}^{n} (\frac{1}{r_i})^k$. Namely[2]:

$$p_t = \sum_{k=1}^{t-1} (-1)^{k-1} p_{t-k} e_k + (-1)^{t-1} t e_t. \tag{23.2}$$

Using this set of relations, together with the fact that each e_k is a sum over subgraphs of size k, we can express each p_t also as a weighted subgraph sum over subgraphs of size $\leq t$.

The key point now is the following combinatorial fact, due to [CF16]. Call a subgraph sum $\sum_{S \sqsubseteq G} \gamma_S$ *additive* if, for any two vertex disjoint graphs G, H, the subgraph sum for the disjoint union $G \uplus H$ is equal to the sum of the sums for G and H. I.e., $\sum_{S \sqsubseteq G \uplus H} \gamma_S = \sum_{S \sqsubseteq G} \gamma_S + \sum_{S \sqsubseteq H} \gamma_S$.

Fact 23.1. *Any additive sugraph sum $\sum_{S \sqsubseteq G} \gamma_S$ is supported only on* connected *subgraphs, i.e., $\gamma_S = 0$ if S is not connected.*

Now it is easy to see that the subgraph sums corresponding to the power sums p_t are all additive. To see this, note that clearly $Z_{G \uplus H}(\mu) = Z_G(\mu) Z_H(\mu)$, and hence the set of roots of $Z_{G \uplus H}$ is the union of the roots of $Z_G(\mu)$ and $Z_H(\mu)$ (with multiplicities). Thus the power sums satisfy $p_t^{G \uplus H} = p_t^G + p_t^H$, as claimed. By Fact 23.1, we can therefore write each p_t as

$$p_t = \sum_{S \sqsubseteq G} \gamma_S^{(t)},$$

where the sum is supported on *connected* subgraphs S of size $\leq t$. (The superscript (t) indicates the fact that each S will in general have a different coefficient for each p_t in which it appears.) If G has maximum

[1] Strictly speaking, we need to include in S not only the vertex set S and its induced edges, but also the cut edges from S to $\bar{S}$; we won't dwell on this detail here.

[2] This set of identities holds when we view the $\{e_k\}$ and $\{p_k\}$ as polynomials in variables $x_1, \ldots, x_n$, so in particular it holds when both are evaluated at the point $(\frac{1}{r_1}, \ldots, \frac{1}{r_n})$.

degree Δ, the number of such subgraphs is easily seen to be at most $n(e\Delta)^t$, and they can be enumerated in time $O(nt^3(e\Delta)^t)$ **[exercise!]**.

Putting all this together, we can use (23.2) along with the above bound on the number of connected S of each size to compute the p_t by dynamic programming in time $\mathrm{poly}(n,t) \times O(\Delta^t)$, which is $\mathrm{poly}(n,\varepsilon^{-1})$ since $t = O(\log(n/\varepsilon))$. We leave the details, which can be found in [LSS19b], as a (somewhat involved) **exercise**.

23.3 The Lee-Yang Theorem

Our discussion so far has reduced the problem of approximating the partition function to that of finding a suitable zero-free region for Z in the complex plane. In the case of the Ising model, this is supplied by a classical theorem from statistical physics, known as the "Lee-Yang circle theorem" [LY52].

Theorem 23.2. *For any graph G and any fixed $\lambda < 1$, the zeros of the ferromagnetic Ising model partition function $Z(\mu)$ lie on the unit circle in the complex plane.*

Before proving this theorem, let us immediately note the following corollary.

Corollary 23.3. *There is a FPTAS for $Z(\mu)$ for all μ with $|\mu| \neq 1$.*

Proof. We saw in the previous lecture that, if there are no zeros in the open disk $B(0,\ell)$, then truncating the Taylor expansion of $\log Z$ at $m = O(\log(n/\varepsilon))$ terms leads to a $1 \pm \varepsilon$ approximation of $Z(\mu)$ for any μ with $|\mu| < \ell$, and in the previous section of this lecture we saw how to compute this truncation in time $\mathrm{poly}(n,\varepsilon^{-1})$. By Theorem 23.2 we can take $\ell = 1$ to get an FPTAS for all $|\mu| < 1$. This also implies an FPTAS for all $|\mu| > 1$ via the symmetry $Z(\mu) = \mu^n Z(\mu^{-1})$. □

Proof of Theorem 23.2. This theorem has many proofs. We present a delightfully combinatorial one due to Asano [Asa70].

First, we generalize $Z(\mu)$ to a multivariate polynomial, with a parameter μ_i for each vertex i:

$$Z_G(\mu_1,\ldots,\mu_n) := \sum_{S\subseteq V} \lambda^{|E(S,\bar{S})|} \prod_{i\in S} \mu_i.$$

We say that Z_G has the *Lee-Yang property*, abbreviated as "is LY", if

$$(\forall\, i\ |\mu_i| > 1) \Rightarrow Z_G(\mu_1,\ldots,\mu_n) \neq 0.$$

Note that it suffices to prove that Z_G is LY, because setting all $\mu_i = \mu$ then implies that $Z_G(\mu)$ is non-zero for all $|\mu| > 1$, and the same holds also for $|\mu| < 1$ via the relation $Z_G(\mu) = \mu^n Z_G(\mu^{-1})$.

We prove that Z_G is LY by induction on the structure of G.

(i) *Base case: G is a single edge.* In this case the partition function can be written explicitly as

$$Z_G(\mu_1,\mu_2) = \mu_1\mu_2 + \lambda(\mu_1+\mu_2) + 1.$$

[**Exercise:** check this!] Now any zero (μ_1,μ_2) of this polynomial must satisfy

$$|\mu_2| = \left|\frac{1+\lambda\mu_1}{\lambda+\mu_1}\right|.$$

For $\lambda \in (0,1)$ this is a Möbius transform that maps the exterior of the unit disk to its interior. (Note that this is where we crucially use the fact that the model is ferromagnetic. In the antiferromagnetic case, $\lambda > 1$, this property fails.) Hence if $|\mu_1| > 1$ then $|\mu_2| \leq 1$. This implies that Z_G is LY.

(ii) *Disjoint union.* For vertex disjoint graphs G, H, if we assume inductively that Z_G, Z_H are both LY, then $Z_{G \uplus H}$ is also LY. This follows immediately from the fact, observed in the previous section, that the roots of $Z_{G \uplus H}$ are the union of the roots of Z_G, Z_H.

(iii) *Merging vertices.* Suppose inductively that G is LY, and let G' be obtained by merging two vertices of G. We claim that G' is also LY. To see this, let μ_1, μ_2 be the parameters associated with the two vertices of G that are being merged, and let μ denote the parameter of the merged vertex in G'. (See Figure 23.1.) Fix all other parameters $\mu_3, \ldots, \mu_n$ in G so that $|\mu_i| > 1$ for $3 \le i \le n$. The partition

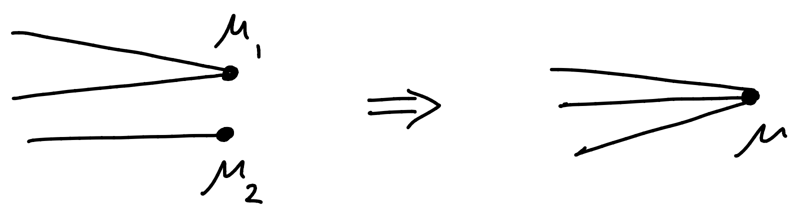

Figure 23.1: Merging two vertices.

function of Z_G can then be written as

$$A\mu_1\mu_2 + B\mu_1 + C\mu_2 + D, \tag{23.3}$$

where A, B, C, D depend on μ_i for $i \ge 3$ (but not on μ_1, μ_2). Since this polynomial is LY, we know that the expression (23.3) cannot be zero if $|\mu_1|, |\mu_2|$ are both > 1. Hence in particular, setting $\mu_1 = \mu_2 = \mu$, the roots of the polynomial

$$A\mu^2 + (B + C)\mu + D$$

have magnitude at most 1, which implies that their product $\left|\frac{D}{A}\right| \le 1$.[3] But note that the partition function of G' (as a function of μ with the other μ_i fixed) is precisely $A\mu + D$ [**exercise:** why?], and its zero is $\mu = \frac{D}{A}$, which we've just shown is ≤ 1 in magnitude. Hence $Z_{G'}$ is indeed LY, as claimed.

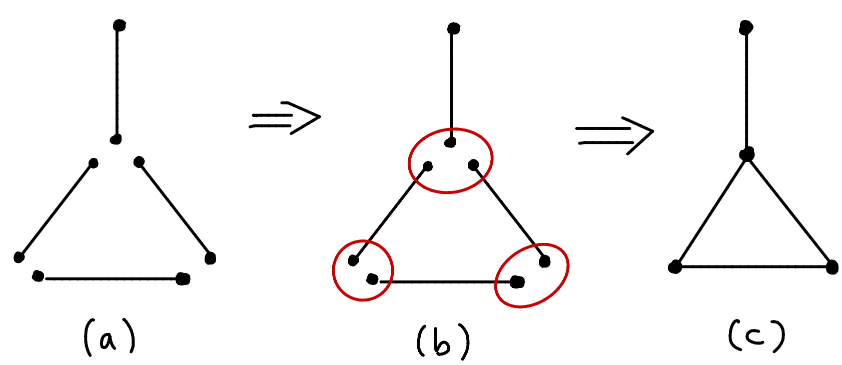

Figure 23.2: Example of proof of Lee-Yang theorem. In (a) we construct a graph consisting of disjoint edges of G. In (b) we merge the endpoints of the edges in pairs. In (c) we arrive at the final graph G.

[3] We are implicitly assuming here that $A \neq 0$. This is true, as can be verified by including this in the induction hypothesis of the present proof; we leave this as an **exercise**.

Finally, we glue the above ingredients together into an inductive proof for any graph G (see Figure 23.2). First, we construct a graph consisting of the disjoint edges of G; by steps (i) and (ii) above, the partition function of this disjoint union is LY. Then we repeatedly merge endpoints of edges to obtain the graph G itself; by step (iii) above the final partition function Z_G will also be LY. This completes the proof. $\square$

23.4 Remarks

1. Recall from the end of the last section that the time complexity of the resulting FPTAS for Z_G includes a factor Δ^t, where $t = O(n/\varepsilon)$, which is $(n/\varepsilon)^{O(\log \Delta)}$. This exponential dependence on $\log \Delta$, which we also saw in Weitz's correlation decay algorithm, means that the running time is no longer polynomially bounded for graphs of unbounded degree.

2. The algorithm presented here *almost* derandomizes the MCMC algorithm for the ferromagnetic Ising model that we saw in Lecture 13. Recall that MCMC provided an FPRAS for all parameter values on any graph. Here, however, we have two restrictions: bounded degree, and non-zero field (i.e., $\mu \neq 1$). This latter restriction is apparently quite significant: recall that the MCMC algorithm also required $\mu \neq 1$ for connectivity of the Markov chain, but was able to handle values of μ very close to 1, specifically $\mu = 1 - \frac{1}{n}$, which is close enough that sampling at this value permits evaluation of the partition function at $\mu = 1$. Here, by contrast, we see from equation (22.7) at the end of the last lecture that the number of terms needed in the Taylor expansion is exponential in $\log(\frac{1}{1-\mu})$, so we need μ to be bounded away from 1. The question of deterministic approximation of $Z_G(\mu)$ at $\mu = 1$ is an important open problem.

3. The approach we have seen above has been used recently to obtain an FPTAS for several other partition functions. Examples include the hard-core model in the uniqueness region $\lambda < \lambda_c(\Delta)$, the antiferromagnetic Ising model in the uniqueness region, and graph colorings for $q \geq 2\Delta$ (and for $q \geq 1.76\Delta$ if the graph is triangle-free) [LSS19a]. (The first two of these applications give alternative deterministic approximation schemes, valid in the same region as the correlation decay-based algorithms discussed previously. The application to colorings is the first known deterministic approximation scheme that works in a range comparable to that of MCMC.) In all cases, the main task is to find a suitable zero-free region in the complex plane. In these other examples, however, we don't have such a powerful and clean location of the zeros as is provided by the Lee-Yang theorem; instead, one can use a tree recurrence of the kind we saw in Weitz's algorithm to inductively establish the existence of a zero-free region, using the fact that the recurrence itself is zero-free on the real line together with a perturbation argument in the complex plane. Numerous other combinatorial examples can be found in [Bar16] and other papers by Barvinok and co-authors, as well as in [PR17].

4. Thanks to the Lee-Yang theorem, we were able to find a zero-free region in the shape of a disk. However, this is not necessary for the method to work. It suffices to find an open, connected region of finite width that includes an interval of the real line containing the easy reference point and the point at which one wants to evaluate the polynomial. Indeed, the other examples mentioned above work this way. One can then use a simple transformation to map such a region to the unit disk and argue as we did above.

23.5 Connection with phase transitions

Loosely speaking, a *phase transition* in a spin system occurs when, for certain parameter values, the local potentials do not fully determine the macroscopic properties of the system as its size tends to infinity (i.e., the system exhibits multiple "phases"). To capture this idea, statistical physicists commonly use two distinct mathematical notions, which we may informally describe as follows. In each case we consider an infinite graph (such as the infinite Δ-regular tree, or the d-dimensional Cartesian lattice $\mathbb{Z}^d$) as the limit of finite subgraphs Λ_n as $n \to \infty$. (E.g., in the case of $\mathbb{Z}^d$, we would take Λ_n to be the cube $[-n, n]^d$ centered at the origin.) We use π_{Λ_n} to denote the Gibbs distribution on the finite graph Λ_n, and Z_{Λ_n} the associated finite partition function.

(i) *Probabilistic.* A phase transition occurs when the infinite-volume Gibbs measure $\pi := \lim_{n\to\infty} \pi_{\Lambda_n}$ (constructed from different boundary conditions) is not unique. (For any given boundary condition, this limit can be shown to exist under mild assumptions.) We have seen this notion in our discussion of the hard-core model and other models on the tree, and we've also mentioned it in the context of the Ising model in $\mathbb{Z}^d$.

(ii) *Analytic.* A phase transition occurs when there is a discontinuity in the *free energy* $\Psi := \lim_{n\to\infty} \frac{1}{|\Lambda_n|} \log Z_{\Lambda_n}$ or one of its derivatives. (Under mild conditions, the free energy can be shown to be well-defined always.)

For most spin systems on amenable graphs, such as $\mathbb{Z}^d$, these two notions can be shown to be equivalent [DS85a, DS85b].

For definiteness, we now specialize our discussion to the classical case of the ferromagnetic Ising model on $\mathbb{Z}^d$. Here, as we have seen earlier, the following picture of the phase diagram is well known:

1. When $\mu = 1$ (i.e., zero field), there is a phase transition at $\lambda = \lambda_c(d)$; i.e., for $\lambda < \lambda_c(d)$ there is a unique Gibbs measure, while for $\lambda > \lambda_c(d)$ there is non-uniqueness (with two extremal Gibbs measures corresponding to the $+$ and $-$ boundary conditions respectively). This transition is also manifested by the fact that the derivative of Ψ (as a function of μ) is *discontinuous* at $\mu = 1$ for $\lambda > \lambda_c(d)$.

2. When $\mu \neq 1$ (non-zero field), there is no phase transition for any value of λ.

How is such behavior established? The quickest route to Property 1 is the probabilistic one. First, using an elementary combinatorial argument, one can show that, at sufficiently low temperatures (large λ) finite correlations persist between the origin (center of Λ_n) and the boundary even as $n \to \infty$ because the probability of a closed "contour" of $+/-$ edges separating the origin from the boundary is bounded away from 1 (note that $+/-$ edges are unlikely at low temperatures); hence the Gibbs measure is not unique. Second, again using an elementary combinatorial argument based this time on the subgraphs-world representation we saw in Lecture 13, one can show that at sufficiently high temperatures such correlations do not persist, by showing that the probability mass on even subgraphs with one odd vertex at the origin and the other on the boundary tends to zero as $n \to \infty$; hence the Gibbs measure is unique. Finally, one can appeal to an easily established monotonicity property to conclude that there must exist a critical value λ_c between these extremes such that the Gibbs measure is unique below λ_c and non-unique above it.

On the other hand, the quickest route to Property 2 is analytic. This is based on the following general result of Yang and Lee [YL52]:

Theorem 23.4. *Let $D \subseteq \mathbb{C}$ be an open, simply connected region of the complex plane such that $D \cap \mathbb{R}$ is an interval of $\mathbb{R}$. If for all Λ_n it is the case that*

$$Z_{\Lambda_n}(\lambda, \mu) \neq 0 \quad \text{for all } \mu \in D, \tag{23.4}$$

then $\Psi(\lambda, \mu) := \lim_{n\to\infty} \frac{1}{|\Lambda_n|} \log Z_{\Lambda_n}(\lambda, \mu)$ admits an analytic *continuation to the whole of D. Thus in particular there is no phase transition for $\mu \in D \cap \mathbb{R}$.*

This theorem is proved via basic techniques in complex analysis; we omit the proof here. Instead, let's see why the theorem is non-trivial and very useful. One might initially assume that there is no need to consider complex arguments, because Z_{Λ_n} is a polynomial with real, non-negative coefficients, so it is never zero and its logarithm is analytic. However, this is no longer necessarily the case when we take the limit as $n \to \infty$ (and indeed, as we mentioned above, the function Ψ is *not analytic* at $\mu = 1$ for $\lambda > \lambda_c$). Theorem 23.4 allows us to take this limit, at the cost of generalizing to the complex plane. The nice thing is that it gives us a *finite* condition (23.4) to check for each n. Note that this condition itself is non-trivial because, once we allow complex arguments, even the finite polynomial Z_{Λ_n} can have zeros.

Now we can see Lee and Yang's motivation for proving the circle theorem, Theorem 23.2. That result tells us that the open unit disk $D = B(0, 1)$ satisfies the hypothesis of Theorem 23.4, and hence there is no phase transition for μ in the *real* interval $(-1, 1)$. In particular, we can conclude Property 2 above, that the Ising model on $\mathbb{Z}^d$ has no phase transition except at the point $\mu = 1$.

Note that the recent flurry of results on polynomial interpolation that we've discussed in the last two lectures can therefore be viewed as an "algorithmicization" of the classical physics theory of phase transitions based on complex analysis.

References

[Asa70] T. Asano. Lee-Yang theorem and the Griffiths inequality for the anisotropic Heisenberg ferromagnet. *Physics Review Letters*, 24:1409–1411, 1970.

[Bar16] A. Barvinok. *Combinatorics and Complexity of Partition Functions*. Springer International, 2016.

[CF16] P. Csikvári and P. Frenkel. Benjamini-Schramm continuity of root moments of graph polynomials. *European Journal of Combinatorics*, 52:302–320, 2016.

[DS85a] R.L. Dobrushin and S.B. Shlosman. Completely analytical Gibbs fields. In J. Fritz, A. Jaffe, and D. Szasz, editors, *Statistical Mechanics and Dynamical Systems: Rigorous Results*, pages 371–403. Birkhäuser, Boston, 1985.

[DS85b] R.L. Dobrushin and S.B. Shlosman. Constructive criterion for the uniqueness of a Gibbs field. In J. Fritz, A. Jaffe, and D. Szasz, editors, *Statistical Mechanics and Dynamical Systems: Rigorous Results*, pages 347–370. Birkhäuser, Boston, 1985.

[LSS19a] J. Liu, A. Sinclair, and P. Srivastava. Correlation decay and partition function zeros: Algorithms and phase transitions. ArXiv preprint 1906.01228, 2020. Preliminary version in *IEEE FOCS*. 2019.

[LSS19b] J. Liu, A. Sinclair, and P. Srivastava. The Ising model partition function: Zeros and deterministic approximation. *Journal of Statistical Physics*, 174:287–315, 2019.

[LY52] T.D. Lee and C.N. Yang. Statistical theory of equations of state and phase transitions. II. Lattice gas and Ising model. *Physics Review*, 87:410–419, 1952.

[PR17] V. Patel and G. Regts. Deterministic polynomial time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46:1893–1919, 2017.

[YL52] C.N. Yang and T.D. Lee. Statistical theory of equations of state and phase transitions. I. Theory of condensation. *Physics Review*, 87:404–409, 1952.