Counting and Sampling	**Fall 2017**

Lecture 13: Correlation Decay: Independence Polynomial

Lecturer: Shayan Oveis Gharan — *Nov 8rd*

In this lecture we study the independence polynomial. Given a graph $G = (V, E)$. The independence polynomial is defined as follows:

$$\text{ind}_G(t) = \sum_I t^{|I|}.$$

For any vertex $v \in V$, the following recurrence is immediate:

$$\text{ind}_G(t) = \text{ind}_{G-v}(t) + t\,\text{ind}_{G-v-N_v}(t).$$

where $N_v = \{u : v \sim u\}$ is the set of neighbors of v in G.

Rewriting the above recurrence we have

$$1 - p_v = \frac{\text{ind}_{G-v}(t)}{\text{ind}_G(t)} = \frac{1}{1 + t\frac{\text{ind}_{G-v-N_v}(t)}{\text{ind}_{G-v}(t)}}. \tag{13.1}$$

The quantity on the LHS is the probability that v is unoccupied, i.e., p_v is the probability that v is occupied. If G is a tree, the quantity $\frac{\text{ind}_{G-v-N_v}(t)}{\text{ind}_{G-v}(t)}$ is naturally interpreted as the probability that none of the neighbors of v is occupied in each of the trees obtained by deleting v.

More concretely, let $\mathbb{T}_n^k$ be the $k-1$-ary tree of depth n which has a single root and every intermediate vertex has $k-1$ children. Let p_n be the probability that the root of $\mathbb{T}_n^k$ is occupied. By (13.1) we can write

$$1 - p_n = \frac{1}{1 + t(1 - p_{n-1})^{k-1}}.$$

Also, $p_0 = \frac{t}{1+t}$. It turns out that the asymptotic behavior of p_n for large and small t are very different. There is a critical t,

$$t_c(k) = \frac{(k-1)^{k-1}}{(k-2)^k}$$

where for every $t < t_c$, there exists $p_\infty = \lim_{n\to\infty} p_n$, and for $t > t_c$ there exists limits

$$p_{even} = \lim_{n\to\infty} p_{2n} \text{ and } p_{odd} = \lim_{n\to\infty} p_{2n+1}.$$

In the terminology of statisticians there is a phase transition at t_c.

Theorem 13.1. *For some $t > 0$ and an integer $k > 2$ consider the transformation*

$$f(x) = \frac{1}{1 + tx^{k-1}}$$

for $0 \le x \le 1$. Let t_c be as defined above. For a positive integer n let f^n denote the n-th iteration of f; e.g., $f^2(x) = f(f(x))$. Then, there exists a unique point x_0 such that $f(x_0) = x_0$ and for all $t < t_c$ and all $0 \le x \le 1$,

$$\lim_{n\to\infty} f^n(x) = x_0.$$

Moreover, $|f^n(x) - x_0| \le \delta^n$ *for some* $\delta > 0$.

If $t > t_c$ *then there exists* x_-, x_+ *such that* $x_- \le x_0 \le x_+$ *and*

$$\lim_{n\to\infty} f^{2n}(x) = \begin{cases} x_- & \text{for all } x < x_0 \\ x_+ & \text{for all } x > x_0. \end{cases}$$

The above theorem directly implies the correlation decay property for $k-1$-ary tree below t_c. The threshold t_c is usually called the uniqueness regime for the infinite tree. Because the occupation probability of the root is independent of weather the leaves are conditioned to be occupied or unoccupied.

On the other hand, observe that for t above the uniqueness regime the marginal probability of the root is no longer unique, and it depends on whether the leaves are conditioned to be occupied.

Weitz [Wei06] in his ground breaking result showed the same decay of correlation property holds at any t below the uniqueness regime. $t_c(\Delta)$, for any graph with maximum degree Δ. He then used this property to design an FPRAS for approximating $\text{ind}_G(t)$ at any such t.

Theorem 13.2 ([Wei06])**.** *There is a deterministic algorithm that for any graph G with maximum degree Δ and for any $t \le \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^{\Delta}}$, approximates $\text{ind}_G(t)$ within $1+\epsilon$ multipicative error in time polynomial in $n, 1/\epsilon$ and exponential in t, Δ.*

After the above theorem it remained an open problem if one can approximate $\text{ind}_G(t)$ for t above $t_c(\Delta)$. More recently, Sly showed that approximating $\text{ind}_G(t)$ for t above $t_c(\Delta)$ NP-hard [Sly10].

In the rest of this lecture we illustrate the main ideas of [Wei06].

Definition 13.3 (Weak Spatial Mixing)**.** *We say that the distribution $I \sim \text{ind}_G(t)$ exhibits weak spatial mixing with rate δ if and only if for any vertex v and any set $S \subseteq V$, and any two configurations σ_S, τ_S,*

$$|p_v^{\sigma_S} - p_v^{\tau_S}| \le \delta(\text{dist}(v, S)).$$

Note that (weak) spatial mixing is useful in algorithm design typically with exponential decay, i.e., when $\delta(\ell) = Ce^{-\ell}$.

It follows from Theorem 13.1 that for large enough n and any k, the tree $\mathbb{T}_n^k$ exhibits weak spatial mixing with exponential decay. To prove Theorem 13.2, Weitz showed that any graph G with maximum degree Δ and any t below $t_c(\Delta)$ exhibits weak spatial mixing with exponential decay. In fact he proves even a stronger property, strong spatial mixing:

Definition 13.4 (Strong Spatial Mixing)**.** *We say that the distribution $I \sim \text{ind}_G(t)$ exhibits strong spatial mixing with rate δ if any only if for any vertex v and any set $S \subseteq V$ and any two configurations σ_S, τ_S*

$$|p_v^{\sigma_S} - p_v^{\tau_S}| \le \delta(\text{dist}(v, U)),$$

where $U \subseteq S$ is the set of vertices of S where σ_S, τ_S differ.

It turns out that one can also prove that the tree $\mathbb{T}_n^k$ exhibits strong spatial mixing with exponential decay below the uniqueness regime. In fact Weitz showed that

Theorem 13.5. *If $\mathbb{T}_n^\Delta$ exhibits weak spatial mixing with rate $\delta(.)$, then it also exhibits strong spatial mixing with rate $\frac{(1+t)(t+(1+t)^\Delta)}{t}\delta(.)$.*

Here we do not prove the above theorem. Instead we focus on the following theorem which is the main contribution of [Wei06]:

Theorem 13.6. *For every positive Δ and t if $\mathbb{T}_n^k$ with activity t exhibits strong spatical mixing with rate δ (as $n \to \infty$), then for the same t every graph with maximum degree Δ also exhibits strong spatial mixing with the same rate $\delta(.)$.*

The proof of this theorem indeed gives a recursive algorithm that can be used to estimate p_v up to a $1+\epsilon/2n$ error. So similar to the algorithm that we discussed last lecture to estimate $\mu_G(t)$ we can recurse and estimate $\text{ind}_G(t)$ within $1+\epsilon$ multiplicative error.

13.1 A (self-avoiding walk) Tree Representation

In this section we prove Theorem 13.6. The proof of the theorem is based on a novel procedure for calculating p_v,

The idea is to construct a tree $\mathbb{T}_{saw}(G,v)$ such that the probability that v is occupied in G is the same as the probability that v is occupied in $\mathbb{T}_{saw}(G,v)$. It follows that we can recursively compute the probability that the root of $\mathbb{T}_{saw}(G,v)$ is occupied using $\log(n/\epsilon)$ depth of recursion similar to the calculations that we did in the last lecture.

Recall that in the tree that we constructed in the last lecture for the matching polynomial we had a vertex for every path starting at the vertex v for which we wanted to estimate the saturation probability. The tree that we will construct will also include vertices closing a cycle; furthermore, these vertices are fixed to be either occupied or unoccupied.

First of all, we need to fix an ordering on the edges adjacent to every vertex of G. So, when we say an edge (u,w) is larger than (u,x) it means that it is larger in the fixed ordering.

Specically, $\mathbb{T}_{saw}(G,v)$ is dened as the tree of all paths originating at v, except that whenever a path closes a cycle the copy (in the tree) of the vertex closing the cycle (in G) is fixed to occupied if the edge closing the cycle is larger than the edge starting the cycle and unoccupied otherwise. See the following figure:

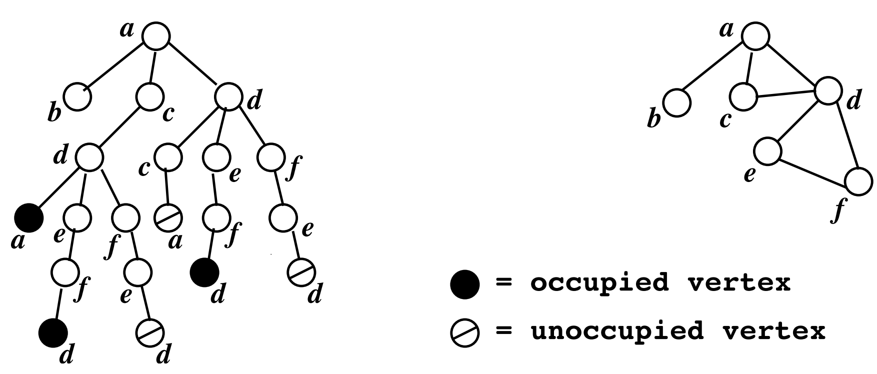

Figure 13.1: The left tree is the tree representation of the right graph. The neighbors of every vertex are ordered alphabetically. So, for example, the occurance of a in the cycle a,c,d is occupied because the edge (a,d) is bigger than the edge (a,c).

The crucial point of the correspondence we establish below between the probability p_v that v is occupied and the probability that the root of $\mathbb{T}_{saw}(G, v)$ is occupied is that it continues to hold when we impose an arbitrary condition on any subset of the vertices of G (and the corresponding condition on the tree). Notice that there is a natural way to correspond a condition on G with a condition on $\mathbb{T}_{saw}(G, v)$. Specifically, since every vertex in the tree $\mathbb{T}_{saw}(G, v)$ corresponds to a vertex in G in a natural way, if a condition on G fixes a vertex u to a certain value, the corresponding condition on $\mathbb{T}_{saw}(G, v)$ fixes all the copies of u to the same value. Note that when we condition a node in the tree to be occupied or unoccupied we can prune the subtree underneath that point because the are independent of the event that the root is occupied.

To make things clear note that we have two types of fixed vertices in $\mathbb{T}_{saw}(G)$. Type I are copies a fixed vertex u in G and type II are structural; these are copies of vertices u that are closing a cycle. Note that if we fix the vertex u to be occupied or unoccupied all structural fixed copies of u will be pruned from the tree automatically because they are grandchildren of another copies of u in the tree.

Theorem 13.7. *For every graph G and every t and every $S \subseteq V$, let $p_v^{\sigma_S}$ be the probability that v is occupied conditioned on σ_S and $\mathbb{P}_v^{\sigma_S}$ be the probability that v is occupied in the tree $\mathbb{T}_{saw}(G)$ when imposing the condition corresponding to σ_S as described above; then,*

$$p_v^{\sigma_S} = \mathbb{P}_v^{\sigma_s}.$$

It is not hard to see that Theorem 13.6 simply follows from the above theorem. First of all, the distance of v to the vertices of S in G is the same as the distance of v to the vertices corresponding to S in $\mathbb{T}_{saw}$ because paths in the tree correspond to paths in G. Secondly, suppose we impose two conditions σ_S, τ_S for a set $S \subseteq V$, and suppose they differ on a set $U \subseteq S$. Then, as we discussed above Type I copies of U in $\mathbb{T}_{saw}$ will be the only locations where the corresponding conditions in $\mathbb{T}_{saw}$ differ. So, if the tree has strong spatial mixing so does G. The only point that we should note is that the tree $\mathbb{T}_s aw$ is not a full $\Delta - 1$-ary tree and it may have significantly less vertices than $\mathbb{T}_n^{\Delta}$ for $n \to \infty$. But this is not of any problem to our calculations because we can assume the vertices of $\mathbb{T}_n^{\Delta}$ which are not present in $\mathbb{T}_s aw$ are conditioned to be unoccupied.

Furthermore, observe that Theorem 13.2 also follows from the above theorem because we can calculate p_v by calculating the probability that the root of $\mathbb{T}_{saw}(G, v)$ is occupied recursively. Here, we do not go to the details as it is essentially similar to the ideas that we discussed last time.

13.2 Proof of Theorem 13.7

In this section we prove Theorem 13.7. This is the most interesting part of the proof. First of all, Weitz works with a change of variables; instead of p_v he works with

$$R_v = \frac{p_v}{1 - p_v}$$

as the ratio of probabilities that v is occupied and occupied. Similarly, we will write $R_v^{\sigma_S}$ to denote the same quantity under conditions on a set S. This is a smart change of variable; it allows him to write a simple recursion for R_{root} for the case of an infinite tree which can seen as an analogous recursion for a general graph G.

We prove the following claim:

Claim 13.8. *Say v is the root of a tree T with degree d. Also, let v_i be the i-th neighbor of v in the tree, and T_i be the subtree rooted at v_i. For a set S of vertices let σ_S be a fixation of vertices in S and let $S_i = T_i \cap S$. Then,*

$$R_v^{\sigma_S} = t \prod_{i=1}^{d} \frac{1}{1 + R_{v_i}^{\sigma_{S_i}}}. \tag{13.2}$$

Note that using the above recursion we can simply estimate R_v for the root by going a few levels down the tree.

Proof. First, by (13.1) we can write

$$1 - p_{T,v} = \frac{1}{1 + t\prod_{i=1}^{d} 1 - p_{T_i,v_i}}.$$

Since $R_v^{\sigma_S} = \frac{p_{T,v}^{\sigma_S}}{1-p_{T,v}^{\sigma_S}}$, we have

$$p_{T,v}^{\sigma_S} = \frac{R_v^{\sigma_S}}{1 + R_v^{\sigma_S}}.$$

So, $1 - p_T^{\sigma_S}(v) = \frac{1}{1+R_v^{\sigma_S}}$. Therefore,

$$\frac{1}{1 - p_{T,v}^{\sigma_S}} = 1 + t\prod_{i=1}^{d} 1 - p_{T_i,v_i}^{\sigma_{S_i}} = 1 + t\prod_{i=1} \frac{1}{1 + R_{v_i}^{\sigma_{S_i}}}.$$

The claim follows from the above equality and the fact that $\frac{1}{1-p_{T,v}^{\sigma_S}} - 1 = \frac{p_{T,v}^{\sigma_S}}{1-p_{T,v}} = R_v^{\sigma_S}$. □

The idea is to write a similar recursion formula for R_v in G, thus relating the marginal probability of v in $\mathbb{T}_{saw}$ with G. Note that if G is a tree then $\mathbb{T}_{saw}$ is the same as G so the proof is obvious. So, the main challenge is to "cancel" the cycle while calculating the marginal probability of the vertices being occupied. To cancel the cycles, we are going to replace v (in G) with d copies $v_1, \ldots, v_d$ where each v_i is incident to a single vertex u_i, the i-th neighbor of v in G. We call this new graph G'.

Roughly speaking, we want to say v is occupied when all v_i's are occupied and it is unoccupied when all of them are unoccupied. But because we have substituted v with d vertices we need to scale down the activity parameter of v_i's. So, we let $t^{1/d}$ be the activity parameter of each v_i. In this way, if all v_i's are occupied it contributes t to the probability of that contribution. Therefore, we can write

$$R_{G,v}^{\sigma_S} = \prod_{i=d}^{1} R_{G',v_i}^{\sigma_S \tau_i}. \tag{13.3}$$

Here τ_i corresponds to fixing all v_j for $j < i$ to be unoccupied and all v_j for $j > i$ to be occupied. Also, by $\sigma_S \tau_i$ we mean concatenating the two fixations. Roughly speaking this is exactly how we cancel the cycles. The above mysterious fixations correspond to the influence of v on itself. The key here is to note that the above fixation is consistent with type II structural fixations that we defined above.

Now, since v_i is the only vertex incident to u_i by (13.1) we have

$$1 - p_{G',v_i}^{\sigma_{S_i}\tau_i} = \frac{1}{1 + t^{1/d}(1 - p_{G'-v,u_i}^{\sigma_{S_i},\tau_i})}.$$

Therefore, similar to the above claim we can write

$$R^{\sigma_{S_i}\tau_i}_{G',v_i} = t^{1/d}\frac{1}{1+R^{\sigma_{S_i},\tau_i}_{G'-v_i,u_i}}.$$

Therefore, by (13.3) we can write

$$R^{\sigma_S}_{G,v} = t\prod_{i=d}^{1}\frac{1}{1+R^{\sigma_{S_i}\tau_i}_{G'-v_i,u_i}}.$$

Note that the above equation also defines a recursion to estimate $R^{\sigma_S}_v$. Furthermore, this recursion will stop after finite number of steps because each time the number unfixed vertices decreases by 1.

To finish the proof of Theorem 13.6 it is enough to show that the above recursion gives the same value for R_v as the (13.2) does when applied to $\mathbb{T}_{saw}(G,v)$. This can be proved by induction. Note that both recursions have the same structure. The only nontrivial is to see that the tree $\mathbb{T}_{saw}(G'-v_i,u_i)$ with the condition corresponding to σ_S,τ_i imposed on it is exactly the same as the the subtree of $\mathbb{T}_{saw}(G,v)$ rooted at u_i with the condition corresponding to σ_S imposed on it. Here is exactly when the structural fixations of type II matter. For every occurrence of v in $\mathbb{T}_{saw}(G,v)$ rooted at u_i, if the edge leading to v is (u_j,v) for $j>i$, then that occurrence of v is occupied and this is consistent with τ_i in the tree $\mathbb{T}_{saw}(G,v)$ rooted at u_i, and similarly if $j<i$, v is unoccupied which is also consistent.

References

[Sly10] Allan Sly. Computational transition at the uniqueness threshold. In *FOCS*, pages 287–296, Washington, DC, USA, 2010. IEEE Computer Society. 13-2

[Wei06] Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149, New York, NY, USA, 2006. ACM. 13-2, 13-3