CS37101-1 Foundations of Markov chain Monte Carlo methods **Spring 2002**

Lecture 1: March 29

Lecturer: Eric Vigoda *Scribe: Varsha Dani & Tom Hayes*

1.1 Introduction

The aim of this course is to address the complexity of counting and sampling problems from an algorithmic perspective. Typically we will be interested in counting the size of a collection of combinatorial structures of a graph, e.g., the number of spanning trees of a graph. We will soon see that this style of counting problem is intimately related to the sampling problem which asks for a random structure from this collection, e.g., generate a random spanning tree. The course will demonstrate that the class of counting and sampling problems can be addressed in a very cohesive and elegant (at least to my tastes) manner.

In this lecture we study some classical algorithms for exact counting. There are few problems which admit efficient algorithms for the exact counting version. In most cases we will have to settle for approximation algorithms. It is interesting to note that both of the algorithms presented in this lecture rely on a reduction to the determinant. This is the case for virtually all (of the very few known) exact counting algorithms.

1.2 Spanning Trees

Our first theorem is known as Kirchoff's Matrix-Tree Theorem [2], and dates back over 150 years.

We are interested in counting the number of spanning trees of an arbitrary undirected graph $G = (V, E)$ with no self-loops. Assume the graph is given by its adjacency matrix $A = (a_{i,j})$ where

$$a_{ij} = \begin{cases} 1 & \text{if } (i,j) \in E \\ 0 & \text{otherwise} \end{cases}$$

Associate the vertex set V with $[n] = \{0, \ldots, n-1\}$, and let D be the diagonal matrix with entry $D(i,i)$ equal to the degree of vertex i for all $i \in [n]$.

It will be useful to consider a corresponding problem on directed graphs.

Definition 1.1 *For a directed graph* $\rightarrow G = (V, \rightarrow E)$*, a set* $\vec{T} \subseteq \vec{E}$ *is an out-tree rooted at* $r \in V$ *if the subgraph* $(V, \vec{T})$ *has the following properties:*

- *it is acyclic (no directed cycles),*
- *the in-degree of* r *is 0, and*
- *the in-degree of every vertex except* r *is exactly 1.*

Taken together these conditions imply that, given an out-tree $\vec{T}$, ignoring orientations the resulting undirected graph is still acyclic. We leave the proof of this observation to the reader. We are interested in the natural directed graph version of our undirected graph.

Let $\overleftrightarrow{G} = (V, \overleftrightarrow{E})$ denote the directed graph formed by replacing each edge $(i,j) \in E$ by a pair of anti-symmetric directed edges $\overrightarrow{ij}$ and $\overrightarrow{ji}$. The following claim illustrates that we can reduce the problem of counting spanning trees to counting out-trees.

Claim 1.2 *For any $r \in V$, the number of spanning trees of G equals the number of out-trees rooted at r in $\overleftrightarrow{G}$.*

Proof: We prove the claim by showing a bijection between the collections. Fix a root $r \in V$. From an out-tree $\overrightarrow{T}$ with root r, form the associated spanning tree by simply removing the orientations on the edges. As we claimed earlier, the resulting undirected graph is acyclic and thus a snapping tree. This mapping can be reversed, hence it is a bijection. Given a spanning tree simply orient each edge away from r (there is a unique path from r to any edge). ■

We need some additional notation before stating the main theorem. Let $K = D - A$, and let $K_{i,i}$ be the minor of K obtained by deleting row i and column i. The matrix K is known as the Kirchoff in-degree matrix.

Theorem 1.3 ([2]) *For any $G = (V, E)$ and $r \in V$, the number of out-trees rooted at r in $\overleftrightarrow{G}$ equals* $\det(K_{r,r})$.

Proof: Without loss of generality, let $r = n - 1$. Recalling the definition of the determinant, we have

$$\det(K_{n-1,n-1}) = \sum_{\sigma \in S_{n-1}} \operatorname{sgn}(\sigma) \prod_i K_{i,\sigma(i)},$$

where

$$\operatorname{sgn}(\sigma) = (-1)^{n-1+\#(\text{cycles})} = (-1)^{\#(\text{even cycles})},$$

and S_{n-1} is the set of permutations of $[n-1] = \{0, 1, \ldots, n-2\}$.

We proceed by decomposing σ based on its fixed points, i.e., self-loops, let

$$L = L(\sigma) = \{i : \sigma(i) = i\},$$

and $\overline{L} = [n-1] \setminus L$.

Using the fact that the number of cycles in σ equals the cardinality of L plus the number of cycles in $\overline{L}$, we have

$$\begin{aligned}
\operatorname{sgn}(\sigma) \prod_i K_{i,\sigma(i)} &= (-1)^{n-1+\#(\text{cycles})} \left(\prod_{i \in L} \sum_{1 \le j \le n} a_{j-i} \right) \left((-1)^{|\overline{L}|} \prod_{i \in \overline{L}} a_{i,\sigma(i)} \right) \\
&= (-1)^{n-1+|L|+|\overline{L}|+\#(\text{cycles in } \overline{L})} \left(\prod_{i \in L} \sum_{1 \le j \le n} a_{j,i} \right) \left(\prod_{i \in \overline{L}} a_{i,\sigma(i)} \right) \\
&= (-1)^{\#(\text{cycles in } \overline{L})} \left(\sum_{f : L \to [n]} \prod_{i \in L} a_{f(i),i} \right) \left(\prod_{i \in \overline{L}} a_{i,\sigma(i)} \right)
\end{aligned}$$

For $i \in \overline{L}$, we can consider $(i, \sigma(i))$ as a possible directed edge. We only need to consider permutations σ where all of these edges exist, otherwise the term is 0. These edges define a directed cycle cover of the vertices $\overline{L}$. Similarly, we can consider the directed graph defined by the edges $\overrightarrow{f(i), i}$ for all $i \in L$. This is a subgraph where each vertex $i \in L$ is assigned one parent.

We can rewrite the above sum over permutations as a sum over the appropriate subgraphs. The corresponding subgraphs are $S = S^f_L \subseteq \overleftrightarrow{E}$ where every vertex $i \in L$ has in-degree 1, and $S' = S^{\sigma}_{\overline{L}} \subseteq \overleftrightarrow{E}$ which is a cycle cover of $\overline{L}$.

For $T \subseteq E$, let $w(T) = \prod_{(i,j)\in T} a_{i,j}$. Then we have

$$\det(K_{n-1,n-1}) = \sum_{(L,S,S')} (-1)^{\#(\text{cycles in } S')} w(S)w(S').$$

When $|L| = n-1$ and S is acyclic, notice that S is an out-tree of $\overleftrightarrow{G}$ rooted at r. We can partition the previous sum with the aim of eliminating any terms in which $|L| \neq n-1$ and/or S contains a cycle. Let c and c' denote the number of cycles in S and S', respectively, and let $C_{\min}$ denote the cycle in $S \cup S'$ which contains the smallest numbered vertex.

$$\begin{aligned}
\det(K_{n-1,n-1}) &= \sum_{\substack{(L,S,S'):\\ |L|=n-1, c=0}} w(S) + \sum_{\substack{(L,S_L,S_{\overline{L}}):\\ c+c'>0}} (-1)^{c'} w(S)w(S') \\
&= \sum_{\substack{(L,S,S'):\\ |L|=n-1,\\ c=0}} w(S) + \sum_{\substack{(L,S,S'):\\ c+c'>0,\\ C_{\min}\in S}} (-1)^{c'} w(S)w(S') + \sum_{\substack{(L,S,S'):\\ c+c'>0,\\ C_{\min}\in S'}} (-1)^{c'} w(S)w(S')
\end{aligned}$$

There is a sign-reversing, magnitude-preserving bijection between the terms of the last two summations. In particular, move the cycle $C_{\min}$ between S and S', simultaneously moving all of the vertices of $C_{\min}$ between L and $\overline{L}$. This bijection implies these two sums cancel each other out, and we are left with the sum over possible out-trees. Since the term $w(S)$ is 1 if all the edges exist and 0 otherwise, we have proven $\det(K_{n-1,n-1})$ equals the number of out-trees rooted at r. ■

1.3 Permanent of planar graphs

The permanent is a natural variant of the determinant without the alternating sign. The formal definition is as follows.

Definition 1.4 *For a $n \times n$ matrix A, let* $\text{per}(A) = \sum_{\sigma\in S_n} \prod_i a_{i,\sigma(i)}$.

If A is a 0/1 matrix, then the permanent equals the number of perfect matchings in the bipartite graph G with incidence matrix A. To be precise, the bipartite graph has $2n$ vertices where each vertex is associated with a specific row or column of A. The vertex for row i and the vertex for column j if $a_{i,j} = 1$. A permutation corresponds to a pairing of row and column vertices. If all of the edges in this pairing exist, and therefore defines a perfect matching, then the permutation contributes 1 to $\text{per}(A)$, and contributes 0 if any edge in the pairing does not exist in G.

Our goal is to compute the permanent for a large class of matrices. For simplicity, we will focus on 0/1 matrices, however everything generalizes to arbitrary non-negative weights on the edges. We shift attention to the graph theoretic problem of perfect matchings and consider the class of planar (not necessarily bipartite) graphs.

We present Kasteleyn's algorithm for computing the number of perfect matchings of a planar graph G. The main idea is to orient the edges of G, resulting in some $\overrightarrow{G}$, so that the number of perfect matchings equals the square root of the determinant of the adjacency matrix of $\overrightarrow{G}$.

Before embarking, we begin with a basic, critical observation about perfect matchings.

Claim 1.5 *For a pair of perfect matchings M, M', the collection $M \cup M'$ consists of disjoint even length cycles and edges.*

An orientation $\vec{G}$ of G is an assignment of a unique direction to each edge. We are aiming for a Pfaffian orientation, which we now formally define.

Definition 1.6 *An even length cycle is oddly oriented in $\vec{G}$, if there are an odd # of clockwise edges.*

Definition 1.7 *An orientation $\vec{G}$ of an undirected graph G is Pfaffian if for all perfect matchings $M, M^{'}$, all cycles of $M \cup M^{'}$ are oddly oriented.*

The previous definition is well-defined since all cycles in $M \cup M'$ are even length. The main result considers the adjacency matrix of an orientation $\vec{G}$.

Definition 1.8 *For an orientation $\vec{G} = (V, \vec{E})$, let $A = (a_{ij})$ where*

$$a_{i,j} = \begin{cases} 1 & \text{if } \vec{ij} \in \vec{E} \\ -1 & \text{if } \vec{ji} \in \vec{E} \\ 0 & \text{otherwise} \end{cases}$$

The following theorem is a major component of Kasteleyn's algorithm.

Theorem 1.9 ([1]) *For any Pfaffian orientation $\overrightarrow{G}$ of G, the square of the number of perfect matchings of G equals the determinant of $A(\overrightarrow{G})$.*

Before proving the theorem, we convert the problem to an associated problem on directed graphs. Recall the definition of $\overleftrightarrow{G} = (V, \overleftrightarrow{E})$ from our earlier work on counting spanning trees. In the following lemma, an even cycle cover is a set of (vertex) disjoint directed even length cycles.

Lemma 1.10 *The square of the number of perfect matchings of G equals the number of even cycle covers of $\overleftrightarrow{G}$.*

Proof: We define a bijection between ordered pairs of perfect matchings and even cycle covers. Consider an ordered pair of perfect matchings M, M' and we will define their associated cycle cover. Each cycle C in $M \cup M^{'}$ is oriented in a canonical way in order to insure the mapping is invertible. Specifically, take the smallest vertex v in C and orient the edge of M incident to v away from v. This forces a unique orientation for the remaining edges of C. Observe that this mapping can easily be reversed, and the lemma is proven. ■

Proof of Theorem: Recall the definition of the determinant,

$$\det(A(\overrightarrow{G})) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_i a_{i,\sigma(i)}$$

Decompose each permutation into its constituent cycles $\sigma = \gamma_1 \ldots \gamma_k$, and let $V_i \subseteq V$ denote the vertices in cycle γ_i.

We claim the set of permutations which contain some odd-length cycle cancel each other out. To prove this claim, consider a permutation σ and its first odd-length cycle γ_i. Let $\sigma' = \gamma_1 \cdots \gamma_{i-1}\gamma_i^{-1}\gamma_{i+1} \cdots \gamma_k$ (we have simply reversed the i-th cycle).

Since σ and σ' contain the same number of cycles, $\text{sgn}(\sigma) = \text{sgn}(\sigma')$. Since γ_i is odd-length, observe that the number of clockwise edges is of opposite parity to the number of counter-clockwise edges, hence

$$\prod_{i \in V_i} a_{i,\gamma(i)} = -\prod_{i \in V_i} a_{i,\gamma_i^{-1}}.$$

Therefore, the contributions from σ and σ' sum to 0.

Consider a permutation σ where every cycle γ_i is even length. We have $\prod_{i \in V_i} a_{i,\sigma(i)} = -1$, for all i, since the orientation is Pfaffian. Thus,

$$\prod_{i \in V} a_{i,\sigma(i)} = (-1)^k = sgn(\sigma),$$

where k is the number of cycles in σ. Therefore, each permutation consisting of only even length cycles contributes 1 to the summation. Since every such permutation corresponds to an even cycle cover of $\overleftrightarrow{G}$, we have proven $\det(A(\overrightarrow{G}))$ equals the number of even cycle covers of $\overleftrightarrow{G}$. This completes the proof of the theorem. ■

We still need to find a Pfaffian orientation for planar graphs. The following structural lemma is the essential ingredient.

Lemma 1.11 *For a planar graph G with orientation $\overrightarrow{G}$, (i) if every face (except possibly the outer face) has an odd number of clockwise edges, then (ii) for all cycles $\overrightarrow{C}$ the number of clockwise edges of $\overrightarrow{C}$ is of opposite parity to the number of vertices in the interior of C, which implies (iii) $\overrightarrow{G}$ is a Pfaffian orientation.*

Proof: We begin by proving (ii) $\implies$ (iii). Consider a cycle C contained in $M \cup M'$ for some pair of perfect matchings M, M'. If any vertex in the interior of C is matched in M with a vertex outside C, then planarity is contradicted. Therefore, there must be an even number of vertices inside C, and (iii) then follows.

We now prove (i) implies (ii). Consider a cycle $\overrightarrow{C}$ of $\overrightarrow{G}$, and the induced subgraph of the vertices on and inside $\overrightarrow{C}$. Let F_{IN} denote the number of non-outer faces in this subgraph, and let $f_1, \ldots, f_{F_{IN}}$ denote the associated faces. Let V_{ON} (respectively, E_{ON}) denote the number of vertices (edges) on $\overrightarrow{C}$. Similarly define V_{IN} and E_{IN} for the vertices and edges strictly inside $\overrightarrow{C}$. Finally, let $E_{ON}^{CLOCK}(f_i)$ (respectively, $E_{ON}^{CLOCK}(C)$) denote the number of clockwise edges on face f_i (cycle $\overrightarrow{C}$).

Plugging the terms into Euler's formula we have

$$(V_{ON} + V_{IN}) - (E_{ON} + E_{IN}) + (F_{IN} + 1) = 2.$$

Since $\overrightarrow{C}$ is a cycle, $V_{ON} = E_{ON}$. This implies,

$$V_{IN} - E_{IN} + F_{IN} = 1. \tag{1.1}$$

By (i), for every face f_i,

$$E_{ON}^{CLOCK}(f_i) \equiv 1 \pmod 2.$$

Summing over all faces inside $\vec{C}$,

$$F_{IN} \equiv \sum_i E_{ON}^{CLOCK}(f_i) \pmod 2. \tag{1.2}$$

Since each edge inside the cycle is clockwise to exactly one of the two faces incident it and each clockwise edge on C is clockwise to the inside face incident it, we have

$$E_{IN} + E_{ON}^{CLOCK}(C) = \sum_i E_{ON}^{CLOCK}(f_i) \tag{1.3}$$

Combining (1.2) with (1.3) yields

$$F_{IN} \equiv E_{ON}^{CLOCK}(C) + E_{IN} \pmod 2.$$

Applying (1.1) we have

$$F_{IN} \equiv E_{ON}^{CLOCK}(C) + F_{IN} + V_{IN} - 1,$$

or equivalently

$$E_{ON}^{CLOCK}(C) + V_{IN} \equiv 1 \pmod 2.$$

This proves the lemma. ■

We are now in position to easily find an Pfaffian orientation of a planar graph. The algorithm constructs the orientation inductively on the number of edges (conditioned on the graph being connected). The base case is a tree. For a tree, the only face is the outer face, hence every edge can be oriented arbitrarily. For a general planar graph, consider an edge e lying on the outer face. Inductively orient $G \setminus e$. Adding e creates one new face f. One of the two orientations of e must guarantee the correct parity for the number of clockwise edges on f. This completes the algorithm for constructing a Pfaffian orientation.

References

[1] P. W. Kasteleyn. Graph theory and crystal physics. In *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, London, 1967.

[2] G. Kirchoff. Über die ausflösung der gleichungen auf welche man bei der untersuchungen der linearen verteilung galvanisher ströme geführt wird. *Poggendorf Ann. Physik*, 72:497–508, 1847.