

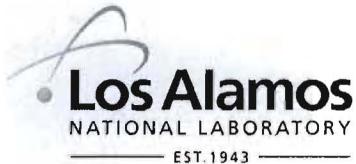
LA-UR- 10-07656

*Approved for public release;
distribution is unlimited.*

Title: Learning Planar Ising Models

Author(s): Jason K. Johnson, 227983, T-4/CNLS LANL
Praneeth netrapalli, Student UT Austin
Michael Chertkov, 149600, T-4 LANL

Intended for: Proceedings: Fourteenth International Conference on Artificial Intelligence and Statistics



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Learning Planar Ising Models

Jason K. Johnson

Center for Nonlinear Studies
& Theoretical Division T-4

Los Alamos National Laboratory, The University of Texas at Austin, Los Alamos National Laboratory,
Los Alamos NM

Praneeth Netrapalli

Dept. of Electrical and
Computer Engineering

Austin TX

Michael Chertkov

Center for Nonlinear Studies
& Theoretical Division T-4

Los Alamos NM

Abstract

Inference and learning of graphical models are both well-studied problems in statistics and machine learning that have found many applications in science and engineering. However, exact inference is intractable in general graphical models, which suggests the problem of seeking the best approximation to a collection of random variables within some tractable family of graphical models. In this paper, we focus our attention on the class of planar Ising models, for which inference is tractable using techniques of statistical physics [Kac and Ward; Kasteleyn]. Based on these techniques and recent methods for planarity testing and planar embedding [Chrobak and Payne], we propose a simple greedy algorithm for learning the best planar Ising model to approximate an arbitrary collection of binary random variables (possibly from sample data). Given the set of all pairwise correlations among variables, we select a planar graph and optimal planar Ising model defined on this graph to best approximate that set of correlations. We present the results of numerical experiments evaluating the performance of our algorithm.

1 Introduction

Graphical models [Lau96, Jor99, Mac03] are widely used to represent the statistical relations among a set of random variables. Nodes of the graph correspond to random variables and edges of the graph represent

statistical interactions among the variables. The problems of inference and learning on graphical models are encountered in many practical applications. The problem of inference is to deduce certain statistical properties (such as marginal probabilities, modes etc.) of a given set of random variables whose graphical model is known. It has wide applications in areas such as error correcting codes, statistical physics and so on. The problem of learning on the other hand is to deduce the graphical model of a set of random variables given statistics (possibly from samples) of the random variables. Learning is also a widely encountered problem in areas such as biology, anthropology and so on.

A certain class of binary-variable graphical models with pairwise interactions known as the *Ising model* has been studied by physicists as a simple model of order-disorder transitions in magnetic materials [Ons44, Bru67]. Remarkably, it was found that in the special case of an Ising model with zero-mean $\{-1, +1\}$ binary random variables and pairwise interactions defined on a planar graph, calculation of the partition function (which is closely tied to inference) is tractable, essentially reducing to calculation of a matrix determinant ([KW52, She60, Kas63, Fis66]). These methods have recently found uses in machine learning [SK08, GJ07].

In this paper, we address the problem of approximating a collection of binary random variables (given their pairwise marginal distributions) by a zero-mean planar Ising model. We also consider the related problem of selecting a non-zero mean Ising model defined on an outer-planar graph (these models are also tractable, being essentially equivalent to a zero-field model on a related planar graph).

There has been a great deal of work on learning graphical models. Much of these has focused on learning over the class of thin graphical models [BJ01, KS01, SCG09] for which inference is tractable by converting the model to tree-structured model. The simplest case of this is learning tree models (treewidth

one graphs) for which it is tractable to find the best tree model by reduction to a max-weight spanning tree problem [CL68]. However, the problem of finding the best bounded-treewidth model is NP-hard for treewidths greater than two [KS01], and so heuristic methods are used to select the graph structure. One popular method is to use convex optimization of the log-likelihood penalized by ℓ_1 norm of parameters of the graphical model so as to promote sparsity [BEd08, LGK06]. To go beyond low treewidth graphs, such methods either focus on Gaussian graphical models or adopt a tractable approximation of the likelihood. Other methods seek only to learn the graph structure itself [RWL10, AKN06] and are often able to demonstrate asymptotic correctness of this estimate under appropriate conditions. One useful application of learning Ising models is for modeling interactions among neurons [CLM07].

The rest of the paper is organized as follows: We present the requisite mathematical preliminaries in Section 2. Section 3 contains our algorithm along with estimates of its computational complexity. We present simulation results in Section 4 and conclude in Section 5 with promising directions for further research and development. All the proofs of propositions are delegated to an appendix to the extended version of the paper (provided as supplementary material to the reviewers; to be made available online upon publication).

2 Preliminaries

In this section, we develop our notation and briefly review the necessary background theory. We will be dealing with binary random variables throughout the paper. We write $P(x)$ to denote the probability distribution of a collection of random variables $x = (x_1, \dots, x_n)$. Unless otherwise stated, we work with undirected graphs $G = (V, E)$ with vertex (or node) set V and edges $\{i, j\} \in E \subset \binom{V}{2}$. For vertices $i, j \in V$ we write $G + ij$ to denote the graph $(V, E \cup \{i, j\})$.

A (pairwise) graphical model is a probability distribution $P(x) = P(x_1, \dots, x_n)$ that is defined on a graph $G = (V, E)$ with vertices $V = \{1, \dots, n\}$ as

$$\begin{aligned} P(x) &\propto \prod_{i \in V} \psi_i(x_i) \prod_{\{i, j\} \in E} \psi_{ij}(x_i, x_j) \\ &\propto \exp\left\{\sum_{i \in V} f_i(x_i) + \sum_{\{i, j\} \in E} f_{ij}(x_i, x_j)\right\} \quad (1) \end{aligned}$$

where $\psi_i, \psi_{ij} \geq 0$ are non-negative node and edge compatibility functions. For positive ψ 's, we may also represent $P(x)$ as a Gibbs distribution with potentials $f_i = \log \psi_i$ and $f_{ij} = \log \psi_{ij}$.

2.1 Entropy, Divergence and Likelihood

For any probability distribution P on some sample space χ , its *entropy* is defined as [CT06]

$$H(P) = - \sum_{x \in \chi} P(x) \log P(x)$$

Suppose we want to calculate how well a probability distribution Q approximates another probability distribution P (on the same sample space χ). For any two probability distributions P and Q on some sample space χ , we denote by $D(P, Q)$ the *Kullback-Leibler divergence* (or *relative entropy*) between P and Q .

$$D(P, Q) = \sum_{x \in \chi} P(x) \log \frac{P(x)}{Q(x)}$$

The *log-likelihood function* is defined as follows:

$$LL(P, Q) = \sum_{x \in \chi} P(x) \log Q(x)$$

The probability distribution in a family \mathcal{F} that maximizes the log-likelihood of a probability distribution P is called the *maximum-likelihood estimate* of P in \mathcal{F} , and this is equivalent to the *minimum-divergence projection* of P to \mathcal{F} :

$$P_{\mathcal{F}} = \arg \max_{Q \in \mathcal{F}} LL(P, Q) = \arg \min_{Q \in \mathcal{F}} D(P, Q)$$

2.2 Exponential Families

A set of random variables $x = (x_1, \dots, x_n)$ are said to belong to an *exponential family* [BN79, WJ08] if there exist functions ϕ_1, \dots, ϕ_m (the *features* of the family) and scalars (*parameters*) $\theta_1, \dots, \theta_m$ such that the joint probability distribution on the variables is given by

$$P(x) = \frac{1}{Z(\theta)} \exp\left(\sum_{\alpha} \theta_{\alpha} \phi_{\alpha}(x)\right)$$

where $Z(\theta)$ is a normalizing constant called the *partition function*. This corresponds to a graphical model if ϕ_{α} happen to be functions on small subsets (e.g., pairs) of all the n variables. The graph corresponding to such a probability distribution is the graph where two nodes have an edge between them if and only if there exists α such that ϕ_{α} depends on both variables. If the functions $\{\phi_{\alpha}\}$ are non-degenerate (please refer to [WJ08] for details), then for any achievable moment parameters $\mu = \mathbb{E}_P[\phi]$ (for an arbitrary distribution P) there exists a unique parameter vector $\theta(\mu)$ that realizes these moments within the exponential family.

Let $\Phi(\theta)$ denote the *log-partition function*

$$\Phi(\theta) \triangleq \log Z(\theta)$$

For an exponential family, we have the following important relation (of *conjugate duality*) between the log-partition function and negative entropy of the corresponding probability distribution $H(\mu) \triangleq H(P_{\theta(\mu)})$ as follows [WJ08]

$$\Phi^*(\mu) \triangleq \max_{\theta \in \mathbb{R}^m} \{ \mu^T \theta - \Phi(\theta) \} = -H(\mu) \quad (2)$$

if the mean parameters μ are achievable under some probability distribution. In fact, this corresponds to the problem of maximizing the log-likelihood relative to an arbitrary distribution P with moments $\mu = \mathbb{E}_P[\phi]$ over the exponential family. The optimal choice of θ realizes the given moments ($\mathbb{E}_\theta[\phi] = \mu$) and this solution is unique for non-degenerate choice of features.

2.3 Ising Graphical Model

The Ising model is a famous model in statistical physics that has been used as simple model of magnetic phenomena and of phase transitions in complex systems.

Definition 1. An Ising model on binary random variables $x = (x_1, \dots, x_n)$ and graph $G = (V, E)$ is the probability distribution defined by

$$P(x) = \frac{1}{Z(\theta)} \exp \left(\sum_{i \in V} \theta_i x_i + \sum_{\{i, j\} \in E} \theta_{ij} x_i x_j \right)$$

where $x_i \in \{-1, 1\}$. Thus, the model is specified by vertex parameters θ_i and edge parameter θ_{ij} .

This defines an exponential family with non-degenerate features ($\phi_i(x) = x_i, i \in V$) and ($\phi_{ij}(x) = x_i x_j, \{i, j\} \in E$) and with corresponding moments ($\mu_i = \mathbb{E}[x_i], i \in V$) and ($\mu_{ij} = \mathbb{E}[x_i x_j], \{i, j\} \in E$).

In fact, any graphical model with binary variables and soft pairwise potentials can be represented as an Ising model with binary variables $x_i = \{-1, +1\}$ and with parameters

$$\theta_i = \frac{1}{2} \sum_{x_i} x_i f_i(x_i) + \frac{1}{4} \sum_{\{i, j\} \in E} \sum_{x_i, x_j} x_i x_j f_{ij}(x_i, x_j)$$

$$\theta_{ij} = \frac{1}{4} \sum_{x_i, x_j} x_i x_j f_{ij}(x_i, x_j).$$

There is also a simple correspondence between the moment parameters of the Ising model and the node and edge-wise marginal distributions. Of course, it is trivial to compute the moments given these marginals: $\mu_i = \sum_{x_i} x_i P(x_i)$ and $\mu_{ij} = \sum_{x_i, x_j} x_i x_j P(x_i, x_j)$. The marginals are recovered from the moments by:

$$P(x_i) = \frac{1}{2}(1 + \mu_i x_i)$$

$$P(x_i, x_j) = \frac{1}{4}(1 + \mu_i x_i + \mu_j x_j + \mu_{ij} x_i x_j)$$

We will be especially concerned with the following subfamily of Ising models:

Definition 2. An Ising model is said to be zero-field if $\theta_i = 0$ for all $i \in V$. It is zero-mean if $\mu_i = 0$ ($P(x_i = \pm 1) = \frac{1}{2}$) for all $i \in V$.

It is simple to verify that the Ising model is zero-field if and only if it is zero-mean. Although the assumption of zero-field appears very restrictive, a general Ising model can be represented as a zero-field model by adding one auxiliary variable node connected to every other node of the graph [GJ07]. The parameters and moments of the two models are then related as follows:

Proposition 1. Consider the Ising model on $G = (V, E)$ with $V = \{1, \dots, n\}$, parameters $\{\theta_i\}$ and $\{\theta_{ij}\}$, moments $\{\mu_i\}$ and $\{\mu_{ij}\}$ and partition function Z . Let $\widehat{G} = (\widehat{V}, \widehat{E})$ denote the extended graph based on nodes $\widehat{V} = V \cup \{n+1\}$ with edges $\widehat{E} = E \cup \{\{i, n+1\}, i \in V\}$. We define a zero-field Ising model on \widehat{G} with parameters $\{\widehat{\theta}_{ij}\}$, moments $\{\widehat{\mu}_{ij}\}$ and partition function \widehat{Z} . If we set the parameters according to

$$\widehat{\theta}_{ij} = \begin{cases} \theta_i & \text{if } j = n+1 \\ \theta_{ij} & \text{otherwise} \end{cases}$$

then $\widehat{Z} = 2Z$ and

$$\widehat{\mu}_{ij} = \begin{cases} \mu_i & \text{if } j = n+1 \\ \mu_{ij} & \text{otherwise} \end{cases}$$

Thus, inference on the corresponding zero-field Ising model on the extended graph \widehat{G} is essentially equivalent to inference on the (non-zero-field) Ising model defined on G .

2.4 Inference for Planar Ising Models

The motivation for our paper is the following result on tractability of inference for the *planar zero-field Ising model*.

Definition 3. A graph is planar if it may be embedded in the plane without any edge crossings.

Moreover, it is known that any planar graph can be embedded such that all edges are drawn as straight lines.

Theorem 1. [KW52]/[She60] Let Z denote the partition function of a zero-field Ising model defined on a planar graph $G = (V, E)$. Let G be embedded in the plane (with edges drawn as straight lines) and let $\phi_{ijk} \in [-\pi, +\pi]$ denote the angular difference (turning angle) between directed edges (i, j) and (j, k) . We define the matrix $W \in \mathbb{C}^{2|E| \times 2|E|}$ indexed by directed

edges of the graph as follows: $W = AD$ where D is the diagonal matrix with $D_{ij,ij} = \tanh \theta_{ij} \triangleq w_{ij}$ and

$$A_{ij,kl} = \begin{cases} \exp(\frac{1}{2}\sqrt{-1}\phi_{ijl}), & j = k \text{ and } i \neq l \\ 0, & \text{otherwise} \end{cases}$$

Then, the partition function of the zero-field planar Ising model is given by:

$$Z = 2^n \left(\prod_{\{i,j\} \in E} \cosh \theta_{ij} \right) \det(I - W)^{\frac{1}{2}}$$

We briefly remark the combinatorial interpretation of this theorem: W is the generating matrix of non-reversing walks of the graph with the weight of a walk γ being

$$w_\gamma = \prod_{(i,j) \in \gamma} w_{ij} \prod_{(i,j,k) \in \gamma} \exp(\frac{1}{2}\sqrt{-1}\phi_{ijk}).$$

The determinant can be interpreted as the (inverse) graph zeta function: $\det(I - W) = \prod_\gamma (1 - w_\gamma)$ where the product is taken over all equivalence classes of aperiodic closed non-reversing walks [She60, Loe10]. A related method for computing the Ising model partition function is based on counting perfect matching of planar graphs [Kas63, Fis66]. We favor the Kac-Ward approach only because it is somewhat more direct.

Since calculation of the partition function reduces to calculating the determinant of a matrix, one may use standard Gaussian elimination methods to evaluate this determinant with complexity $O(n^3)$. In fact, using the generalized nested dissection algorithm to exploit sparsity of the matrix, the complexity of these calculations can be reduced to $O(n^{3/2})$. Thus, inference of the zero-field planar Ising model is tractable and scales well with problem size.

It also turns out that the gradient and Hessian of the log-partition function $\Phi(\theta) = \log Z(\theta)$ can be calculated efficiently from the Kac-Ward determinant formula. We recall that derivatives of $\Phi(\theta)$ recover the moment parameters of the exponential family model [BN79, WJ08]:

$$\nabla \Phi(\theta) = \mathbb{E}_\theta[\phi] = \mu.$$

Thus, inference of moments (and node and edge marginals) are likewise tractable for the zero-field planar Ising model.

Proposition 2. *Let $\mu = \nabla \Phi(\theta)$, $H = \nabla^2 \Phi(\theta)$. Let $S = (I - W)^{-1}A$ and $T = (I + P)(S \circ S^T)(I + P^T)$ where A and W are defined as in Theorem 1, \circ denotes the element-wise product and P is the permutation matrix swapping indices of directed edges (i, j) and (j, i) . Then,*

$$\mu_{ij} = w_{ij} - \frac{1}{2}(1 - w_{ij}^2)(S_{ij,ij} + S_{ji,ji})$$

$$H_{ij,kl} = \begin{cases} 1 - \mu_{ij}^2 & \text{if } ij = kl, \text{ else} \\ -\frac{1}{2}(1 - w_{ij}^2)T_{ij,kl}(1 - w_{kl}^2) & \end{cases}$$

Note, calculating the full matrix S requires $O(n^3)$ calculations. However, to compute just the moments μ only the diagonal elements of S are needed. Then, using the generalized nested dissection method, inference of moments (edge-wise marginals) of the zero-field Ising model can be achieved with complexity $O(n^{3/2})$ [LRT79, LT79, GLV00]. However, computing the full Hessian is more expensive, requiring $O(n^3)$ calculations.

Inference for Outer-Planar Graphical Models

We emphasize that the above calculations require both a planar graph G and a zero-field Ising model. Using the graphical transformation of Proposition 1, the latter zero-field condition may be relaxed but at the expense of adding an auxiliary node connected to all the other nodes. In general planar graphs G , the new graph \widehat{G} may not be planar and hence may not admit tractable inference calculations. However, for the subset of planar graphs where this transformation does preserve planarity inference is still tractable.

Definition 4. *A graph G is said to be outer-planar if there exists an embedding of G in the plane where all the nodes are on the outer face.*

In other words, the graph G is outer-planar if the extended graph \widehat{G} (defined by Proposition 1) is planar. Then, from Proposition 1 and Theorem 1 it follows that:

Proposition 3. *[GJ07] The partition function and moments of any outer-planar Ising graphical model (not necessarily zero-field) can be calculated efficiently. Hence, inference is tractable for any binary-variable graphical model with pairwise interactions defined on an outer-planar graph.*

This motivates the problem of learning outer-planar graphical models for a collection of (possibly non-zero mean) binary random variables.

3 Learning Planar Ising Models

This section addresses the main goals of the paper, which are two-fold:

1. Solving for the maximum-likelihood Ising model on a given planar graph to best approximate a collection of zero-mean random variables.
2. How to select (heuristically) the planar graph to obtain the best approximation.

We address these respective problems in the following two subsections. The solution of the first problem is an integral part of our approach to the second. Both solutions are easily adapted to the context of learning outer-planar graphical models of (possibly non-zero mean) binary random variables.

3.1 ML Parameter Estimation

As discussed in Section 2.2, maximum-likelihood estimation over an exponential family is a convex optimization problem (2) based on the log-partition function $\Phi(\theta)$. In the case of the zero-field Ising model defined on a given planar graph it is tractable to compute $\Phi(\theta)$ via a matrix determinant described in Theorem 1. Thus, we obtain an unconstrained, tractable, convex optimization problem for the maximum-likelihood zero-field Ising model on the planar graph G to best approximate a probability distribution $P(x)$:

$$\max_{\theta \in \mathbb{R}^{|E|}} \left\{ \sum_{ij} (\mu_{ij} \theta_{ij} - \log \cosh \theta_{ij}) - \frac{1}{2} \log \det(I - W(\theta)) \right\}$$

Here, $\mu_{ij} = \mathbb{E}_P[x_i x_j]$ for all edges $\{i, j\} \in G$ and the matrix $W(\theta)$ is as defined in Theorem 1. If P represents the empirical distribution of a set of independent identically-distributed (iid) samples $\{x^{(s)}, s = 1, \dots, S\}$ then $\{\mu_{ij}\}$ are the corresponding empirical moments $\mu_{ij} = \frac{1}{S} \sum_s x_i^{(s)} x_j^{(s)}$.

Newton's Method We solve this unconstrained convex optimization problem using Newton's method with step-size chosen by back-tracking line search [BV04]. This produces a sequence of estimates $\theta^{(s)}$ calculated as follows:

$$\theta^{(s+1)} = \theta^{(s)} + \lambda_s H(\theta^{(s)})^{-1} (\mu(\theta^{(s)}) - \mu)$$

where $\mu(\theta^{(s)})$ and $H(\theta^{(s)})$ are calculated using Proposition 2 and $\lambda_s \in (0, 1]$ is a step-size parameter chosen by backtracking line search (see [BV04] for details). The per iteration complexity of this optimization is $O(n^3)$ using explicit computation of the Hessian at each iteration. This complexity can be offset somewhat by only re-computing the Hessian a few times (reusing the same Hessian for a number of iterations), to take advantage of the fact that the gradient computation only requires $O(n^{\frac{3}{2}})$ calculations. As Newton's method has quadratic convergence, the number of iterations required to achieve a high-accuracy solution is typically 8-16 iterations (essentially independent of problem size). We estimate the computational complexity of solving this convex optimization problem as roughly $O(n^3)$.

3.2 Greedy Planar Graph Selection

We now consider the problem of selection of the planar graph G to best approximate a probability distribution $P(x)$ with pairwise moments $\mu_{ij} = \mathbb{E}_P[x_i x_j]$ given for all $i, j \in V$. Formally, we seek the planar graph that maximizes the log-likelihood (minimizes the divergence) relative to P :

$$\widehat{G} = \arg \max_{G \in \mathcal{P}_V} LL(P, P_G) = \arg \max_{G \in \mathcal{P}_V} \max_{Q \in \mathcal{F}_G} LL(P, Q)$$

where \mathcal{P}_V is the set of planar graphs on the vertex set V , \mathcal{F}_G denotes the family of zero-field Ising models defined on graph G and $P_G = \arg \max_{Q \in \mathcal{F}_G} LL(P, Q)$ is the maximum-likelihood (minimum-divergence) approximation to P over this family.

We obtain a heuristic solution to this graph selection problem using the following greedy edge-selection procedure. The input to the algorithm is a probability distribution P (which could be empirical) on n binary $\{-1, 1\}$ random variables. In fact, it is sufficient to summarize P by its pairwise correlations $\mu_{ij} = \mathbb{E}_P[x_i x_j]$ on all pairs $i, j \in V$. The output is a maximal planar graph G and the maximum-likelihood approximation θ_G to P in the family of zero-field Ising models defined on this graph.

Algorithm 1 GreedyPlanarGraphSelect(P)

```

1:  $G = \emptyset, \theta_G = 0$ 
2: for  $k = 1 : 3n - 6$  do
3:    $\Delta = \{\{i, j\} \subset V | \{i, j\} \notin G, G + ij \in \mathcal{P}_V\}$ 
4:    $\tilde{\mu}_\Delta = \{\tilde{\mu}_{ij} = \mathbb{E}_{\theta_G}[x_i x_j], \{i, j\} \in \Delta\}$ 
5:    $G \leftarrow G \cup \arg \max_{e \in \Delta} D(P_e, \tilde{P}_e)$ 
6:    $\theta_G = \text{PlanarIsing}(G, P)$ 
7: end for

```

The algorithm starts with an empty graph and then sequentially adds edges to the graph one at a time so as to (heuristically) increase the log-likelihood (decrease the divergence) relative to P as much as possible at each step. Here is a more detailed description of the algorithm along with estimates of the computational complexity of each step:

- **Line 3.** First, we enumerate the set Δ of all edges one might add (individually) to the graph while preserving planarity. This is accomplished by an $O(n^3)$ algorithm in which we iterate over all pairs $\{i, j\} \notin G$ and for each such pair we form the graph $G + ij$ and test planarity of this graph using known $O(n)$ algorithms [CP95].
- **Line 4.** Next, we perform tractable inference calculations with respect to the Ising model on G to

calculate the pairwise correlations $\tilde{\mu}_{ij}$ for all pairs $\{i, j\} \in \Delta$. This is accomplished using $O(n^{3/2})$ inference calculations on augmented versions of the graph G . For each inference calculation we add as many edges to G from Δ as possible (setting $\theta = 0$ on these edges) while preserving planarity and then calculate all the edge-wise moments of this graph using Proposition 2 (including the zero-edges). This requires at most $O(n)$ iterations to cover all pairs of Δ , so the worst-case complexity to compute all required pairwise moments is $O(n^{5/2})$.

- *Line 5.* Once we have these moments, which specify the corresponding pairwise marginals of the current Ising model, we compare these moments (pairwise marginals) to those of the input distribution P by evaluating the pairwise KL-divergence between the Ising model and P . As seen by the following proposition, this gives us a lower-bound on the improvement obtained by adding the edge:

Proposition 4. *Let P_G and P_{G+ij} be the projections of P on G and $G + ij$ respectively. Then,*

$$D(P, P_G) - D(P, P_{G+ij}) \geq D(P(x_i, x_j), P_G(x_i, x_j))$$

where $P(x_i, x_j)$ and $P_G(x_i, x_j)$ represent the marginal distributions on x_i, x_j of probabilities P and P_G respectively.

Thus, we greedily select the next edge $\{i, j\}$ to add so as to maximize this lower-bound on the improvement measured by the increase on log-likelihood (this being equal to the decrease in KL-divergence).

- *Line 6.* Finally, we calculate the new maximum-likelihood parameters θ_G on the new graph $G \leftarrow G + ij$. This involves solution of the convex optimization problem discussed in the preceding subsection, which requires $O(n^3)$ complexity. This step is necessary in order to subsequently calculate the pairwise moments $\tilde{\mu}$ which guide further edge-selection steps, and also to provide the final estimate.

We continue adding one edge at a time until a maximal planar graph (with $3n - 6$ edges) is obtained. Thus, the total complexity of our greedy algorithm for planar graph selection is $O(n^4)$.

Non-Maximal Planar Graphs Since adding an edge always gives an improvement in the log-likelihood, the greedy algorithm always outputs a maximal planar graph. However, this might lead to over-

fitting of the data especially when the input probability distribution corresponds to an empirical distribution. In such cases, to avoid over-fitting, we might modify the algorithm so that an edge is added to the graph only if the improvement in log-likelihood is more than some threshold γ . An experimental search can be performed for a suitable value of this threshold (e.g. so as to minimize some estimate of the generalization, such as in cross validation methods [Zha93]).

Outer-Planar Graphs and Non-Zero Means

The greedy algorithm returns a zero-field Ising model (which has zero mean for all the random variables) defined on a planar graph. If the actual random variables are non-zero mean, this may not be desirable. For this case we may prefer to exactly model the means of each random variable but still retain tractability by restricting the greedy learning algorithm to select outer-planar graphs. This model faithfully represents the marginals of each random variable but at the cost of modeling fewer pairwise interactions among the variables.

This is equivalent to the following procedure. First, given the sample moments $\{\mu_i\}$ and $\{\mu_{ij}\}$ we convert these to an equivalent set of zero-mean moments $\hat{\mu}$ on the extended vertex set $\hat{V} = V \cup \{n + 1\}$ according to Proposition 1. Then, we select a zero-mean planar Ising model for these moments using our greedy algorithm. However, to fit the means of each of the original n variables, we initialize this graph to include all the edges $\{i, n + 1\}$ for all $i \in V$. After this initialization step, we use the same greedy edge-selection procedure as before. This yields the graph \hat{G} and parameters $\theta_{\hat{G}}$. Lastly, we convert back to a (non-zero field) Ising model on the subgraph of \hat{G} defined on nodes V , as prescribed by Proposition 1. The resulting graph G and parameters θ_G is our heuristic solution for the maximum-likelihood outer-planar Ising model.

Lastly, we remark that it is not essential that one chooses between the zero-field planar Ising model and the outer-planar Ising model. We may allow the greedy algorithm to select something in between—a partial outer-planar Ising model where only nodes of the outer-face are allowed to have non-zero means. This is accomplished simply by omitting the initialization step of adding edges $\{i, n + 1\}$ for all $i \in V$.

4 Simulations

In this section, we present the results of numerical experiments evaluating our algorithm.

Counter Example The first result, presented in Figure 1 illustrates the fact that our algorithm does

not always recover the exact structure even when the underlying graph is planar and the algorithm is given exact moments as inputs.

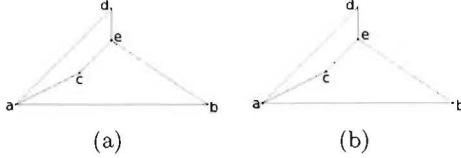


Figure 1: Counter example : (a) Original graphical model (b) Recovered graphical model. The recovered graphical model has one spurious edge $\{a, e\}$ and one missing edge $\{c, d\}$. It is clear from this example that our algorithm is not always optimal.

We define a zero-field Ising model on the graph in Figure 1(a) with the edge parameters as follows: $\theta_{bc} = \theta_{cd} = \theta_{bd} = 0.1$ and $\theta_{ij} = 1$ for all the other edges. Figure 1(a) shows the edge parameters in the graph pictorially using the intensity of the edges - higher the intensity of an edge, higher the corresponding edge parameter. When the edge parameters are as chosen above, the correlation between nodes a and e is greater than the correlation between any other pair of nodes. This leads to the edge between a and e to be the first edge added in the algorithm. However, since $K5$ (the complete graph on 5 nodes) is not planar, one of the actual edges is missed in the output graph. Figure 1(b) shows the edge weighted recovered graph.

Recovery of Zero-Field Planar Ising Model We now present the results of our experiments on a zero field Ising model on a 7×7 grid. The edge parameters are chosen to be uniformly random between -1 and 1 with the condition that the absolute value be greater than a threshold (chosen to be 0.05) so as to avoid edges with negligible interactions. We use Gibbs sampling to obtain samples from this model and calculate empirical moments from these samples which are then passed as input to our algorithm. The results are seen in Figure 2 (see caption for details).

Recovery of Non-Zero-Field Outer Planar Ising Model As explained in Section 3.2, our algorithm can also be used to find the best outer planar graphical model describing a given empirical probability distribution. In this section, we present the results of our numerical experiments on a 12 node outer planar binary pairwise graphical model where the nodes have non-zero mean. Though our algorithm gives perfect reconstruction on graphs with many more nodes, we choose a small example to illustrate the result effectively. We again use Gibbs sampling to obtain samples and calculate empirical moments from those samples. Figure 3(a) presents the original graphical model. Fig-

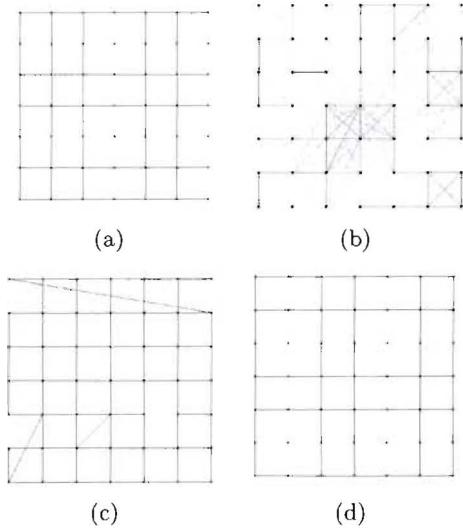


Figure 2: 7×7 grid : (a) Original graphical model (b) Input moments (calculated empirically from 10^5 samples) (c) Recovered graphical model (10^4 samples) (d) Recovered graphical model (10^5 samples). The inputs to the algorithm are the empirical moments obtained from the samples. The algorithm is stopped when the recovered graph has the same number of edges as the original graphical model. With 10^4 samples, there are some errors in the recovered graphical model. When the number of samples is increased to 10^5 , we see perfect recovery.

ures 3(b) and 3(c) present the output graphical models for 10^3 and 10^4 samples respectively. We make sure that the first moments of all the nodes are satisfied by starting with the auxiliary node connected to all other nodes. When the number of samples is 10^3 , the number of erroneous edges in the output as depicted by Figure 3(b) is 0.18. However, as the number of samples increases to 10^4 , the recovered graphical model in Figure 3(c) is exactly the same as the original graphical model.

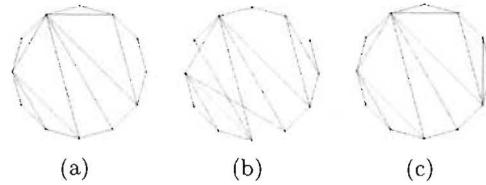


Figure 3: Outer planar graphical model : (a) Original graphical model (b) Recovered graphical model (10^3 samples) (c) Recovered graphical model (10^4 samples). Even in this case, the number of erroneous edges in the output of our algorithm decreases with increasing number of samples. With 10^4 samples, we recover the graphical model exactly.

5 Conclusion

We have proposed a greedy heuristic to obtain the maximum-likelihood planar Ising model approximation to a collection of binary random variables with known pairwise marginals. The algorithm is simple to implement with the help of known methods for tractable inference in planar Ising models, efficient methods for planarity testing and embedding of planar graphs. We have presented simulation results of our algorithm on sample data sets to evaluate its performance.

Future Work Many directions for further work are suggested by the methods and results of this paper. Firstly, we know that the greedy algorithm is not guaranteed to find the best planar graph. Hence, there are several strategies one might consider to further refine the estimate. One is to allow the greedy algorithm to also remove previously-added edges which prove to be less important than some other edge. It may also be possible to use some more generalized notion of local search, such as adding/removing multiple edges at a time such as searching the space of maximal planar graphs by considering “edge flips”, that is, replacing an edge by an orthogonal edge connecting opposite vertices of the two adjacent faces. Also, rather than insisting on the greedy condition of only accepting graphical changes that increase the log-likelihood, one could also consider randomized search strategies such as simulated annealing or genetic programming in the hope of escaping local minima. Another limitation of our current framework is that it only allows learning planar graphical models on the set of observed random variables and, moreover, requires that all variables are observed in each sample. One could imagine extensions of our approach to handle missing samples (using the expectation-maximization approach) or to try to identify hidden variables that were not seen in the data.

References

- [AKN06] P. Abbeel, D. Koller, and A.Y. Ng. Learning factor graphs in polynomial time and sample complexity. *J. Mach. Learn. Res.*, 7, 2006.
- [BEd08] O. Banerjee, L. El Ghaoui, and A. d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *J. Mach. Learn. Res.*, 9, 2008.
- [BJ01] F. Bach and M. Jordan. Thin junction trees. In *NIPS*, 2001.
- [BN79] O. Barndorff-Nielsen. Information and exponential families in statistical theory. *Bull. Amer. Math. Soc.*, 1979.
- [Bru67] S. Brush. History of the Lenz-Ising model. *Rev. Mod. Phys.*, 39(4), 1967.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge U. Press, 2004.
- [CL68] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Trans. on Information Theory*, 14, 1968.
- [CLM07] S. Cocco, S. Leibler, and R. Monasson. On the inverse Ising problem: application to neural networks data. In *XXIII IUPAP Inter. Conf. on Stat. Phys.* Genova, Italy, 2007. URL http://st23.statphys23.org/webservices/abstract/preview_pop.php?ID_PAP%_ER=650.
- [CP95] M. Chrobak and T. Payne. A linear-time algorithm for drawing a planar graph on a grid. *Infor. Processing Letters*, 54(4), 1995.
- [CT06] T. Cover and J. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- [Fis66] M. Fisher. On the dimer solution of planar Ising models. *J. Math. Phys.*, 7(10), 1966.
- [GJ07] A. Globerson and T. Jaakkola. Approximate inference using planar graph decomposition. In *NIPS*, 2007.
- [GLV00] A. Galluccio, M. Loeb, and J. Vondrak. New algorithm for the ising problem: Partition function for finite lattice graphs. *Physical Review Letters*, 84(26), 2000.
- [Jor99] M. Jordan, editor. *Learning in Graphical Models*. MIT Press, 1999.
- [Kas63] P. Kasteleyn. Dimer statistics and phase transitions. *J. Math. Phys.*, 4(2), 1963.
- [KS01] D. Karger and N. Srebro. Learning Markov networks: Maximum bounded tree-width graphs. In *SODA*, 2001.
- [KW52] M. Kac and J. Ward. A combinatorial solution of the two-dimensional Ising model. *Phys. Rev.*, 88(6), 1952.
- [Lau96] S. Lauritzen. *Graphical Models*. Oxford U. Press, 1996.

- [LGK06] S. Lee, V. Ganapathi, and D. Koller. Efficient structure learning of Markov networks using l_1 -regularization. In *NIPS*, 2006.
- [Loe10] M. Loebl. *Discrete Mathematics in Statistical Physics*. Vieweg + Teubner, 2010.
- [LRT79] R. Lipton, D. Rose, and R. Tarjan. Generalized nested dissection. *SIAM J. Numer. Analysis*, 16(2), 1979.
- [LT79] R. Lipton and R. Tarjan. A separator theorem for planar graphs. *SIAM J. Applied Math.*, 36(2), 1979.
- [Mac03] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge U. Press, 2003.
- [Ons44] L. Onsager. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Phys. Rev.*, 65(3-4), 1944.
- [RWL10] P. Ravikumar, M. Wainwright, and J. Lafferty. High-dimensional graphical model selection using ℓ_1 -regularized logistic regression. *Annals of Statistics*, 38(3), 2010.
- [SCG09] D. Shahaf, A. Checketka, and C. Guestrin. Learning thin junction trees via graph cuts. In *AISTATS*, 2009.
- [She60] S. Sherman. Combinatorial aspects of the Ising model for ferromagnetism. i. a conjecture of Feynman on paths and graphs. *J. Math. Phys.*, 1(3), 1960.
- [SK08] N. Schraudolph and D. Kamenetsky. Efficient exact inference in planar Ising models. In *NIPS*, 2008.
- [WJ08] M. Wainwright and M. Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Now Publishers Inc., 2008.
- [Zha93] P. Zhang. Model selection via multifold cross validation. *Annals of Stat.*, 21(1), 1993.