

Sparse Estimation in Ising Model via Penalized Monte Carlo Methods

Błażej Miasojedow

BMA@MIMUW.EDU.PL

Institute of Applied Mathematics and Mechanics

University of Warsaw

ul. Banacha 2, 02-097, Warszawa, Poland

and

Institute of Mathematics

Polish Academy of Sciences

ul. Śniadeckich 8, 00-656 Warszawa

Wojciech Rejchel

WREJCHEL@GMAIL.COM

Faculty of Mathematics and Computer Science

Nicolaus Copernicus University

ul. Chopina 12/18, 87-100 Toruń, Poland

and

Institute of Applied Mathematics and Mechanics

University of Warsaw

ul. Banacha 2, 02-097, Warszawa, Poland

Editor: Sara van de Geer

Abstract

We consider a model selection problem in high-dimensional binary Markov random fields. The usefulness of the Ising model in studying systems of complex interactions has been confirmed in many papers. The main drawback of this model is the intractable norming constant that makes estimation of parameters very challenging. In the paper we propose a Lasso penalized version of the Monte Carlo maximum likelihood method. We prove that our algorithm, under mild regularity conditions, recognizes the true dependence structure of the graph with high probability. The efficiency of the proposed method is also investigated via numerical studies.

Keywords: Ising model, Monte Carlo Markov chain, Markov random field, model selection, Lasso penalty

1. Introduction

A Markov random field is an undirected graph (V, E) , where $V = \{1, \dots, d\}$ is a set of vertices and $E \subset V \times V$ is a set of edges. The structure of this graph describes conditional independence among subsets of a random vector $Y = (Y(1), \dots, Y(d))$, where a random variable $Y(s)$ is associated with a vertex $s \in V$. Finding interactions between random variables is a central element of many branches of science, for example biology, genetics, physics or social network analysis. The goal of the current paper is to recognize the structure of a graph on the basis of a sample consisting of n independent observations. We consider

the high-dimensional setting, i.e. the number of vertices d can be comparable or larger than the sample size n . It is motivated by contemporary applications of Markov random fields in the above-mentioned places, for instance gene microarray data.

The Ising model (Ising, 1925) is an important example of a mathematical model that is often used to explain relations between discrete random variables. In the literature one can find many papers that argue for its effectiveness in recognizing the structure of a graph (Ravikumar et al., 2010; Höfling and Tibshirani, 2009; Guo et al., 2010; Xue et al., 2012; Jalali et al., 2011). This model also plays a key role in our paper. On the other hand, the Ising model is an example of an intractable constant model that is the joint distribution of Y is known only up to a norming constant and this constant cannot be calculated in practice.

Thus, there are two main difficulties in the considered model. The first one is the high-dimensionality of the problem. The second one is the intractable norming constant. To overcome the first obstacle we apply a well-known Lasso method (Tibshirani, 1996). The properties of this method in model selection are deeply studied in many papers that mainly investigate linear models or generalized linear models (Bickel et al., 2009; Bühlmann and van de Geer, 2011; Huang and Zhang, 2012; van de Geer, 2008; Ye and Zhang, 2010; Zhao and Yu, 2006; Zhou, 2009). However, it is not difficult to find papers that describe properties of Lasso estimators in more complex models, for instance Markov random fields (Banerjee et al., 2008; Bühlmann and van de Geer, 2011; Ravikumar et al., 2010; Höfling and Tibshirani, 2009; Guo et al., 2010; Xue et al., 2012) that are considered in this paper.

There are many approaches trying to overcome the second obstacle that is the intractable norming constant. For instance, in Ravikumar et al. (2010) one proposes to perform d regularized logistic regression problems. This idea is based on the fact that the norming constant reduces, if one considers the conditional distribution instead of the joint distribution in the Ising model. This simple fact is at the heart of the pseudolikelihood approach (Besag, 1974) that is replacing the likelihood (that contains the norming constant) by the product of conditionals (that do not contain the norming constant). This idea is widely applied in the literature (Höfling and Tibshirani, 2009; Guo et al., 2010; Xue et al., 2012; Jalali et al., 2011) to study model selection properties of high-dimensional Ising models. However, this approach works well only if the pseudolikelihood is a good approximation of the likelihood. In general, it depends on the true structure of a graph. Namely, if this structure of the graph is sufficiently simple (examples of different structures can be found in section 5.1), then the product of conditionals should be close to the joint distribution. However, in practice this knowledge is unavailable. Another approach is described in Banerjee et al. (2008). It adapts the method that estimates the precision matrix in gaussian graphical models to the binary case. In the current paper we propose the approach to the norming constant problem that relates to Markov chain Monte Carlo (MCMC) methods. Namely, the norming constant is approximated using the importance sampling technique. This method is independent of the unknown complexity of the estimated graph. It is sufficient that the size of a sample used in importance sampling is sufficiently large to have good approximation of the likelihood.

The MCMC method is a well-known approach to overcome the problem with the intractable norming constant in classical (low-dimensional) estimation of graphs. For instance, its properties are investigated in influential papers Geyer and Thompson (1992); Geyer (1994). In the high-dimensional Ising model these algorithms were also studied. For

example, Honorio (2012) and Atchadé et al. (2017) analyzed stochastic versions of proximal gradient algorithms. Both papers derive nonasymptotic bounds between the output of the algorithm and the true minimizer of the cost function. However, in the current paper we focus on model selection properties of MCMC methods. We investigate them in the high-dimensional scenario and compare to the existing methods that are mentioned above. Model selection for undirected graphical models means finding the existing edges in the “sparse” graph that is a graph having relatively few edges (comparing to the total number of possible edges $\frac{d(d-1)}{2}$ and the sample size n).

The paper is organized as follows: in the next section we describe the Ising model and our approach to the problem that relates to minimization of the penalized MCMC approximation of the likelihood. The literature concerning this topic is also discussed. In section 3 we state main theoretical results. Details of efficient implementation are given in section 4, while the results of numerical studies are presented in section 5. The conclusions can be found in section 6. Finally, the proofs are postponed to appendices A and B.

2. Model description and related works

In this section we introduce the Ising model and the proposed method. It also contains a review of the literature relating to this problem.

2.1. Ising Model and undirected graphs

Let (V, E) be an undirected graph that consists of a set of vertices V and a set of edges E . The random vector $Y = (Y(1), Y(2), \dots, Y(d))$, that takes values in \mathcal{Y} , is associated with this graph. In the paper we consider a special case of the Ising model that $Y(s) \in \{-1, 1\}$ and the joint distribution of Y is given by the formula

$$p(y|\theta^*) = \frac{1}{C(\theta^*)} \exp \left(\sum_{r < s} \theta_{rs}^* y(r)y(s) \right), \tag{1}$$

where the sum in (1) is taken over such pairs of indices $(r, s) \in \{1, \dots, d\}^2$ that $r < s$. The vector $\theta^* \in \mathbb{R}^{d(d-1)/2}$ is a true parameter and $C(\theta^*)$ is a norming constant, i.e.

$$C(\theta^*) = \sum_{y \in \mathcal{Y}} \exp \left(\sum_{r < s} \theta_{rs}^* y(r)y(s) \right).$$

The norming constant is a finite sum but it consists of 2^d elements that makes it intractable even for a moderate size of d .

For convenience, we denote $J(y) = (y(r)y(s))_{r < s}$, so

$$p(y|\theta^*) = \frac{1}{C(\theta^*)} \exp [(\theta^*)' J(y)].$$

Remark 1 *The model (1) is a simplified version of the Ising model, for instance we omit an external field in (1). We have decided to restrict to the model containing only parameters θ_{rs}^* , because interactions between random variables is what we focus on in the current paper. However, our results can be relatively easily extended.*

The Ising model has the following property: vertices r and s are not connected by an edge (i.e. $\theta_{rs}^* = 0$) means that variables $Y(r)$ and $Y(s)$ are conditionally independent given the other vertices. Therefore, we recognize the structure of the graph (its edges) by estimating the parameter θ^* . Assume that Y_1, \dots, Y_n are independent random vectors from the model (1). Then the negative log-likelihood is

$$\ell_n(\theta) = -\frac{1}{n} \sum_{i=1}^n \theta' J(Y_i) + \log C(\theta). \quad (2)$$

The second term in (2) contains the norming constant so we cannot use (2) to estimate θ^* . To overcome this problem one usually replaces the negative log-likelihood by its approximation and estimates θ^* using the minimizer of this approximation. In the current paper the approximation of (2) is based on Monte Carlo (MC) methods. Suppose that $h(y)$ is an importance sampling distribution and note that

$$C(\theta) = \sum_{y \in \mathcal{Y}} \exp[\theta' J(y)] = \sum_{y \in \mathcal{Y}} \frac{\exp[\theta' J(y)]}{h(y)} h(y) = \mathbb{E}_{Y \sim h} \frac{\exp[\theta' J(Y)]}{h(Y)} \quad (3)$$

for each θ . An MC approximation of the norming constant is

$$\frac{1}{m} \sum_{k=1}^m \frac{\exp[\theta' J(Y^k)]}{h(Y^k)}, \quad (4)$$

where Y^1, \dots, Y^m is a sample drawn from h or, which is more realistic and is considered in the current paper, Y^1, \dots, Y^m is a Markov chain with h being a density of its stationary distribution. Thus, the MCMC approximation of (2) is

$$\ell_n^m(\theta) = -\frac{1}{n} \sum_{i=1}^n \theta' J(Y_i) + \log \left(\frac{1}{m} \sum_{k=1}^m \frac{\exp[\theta' J(Y^k)]}{h(Y^k)} \right). \quad (5)$$

A natural choice of the importance sampling distribution is $h(y) = p(y|\psi)$ for some parameter ψ . It leads to

$$\ell_n^m(\theta) = -\frac{1}{n} \sum_{i=1}^n \theta' J(Y_i) + \log \left(\frac{1}{m} \sum_{k=1}^m \exp[(\theta - \psi)' J(Y^k)] \right) + \log(C(\psi)). \quad (6)$$

The last term in (6), that contains the unknown constant $C(\psi)$, does not depend on θ , so it can be ignored while minimizing (6).

Our goal is selecting the true model (recognizing edges of a graph) in the high-dimensional setting. It means that the number of vertices d can be large. In fact, it can be greater than the sample size, i.e. $d = d_n \gg n$. To estimate the vector θ^* we use penalized empirical risk minimization. The natural choice of the penalty would be the l_0 -penalty, but it makes the procedure nonconvex and computationally expensive even for moderate values of d . To avoid such problems, we use the Lasso penalty and minimize a function

$$\ell_n^m(\theta) + \lambda_n^m |\theta|_1, \quad (7)$$

where $|\theta|_1 = \sum_{r < s} |\theta_{rs}|$ and $\lambda_n^m > 0$ is a smoothing parameter, that is a balance between minimizing the MCMC approximation and the penalty. We denote the minimizer of (7) by $\hat{\theta}_n^m$. Notice that the function (7), that we minimize, is convex in θ , because the Lasso penalty as well as the MCMC approximation (6) are convex function in θ . The latter follows from the fact that the Hessian of $\ell_n^m(\theta)$, that is given explicitly in (19), is a weighted covariance matrix with positive weights that sum up to one. Convexity of the problem is important from the practical and theoretical point of view. First, every minimum of a convex function is the global minimum, so there are no local minimum problems. Second, convexity is also utilized in the proofs of the results contained in the paper. In further parts of the paper we study properties of $\hat{\theta}_n^m$ in model selection.

2.2. Related works

Model selection in the high-dimensional Ising model is a popular topic and many papers investigating this problem using different methods can be found in the literature (Banerjee et al., 2008; Bresler et al., 2008; Lee et al., 2006; Ravikumar et al., 2010; Höfling and Tibshirani, 2009; Guo et al., 2010; Xue et al., 2012; Jalali et al., 2011). The significant part of them uses the pseudolikelihood approximation with the Lasso penalty. For instance, Ravikumar et al. (2010) applies it by considering d logistic regression problems. They prove that this algorithm is model selection consistent, if some regularity conditions are satisfied. These conditions are similar to the “irrepresentable conditions” (Zhao and Yu, 2006), that are sufficient to prove an analogous property in the linear model. The pseudolikelihood method with the Lasso as a “joint” procedure is proposed in Höfling and Tibshirani (2009). Moreover, in the same paper one also proposes an “exact” algorithm, that minimizes the negative log-likelihood with the Lasso penalty. However, this procedure also bases on the pseudolikelihood approximation. Model selection consistency of the latter algorithm has not been studied yet. The former procedure has this property, that is showed in Guo et al. (2010) provided that conditions similar to Ravikumar et al. (2010) are satisfied. In Xue et al. (2012) the Lasso penalty is replaced by the SCAD penalty (Fan and Li, 2001) and theoretical properties of this algorithm are studied. In Jalali et al. (2011) one replaces restrictive irrepresentable conditions by weaker restricted strong convexity and smoothness conditions (Negahban et al., 2009) and proves model selection consistency of an algorithm that joints ideas from Ravikumar et al. (2010) and Zhang (2009). Namely, it performs d separate logistic regression problems with the forward-backward greedy approach. The algorithm described in Banerjee et al. (2008) is also based on the likelihood approximation with the Lasso penalty. However, it does not apply the pseudolikelihood method. Using the determinant relaxation (Wainwright and Jordan, 2006) it treats the problem of model selection in discrete Markov random fields analogously to the continuous case.

In the current paper we apply the MCMC method to overcome the intractable norming constant problem. Our experimental study (presented in section 5) confirms that estimators based on the MCMC approximation usually perform comparably or better in model selection than their competitors from Banerjee et al. (2008); Höfling and Tibshirani (2009); Ravikumar et al. (2010). Our theoretical results are similar to those described in the previous paragraph, that is we prove model selection consistency. But, in general, our assumptions

are weaker than their analogs from the above-mentioned papers. The detailed comparison is given after Corollary 3 in section 3.

Moreover, the advantage of our algorithm is that the MCMC approach allows us to approximate the norming constant with an arbitrary precision. The approximation error of other methods is given by the problem/data. It depends on the unknown structure of a graph and a user cannot improve it. In our approach a user can improve approximation by increasing the length of simulation, using the MCMC algorithm tailored to the problem. The obvious drawback of our approach is the need of additional simulations to obtain the MCMC sample. It makes our procedure computationally more complex, but at the same time more accurate in selecting the true model.

2.3. Notation

In further parts of the paper we need few notations. Most of them are collected in this subsection.

For simplicity, we write $\hat{\theta}$ and λ instead of $\hat{\theta}_n^m$ and λ_n^m , respectively. Besides, we denote the number of estimated parameters in the model by $\bar{d} = d(d-1)/2$. Nonzero coordinates of θ^* are collected in the set T , and T^c is a completion of T . Besides, $\bar{d}_0 = |T|$ denotes the number of elements in the set T .

For a vector a we denote its l_∞ -norm by $|a|_\infty = \max_k |a_k|$ and $a^{\otimes 2} = aa'$. The vector a_T is the same as the vector a on T and zero otherwise. The l_∞ -norm of a matrix Σ is denoted by $|\Sigma|_\infty = \max_{k,l} |\Sigma_{kl}|$.

Let us consider a Markov chain on a space \mathcal{S} with a transition kernel $P(x, \cdot)$ and a stationary distribution π . We define the Hilbert space $L^2(\pi)$ as a space of functions that $\pi(f^2) < \infty$ and the inner product is given as $\langle f, g \rangle = \int_{\mathcal{S}} f(x)g(x)\pi(dx)$. The linear operator P on $L^2(\pi)$ associated with the transition kernel $P(x, \cdot)$ is defined as follows

$$Pf(x) = \int_{\mathcal{S}} f(y)P(x, dy).$$

We say that the Markov chain has a spectral gap $1 - \kappa$ if and only if

$$\kappa = \sup\{|\rho| : \rho \in \text{Spec}(P) \setminus \{1\}\},$$

where $\text{Spec}(\cdot)$ denotes the spectrum of an operator in $L^2(\pi)$. For reversible chains the spectral gap property is equivalent to geometric ergodicity of the chain (Kontoyiannis and Meyn, 2012; Roberts and Rosenthal, 1997).

In the paper we focus on the Gibbs sampler for the Ising model. However, theoretical results remain true for other MCMC algorithms as long as the spectral gap property is satisfied. The random scan Gibbs sampler for the Ising model with a joint distribution $p(y|\psi)$ is defined as follows: given Y^{k-1} , first we sample uniformly index r and we draw $Y^k(r)$ from the distribution

$$\mathbb{P}(Y^k(r) = 1) = \frac{\exp\{\psi'J(Y^+)\}}{\exp\{\psi'J(Y^+)\} + \exp\{\psi'J(Y^-)\}}, \quad (8)$$

where $Y^+(s) = Y^-(s) = Y^{k-1}(s)$ for $s \neq r$ and $Y^+(r) = 1$, $Y^-(r) = -1$. For $s \neq r$ we set $Y^k(s) = Y^{k-1}(s)$.

Suppose that Y^1, \dots, Y^m is a Markov chain on \mathcal{Y} generated by a random scan Gibbs sampler defined as above. By construction the chain is irreducible, aperiodic and $h(y) = p(y|\psi)$ is the density of its stationary distribution. Therefore, the stationary distribution is defined uniquely and the chain is ergodic for any initial measure ν with the density q . Moreover, there exists a spectral gap $1 - \kappa$, because the state space is finite and the chain is reversible. Actually, κ is the second greatest absolute value of eigenvalues of the transition matrix. We will need three quantities related to this Markov chain :

$$\beta_1 = \sqrt{\sum_{y \in \mathcal{Y}} \frac{q^2(y)}{h(y)}}, \quad \beta_2 = \frac{1 - \kappa}{1 + \kappa}, \quad M = \max_{y \in \mathcal{Y}} \frac{\exp((\theta^*)'J(y))}{h(y)C(\theta^*)}. \quad (9)$$

Roughly speaking, these three values can be viewed as: β_1 – how close the initial density is to the stationary one, β_2 – how fast the chain “mixes”, M – how close the importance sampling density is to the true density (1).

3. Main results

In this section we state key results of the paper. In the first one (Theorem 2) we show that the estimation error of the minimizer of the MCMC approximation with the Lasso penalty can be controlled. In the second result (Corollary 3) we prove model selection consistency for the thresholded Lasso estimator (Zhou, 2009).

First, we introduce the cone invertibility factor that plays an important role in investigating properties of Lasso estimators. It is defined analogously to Ye and Zhang (2010); Huang and Zhang (2012); Huang et al. (2013) that concerns linear regression, generalized linear models and the Cox model, respectively. It is also closely related to the compatibility condition (van de Geer, 2008) or the restricted eigenvalue condition (Bickel et al., 2009). Thus, for $\xi > 1$ and the set T we define a cone as

$$\mathcal{C}(\xi, T) = \{\theta : |\theta_{T^c}|_1 \leq \xi |\theta_T|_1\} .$$

For a nonnegative definite matrix Σ the cone invertibility factor is

$$F(\xi, T, \Sigma) = \inf_{0 \neq \theta \in \mathcal{C}(\xi, T)} \frac{\theta' \Sigma \theta}{|\theta_T|_1 |\theta|_\infty} .$$

Cone invertibility factors of Hessians of two functions are crucial in our argumentation. The first function is the expectation of the negative log-likelihood (2), i.e.

$$\mathbb{E} \ell_n(\theta) = -\theta' \mathbb{E} J(Y) + \log C(\theta) \quad (10)$$

and the second one is the MCMC approximation (5). We denote them as

$$F(\xi, T) = \inf_{0 \neq \theta \in \mathcal{C}(\xi, T)} \frac{\theta' \nabla^2 \log C(\theta^*) \theta}{|\theta_T|_1 |\theta|_\infty} \quad (11)$$

and

$$\bar{F}(\xi, T) = \inf_{0 \neq \theta \in \mathcal{C}(\xi, T)} \frac{\theta' \nabla^2 \ell_n^m(\theta^*) \theta}{|\theta_T|_1 |\theta|_\infty}, \quad (12)$$

respectively. Notice that only the values of $\nabla^2 \log C(\theta)$ and $\nabla^2 \ell_n^m(\theta)$ at the true parameter θ^* are taken into consideration in (11) and (12).

Now we can state main results of the paper.

Theorem 2 *Let $\varepsilon > 0, \xi > 1$ and $\alpha(\xi) = 2 + \frac{\varepsilon}{\xi-1}$. If*

$$n \geq \frac{8(1+\xi)^4 \alpha^2(\xi) \bar{d}_0^2 \log(2\bar{d}/\varepsilon)}{F^2(\xi, T)}, \quad (13)$$

$$m \geq \frac{64(1+\xi)^4 \alpha^2(\xi) \bar{d}_0^2 M^2 \log[2\bar{d}(\bar{d}+1)\beta_1/\varepsilon]}{F^2(\xi, T)\beta_2}, \quad (14)$$

then with probability at least $1 - 4\varepsilon$ we have the inequality

$$\left| \hat{\theta} - \theta^* \right|_\infty \leq \frac{2e \xi \alpha(\xi) \lambda}{(\xi+1)[\alpha(\xi)-2]F(\xi, T)}, \quad (15)$$

where

$$\lambda = \frac{\xi+1}{\xi-1} \max \left(2\sqrt{\frac{2\log(2\bar{d}/\varepsilon)}{n}}, 8M\sqrt{\frac{\log[(2\bar{d}+1)\beta_1/\varepsilon]}{m\beta_2}} \right). \quad (16)$$

Corollary 3 *Suppose that conditions (13) and (14) are satisfied. Let $\theta_{min}^* = \min_{(r,s) \in T} |\theta_{rs}^*|$ and R_n^m denote the right-hand side of the inequality (15). Consider the Lasso estimator with a threshold $\delta > 0$ that is the set of nonzero coordinates of the final estimator is defined as $\hat{T} = \{(r, s) : |\hat{\theta}_{rs}| > \delta\}$. If $\theta_{min}^*/2 > \delta \geq R_n^m$, then*

$$P(\hat{T} = T) \geq 1 - 4\varepsilon.$$

The main results of the paper describe properties of estimators, that are obtained by minimization of the MCMC approximation (5) with the Lasso penalty. Theorem 2 states that the estimation error of the Lasso estimator can be controlled. Roughly speaking, the estimation error is small, if the initial sample size and the MCMC sample size are large enough, the model is sparse and the cone invertibility factor $F(\xi, T)$ is not too close to zero. The influence of the model parameters (n, d, \bar{d}_0) as well as Monte Carlo parameters (m, β_1, β_2, M) on the results are explicitly stated. It is worth to emphasize that our results work in the high-dimensional scenario, i.e. the number of vertices d can be greater than the sample size n provided that the model is sparse. Indeed, the condition (13) is satisfied even if $\bar{d} \sim O(e^{n^{c_1}})$, $\bar{d}_0 \sim O(n^{c_2})$ and $c_1 + 2c_2 < 1$. The condition (14), that relates to the MCMC sample size, is also reasonable. The number β_1 depends on the initial and stationary distributions. In general, its relation to the number of vertices is exponential. However, in (14) it appears with the logarithm. Moreover, β_1 is also reduced using so called burn-in time, i.e. the beginning of the Markov chain trajectory is discarded. Next, the number β_2 is related to the spectral gap of a Markov chain. Under mild conditions the inverse of β_2 depends polynomially on d , and under strong regularity conditions it can be reduced to $O(d \log d)$ as in Mossel and Sly (2013). Finally, there is also the number M in the condition (14) that relates to the distance between the stationary distribution $h(\cdot)$ and $p(\cdot|\theta^*)$. Stating

the explicit relation between M and the model seems to be difficult. However, the algorithm, that we propose to calculate $\hat{\theta}$, is designed in such a way to minimize the impact of M on the results. The detailed implementation of the algorithm is given in section 4.

The estimation error of the Lasso estimator in Theorem 2 is measured in l_∞ -norm. Similarly to Huang et al. (2013), it can be extended to the general l_q -norm, $q \geq 1$. We omit it, because (15) is sufficient to obtain the second main result of the paper (Corollary 3). It states that the thresholded Lasso estimator is model selection consistent, if, additionally to (13) and (14), the nonzero parameters are not too small and the threshold is appropriately chosen. It is a consequence of the fact, which follows from Theorem 2, that the Lasso separates significant parameters from irrelevant ones, i.e. for each $(r, s) \in T$ and $(r', s') \notin T$ we have $|\hat{\theta}_{rs}| > |\hat{\theta}_{r's'}|$ with high probability. However, Corollary 3 does not give a way of choosing the threshold δ , because both endpoints of the interval $[R_n^m, \theta_{min}^*/2]$ are unknown. It is not a surprising fact and has been already observed, for instance, in linear models (Ye and Zhang, 2010, Theorem 8). In section 4 we propose a method of choosing a threshold, that relates to information criteria.

We have already mentioned that there are many approaches to the high-dimensional Ising model. Now we compare conditions, that are sufficient to prove model selection consistency in the current paper to those basing on the likelihood approximation. If we simplify regularity conditions in Theorem 2, Corollary 3 and forget about Monte Carlo parameters in (14), then we have:

- (a) the cone invertibility factor condition is satisfied,
- (b) the sample size should be sufficiently large, that is $n > \bar{d}_0^2 \log d$,
- (c) the nonzero parameters should be sufficiently large, that is $\theta_{min}^* > \sqrt{\frac{\log d}{n}}$.

In Ravikumar et al. (2010, Corollary 1) one needs stronger irrepresentable condition in (a). Their analog of (b) is $n \geq v^3 \log d$, where v is the maximum neighbourhood size. Since v is smaller than \bar{d}_0 , their condition is less restrictive. However, in (c) they require the minimum signal strength to be higher than ours, because it has to be larger than $\sqrt{\frac{v \log d}{n}}$ as distinct from $\sqrt{\frac{\log d}{n}}$ in our paper.

Assumptions in Guo et al. (2010, Theorem 2) are stronger than ours. Indeed, they need irrepresentable condition in (a), \bar{d}_0 in the third power in (b) and additional factor $\sqrt{\bar{d}_0}$ in (c).

In Xue et al. (2012, Corollary 3.1 (2)) model selection consistency of Lasso estimators is also proved with more restrictive conditions than ours. Namely, they are similar to Ravikumar et al. (2010) and Guo et al. (2010) but \bar{d}_0 is reduced in the condition (c). Moreover, they also consider the pseudolikelihood approximation with the SCAD penalty and shows that the condition (a) seems to be superfluous in this case, see Xue et al. (2012, Corollary 3.1 (1)). However, using the SCAD penalty they minimize a nonconvex function to obtain an estimator, so they have to prove that the computed (local) minimizer is the desired theoretic local solution. Their approach can be viewed as a sequence of weighted Lasso problems, so they need auxiliary Lasso procedures to behave well. Therefore, the irrepresentable condition is assumed (Xue et al., 2012, Corollary 3.2).

The conditions sufficient for model selection consistency that are stated in Jalali et al. (2011, Theorem 2) are comparable to ours but also more restrictive. Instead of the condition (a) they consider a similar requirement called the restricted strong convexity condition. It is completed by the restricted strong smoothness condition. Moreover, in the the lower bound in the condition (c) they need an additional factor $\sqrt{d_0}$ as well as the upper bound for θ_{min}^* .

In the proof of Theorem 2 we use methods that are well-known while investigating properties of Lasso estimators as well as some new argumentation. The main novelty (and difficulty) is the use of the Monte Carlo sample, that contains dependent vectors. The first part of our argumentation consists of two steps:

- (i) the first step can be viewed as "deterministic". We apply methods that were developed in Ye and Zhang (2010); Huang and Zhang (2012); Huang et al. (2013) and strongly exploit convexity of the considered problem. These auxiliary results are stated in Lemma 5 and Lemma 6 in the appendix A,
- (ii) the second step is "stochastic". We state a probabilistic inequality that bounds the l_∞ -norm of the derivative of the MCMC approximation (5) at θ^* , that is

$$\nabla \ell_n^m(\theta^*) = -\frac{1}{n} \sum_{i=1}^n J(Y_i) + \frac{\sum_{k=1}^m w_k(\theta^*) J(Y^k)}{\sum_{k=1}^m w_k(\theta^*)} \quad (17)$$

where

$$w_k(\theta) = \frac{\exp[\theta' J(Y^k)]}{h(Y^k)}, \quad k = 1, \dots, m. \quad (18)$$

Notice that (17) contains independent random variables Y_1, \dots, Y_n from the initial sample and the Markov chain Y^1, \dots, Y^m from the MC sample. Therefore, to obtain the exponential inequalities for the l_∞ -norm of (17), which are given in Lemma 7 and Corollary 8 in the appendix A, we apply the MCMC theory. In particular, we frequently use the following Hoeffding's inequality for Markov chains (Miasojedow, 2014, Theorem 1.1).

Theorem 4 *Let Y^1, \dots, Y^m be a reversible Markov chain with a stationary distribution with a density h and a spectral gap $1 - \kappa$. Moreover, let $g : \mathcal{Y} \rightarrow \mathbb{R}$ be a bounded function and $\mu = \mathbb{E}_{Y \sim h} g(Y)$ be a stationary mean value. Then for every $t > 0, m \in \mathbb{N}$ and an initial distribution q*

$$P \left(\left| \frac{1}{m} \sum_{k=1}^m g(Y^k) - \mu \right| > t \right) \leq 2\beta_1 \exp \left(-\frac{\beta_2 m t^2}{|g|_\infty^2} \right),$$

where $|g|_\infty = \sup_{y \in \mathcal{Y}} |g(y)|$ and β_1, β_2 are defined in (9).

The next part of our argumentation relates to the fact that the Hessian of the MCMC approximation at θ^* , that is

$$\nabla^2 \ell_n^m(\theta^*) = \frac{\sum_{k=1}^m w_k(\theta^*) J(Y^k)^{\otimes 2}}{\sum_{k=1}^m w_k(\theta^*)} - \left[\frac{\sum_{k=1}^m w_k(\theta^*) J(Y^k)}{\sum_{k=1}^m w_k(\theta^*)} \right]^{\otimes 2}, \quad (19)$$

is random variable. Similar problems were considered in several papers investigating properties of Lasso estimators in the high-dimensional Ising model (Ravikumar et al., 2010; Guo et al., 2010; Xue et al., 2012) or the Cox model (Huang et al., 2013). We overcome this difficulty by bounding from below the cone invertibility factor $\bar{F}(\xi, T)$ by nonrandom $F(\xi, T)$. Therefore, we need to prove that Hessians of (10) and (5) are close. It is obtained again using the MCMC theory in Lemma 9 and Corollary 10 in the appendix A.

Finally, the proofs of Theorem 2 and Corollary 3 are stated in the appendix B.

4. Details of implementation

In this section we describe in details practical implementation of the algorithm analyzed in the previous section.

The solution of the problem (7) depends on the choice of λ in the penalty term and the parameter ψ in the instrumental distribution. Finding “optimal” λ and ψ is difficult in practice. To overcome this problem we compute a sequence of minimizers $(\hat{\theta}_i)_i$ such that $\hat{\theta}_i$ corresponds to $\lambda = \lambda_i$ and the sequence $(\lambda_i)_i$ is decreasing. In the second step we use the Bayesian Information Criterion (BIC) to choose the final penalty λ . More precisely, we start with the greatest value λ_0 for which the entire vector $\hat{\theta}$ is zero. For each value of λ_i , $i \geq 1$, we set $\psi = \hat{\theta}_{i-1}$ and use the MCMC approximation (7) with Y^1, \dots, Y^m given by a Gibbs sampler with the stationary distribution $p(\cdot | \hat{\theta}_{i-1})$. This scheme exploits warm starts and leads to more stable algorithm. Next, the estimator $\hat{\theta}$ is chosen using BIC that is a popular method of choosing λ in the literature, for instance in Xue et al. (2012). Notice that the function $\ell_n^m(\theta)$ is convex, so we can use proximal gradient algorithms to compute $\hat{\theta}_i$ as a solution of (7) for a given λ_i and ψ . In the current paper we use the FISTA algorithm with backtracking from Beck and Teboulle (2009). The whole procedure is summarized in Algorithm 1.

Algorithm 1 MCMC Lasso for Ising model

Let $\lambda_0 > \lambda_1 > \dots > \lambda_{100}$ and $\psi = 0$.

for $i = 1$ to 100 **do**

 Simulate Y^1, \dots, Y^m using a Gibbs sampler with the stationary distribution $p(y|\psi)$.

 Run the FISTA algorithm to compute $\hat{\theta}_i$ as

$$\arg \min_{\theta} \{ \ell_n^m(\theta) + \lambda_i \|\theta\|_1 \}.$$

 Set $\psi = \hat{\theta}_i$.

end for

Next, set $\hat{\theta} = \hat{\theta}_{i^*}$, where

$$i^* = \arg \min_{1 \leq i \leq 100} \left\{ n \ell_n^m(\hat{\theta}_i) + \log(n) \|\hat{\theta}_i\|_0 \right\}$$

and $\|\theta\|_0$ denotes the number of non-zero elements of θ .

In Algorithm 1 we use 100 values of λ uniformly spaced on the log scale, starting from the largest λ , which corresponds to the empty model. We use $m = 10^3 d$ iteration of the Gibbs sampler. To compute $\ell_n^m(\hat{\theta}_i)$ for $i = 1, \dots, 100$ in BIC we generate one more sample of the size $m = 10^4 d$ using the Gibbs sampler with the stationary distribution $p(\cdot|\hat{\theta}_{50})$.

The important property of our implementation is that the chosen $\psi = \hat{\theta}_{i-1}$ is usually close to $\hat{\theta}_i$, because differences between consecutive λ_i 's are small. In our studies the final estimator $\hat{\theta}$ is the element of the sequence $(\hat{\theta}_i)_i$, that recognizes the true model in the best way, i.e. it minimizes the MCMC approximation and is sparse simultaneously. One believes that the final estimator $\hat{\theta} = \hat{\theta}_i$ is close to θ^* , therefore the chosen $\psi = \hat{\theta}_{i-1}$ should be also similar to θ^* , that makes M in (9) close to one. Finally, notice that conditionally on the previous step our algorithm fits to the framework described in subsection 2.1.

Note that in the first iteration in Algorithm 1 we use an uniform distribution on $\{-1, +1\}^d$ as an instrumental distribution and we can use i.i.d sample Y^1, \dots, Y^m . Therefore, for λ_1 we get $\beta_1 = \beta_2 = 1$. When we compute estimators for λ_i with $i \geq 2$ we use the last sample from the previous step as an initial point, so since the Markov chain generated by the Gibbs sampler is ergodic its initial distribution should be close to $p(y|\hat{\theta}_{i-2})$. Therefore, even without the burn-in time β_1 should be approximately equal to the L^2 -distance between $p(y|\hat{\theta}_{i-1})$ and $p(y|\hat{\theta}_{i-2})$, which seems to be small, because differences between consecutive λ_i 's are small. So, after discarding initial iterations β_1 is further reduced. Our choice of stationary distributions also leads to relatively small variances of importance sampling weights given in (18).

The bounds on β_2 are challenging problem itself and the sharp bounds are available only in very specific cases. Due to that, we do not have the explicit control on β_2 . However, in our procedure the stationary distributions are typically given by sparse Ising models and by Mossel and Sly (2013) the spectral gap depends mostly on the number of existing edges in the graph. Therefore β_2 should not vanish to rapidly with dimensionality of the problem.

Finally, the thresholded estimator is obtained using the Generalized Information Criterion (GIC). For a prespecified set of thresholds Δ we calculate

$$\delta^* = \arg \min_{\delta \in \Delta} \left\{ n\ell_n^m(\hat{\theta}^\delta) + \log(\bar{d})\|\hat{\theta}^\delta\|_0 \right\} ,$$

where $\hat{\theta}^\delta$ is the Lasso estimator $\hat{\theta}$ after thresholding with the level δ . To compute $\ell_n^m(\hat{\theta}^\delta)$ for $\delta \in \Delta$ in GIC we generate the last sample of the size $m = 10^4 d$ using the Gibbs sampler with the stationary distribution $p(\cdot|\hat{\theta})$. To find the optimal threshold we apply GIC that uses larger penalty than BIC. Choosing the threshold in this way should be better in model selection and is recommended, for instance, in Pokarowski and Mielniczuk (2015).

In the rest of this section we discuss computational complexity of our method. The computational cost of a single step of the Gibbs sampler is dominated by computing probability (8), which is of the order $O(r)$, where r is the maximal degree of vertices in a graph related to the stationary distribution. In the paper we focus on estimation of sparse graphs, so the proposed λ_i 's have to be sufficiently large to make $\hat{\theta}_i$'s sparse. Therefore, the degree r is rather small and the computational cost of generating Y^1, \dots, Y^m is of order $O(m)$. Next, we need to compute $\ell_n^m(\theta)$ and its gradient. For an arbitrary Markov chain the cost of these computation is of the order $O(d^2 m)$. But when we use single site updates as in the Gibbs sampler we can reduce it to $O(dm)$ by remembering which coordinate of Y^k where updated. Indeed, if we know that only the r coordinates are updated in the step $Y^k \rightarrow Y^{k+1}$, then

$$\begin{aligned} \theta' J(Y^{k+1}) = \theta' J(Y^k) + \sum_{s: s < r} \left(\theta_{sr} [Y^{k+1}(s)Y^{k+1}(r) - Y^k(s)Y^k(r)] \right) \\ + \sum_{s: s > r} \left(\theta_{rs} [Y^{k+1}(s)Y^{k+1}(r) - Y^k(s)Y^k(r)] \right) . \end{aligned}$$

Finally, it is well-known that FISTA (Beck and Teboulle, 2009) achieves accuracy ϵ in $O(\epsilon^{\frac{1}{2}})$ steps. So, the total cost of computing the solution for single λ_i with precision ϵ is of order $O(\epsilon^{\frac{1}{2}} md)$. The further reduction of the cost can be obtained using sparsity of $\hat{\theta}_{i-1}$ in computing $\ell_n^m(\theta)$ and its gradient, and introducing active variables inside the FISTA algorithm.

5. Numerical experiments

In this section we present efficiency of the proposed method via numerical studies. First, we compare our method to three algorithms, which we have mentioned previously, using simulated data sets. Next we apply our method to the real data example.

5.1. Simulated data

To illustrate the performance of the proposed method we simulate data sets in two scenarios:

- M1: The first 6 vertices are correlated, while the remaining vertices are independent: $\theta_{rs}^* = \pm 2$ for $r < s$ and $s = 2, 3, 4, 5, 6$, other $\theta_{rs}^* = 0$. Thus, the model dimension in this problem is 15. The signs are chosen randomly.




Figure 1: The results for $M1$ model

$M2$: The first 20 vertices have the “chain structure” and the rest are independent: $\theta_{r-1,r}^* = \pm 1$ for $r \leq 20$. Again the signs are chosen randomly and the model dimension is 19.

The model $M1$ corresponds to a dense structure on a small subset of vertices. The model $M2$ is a simple structure which involves relatively large subset of vertices.

We consider the following cases: $d = 20, 50, 100$. So, the considered number of possible edges (parameters of the model) is $\bar{d} = 190, 1225, 4950$, respectively. For $\bar{d} = 190, 1225$ we use $n = 20, 40, 80$ and for $\bar{d} = 4950$ we use $n = 40, 80$.

For each configuration of the model, the number of vertices d and the number of observation n we sample 100 replications of data sets. For sampling a data set we use a final configuration of independent Gibbs samplers of the length 10^6 .

In simulation study we compare our methods to the following methods: the pseudo-likelihood approach, eLASSO proposed by van Borkulo et al. (2014) and the Gaussian approximation from Banerjee et al. (2008). The eLASSO method is based on separate logistic models and it is similar to the method proposed by Ravikumar et al. (2010). For all methods except eLASSO we use two stage procedures, which are analogous to Algorithm 1. Namely, we choose λ by BIC and in the second step we choose the threshold δ by GIC. For the Gaussian approximation we use the same approximate likelihood as in Viallon et al. (2013). For eLASSO for every node we use BIC as in van Borkulo et al. (2014).




Figure 2: The results for $M2$ model

In the comparison we use the following measures of the accuracy. First, to observe the ability of methods to separate true and false edges we compute AUC, where the ROC curve is computed as the threshold δ varies. The estimates are also compared using the false discovery rate (FDR) and the ability of recognizing true edges (denoted by “Power”). The results are summarized in Figure 1 and Figure 2 for models $M1$ and $M2$, respectively.

In the first model we can observe that our algorithm and the Gaussian approximation work very well and comparably. The latter has slightly larger power, but at the price of slightly larger FDR. The dominance of these two methods over the pseudolikelihood estimator and eLASSO is evident. The pseudolikelihood method finds well true edges, but is not able to discard false edges. eLASSO works very poorly in this model.

The second model has a simple structure, so both methods based on the pseudolikelihood approach work much better than in $M1$. The accuracy of the Gaussian approximation is weak in this model. It has substantial problem with finding true edges, especially when d is large and n small. For $n = 20, 40$ our estimator has relatively small FDR and large Power. In this model we can observe that FDR increases as n increases for the Gaussian approximation and for $n = 80$ it reaches about 0.5. The MCML approximation has also “increasing FDR”, but this behaviour is less conspicuous. In fact, for $n = 80$ FDR of our algorithm is still comparable to those of eLASSO and Pseudolikelihood.

Summarizing, it is difficult to indicate the winner algorithm. We can observe that the quality of our algorithm in selecting the true model is satisfactory. Moreover, only this procedure works on a good level in both models and avoids making noticeable mistakes. The Gaussian approximation works well in $M1$, but seems to be the weakest in $M2$. The eLASSO completely fails in $M1$, but its quality in $M2$ is high, especially for large sample sizes. The pseudolikelihood approach in all examples well separates true and false edges, has good power but in comparison with other methods its FDR is too high, so models that it chooses contains many irrelevant edges.

Clearly, by construction the computational cost of our method is larger than its three competitors. However, the whole time that is needed to compute our estimator is reasonable. For instance, for a single data set and $n = 40$ computing the estimator on 3.4 GHZ CPU takes about 15 seconds for $\bar{d} = 190$, 60 seconds for $\bar{d} = 1225$ and 5 minutes for $\bar{d} = 4950$. Since our algorithm uses only sufficient statistics the dependence on n of its computational cost is negligible. Thus, the computational times for $n = 80$ are almost the same.

5.2. CAL500 dataset

We apply our method to “CAL500” dataset (Turnbull et al., 2008). Working with a real data set we consider the Ising model (1) with a linear term (an external field). This modification is motivated by the fact that in practice marginal probabilities of $Y(s)$ being $+1$ or -1 are unknown and should be estimated. The adaptation of our algorithm to this case is straightforward.

For model selection in the Ising model there are no natural measures of the quality of estimates. One would try to compare the prediction ability of obtained estimators, but prediction for the Ising model is challenging itself and results will be biased by the method used to approximate predicted states. Moreover, all considered methods optimize different loss functions, so these loss functions also cannot be used to the honest comparison of the methods. Due to that, we decide to show only the results of our method for the real data example.

The considered data set consists of 174 binary features and 68 numeric features for 502 songs. We skipped the numeric features and apply our method to find the dependence structure between labels. These labels concerning genre, mood or instrument are annotated to songs. We run our algorithm analogously to the case of simulated data and as the result we obtain a sparse graph with 181 edges, see Figure 3. We observe that founded edges are rather intuitive. For instance, among the most positively correlated labels we have labels denoted by 3 and 14 („Emotion-Arousing-Awakening” and „Emotion-Exciting-Thrilling”), 57 and 61 („Song-Like” and „Song-Recommend”) or 22 and 34 („Emotion-Loving-Romantic” and „Emotion-Touching-Loving”). On the other hand, the most negatively correlated labels are: 44 and 45 („Instrument-Female Lead Vocals” and „Instrument - Male Lead Vocals”) or 63 and 64 („Song-Texture Acoustic” and „Song-Texture Electric”).

6. Conclusions

In the paper we consider a problem of structure learning for binary Markov random fields. We base estimation of model parameters on the Lasso penalized Monte Carlo approximation of the likelihood. In the theoretical part of the paper we show that the proposed procedure




Figure 3: The obtained graph for CAL500 dataset. The width of the edges corresponds to the magnitude of $|\hat{\theta}_{rs}^{\delta^*}|$. The edges with higher absolute value are wider. The color denotes the sign of $\hat{\theta}_{rs}^{\delta^*}$: green – positive, red – negative. To improve clarity we do not show edges corresponding to $|\hat{\theta}_{rs}^{\delta^*}| < 0.001$.

reveals the true dependence structure with high probability. The regularity conditions that we need are not restrictive and are weaker than assumptions used in the other approaches based on the likelihood approximation. Moreover, the theoretical results are completed by numerical experiments. They confirm that the MCMC approximation is able to find the true model in a satisfactory way and its quality is comparable or higher than competing algorithms.

The results of the current paper can be easily extended to other discrete Markov random fields. There are also some non-trivial issues that are not discussed in the paper, for instance

investigating the model (1) with predictors (covariates). The evaluation of the prediction error of the procedure is also a difficult task. Clearly, these problems need detailed studies.

Acknowledgments

We would like to thank the associate editor and reviewers for their comments, that have improved the paper. Błażej Miasojedow was supported by the Polish National Science Center grant no. 2015/17/D/ST1/01198. Wojciech Rejchel was supported by the Polish National Science Center grants no. 2014/12/S/ST1/00344 and no. 2015/17/B/ST6/01878.

Appendix A. Auxiliary results

In this section we formulate lemmas that are needed to prove main results of the paper. The roles, that they play, are described in detail at the end of section 3. The first lemma is borrowed from Huang et al. (2013, Lemma 3.1).

Lemma 5 *Let $\tilde{\theta} = \hat{\theta} - \theta^*$, $z^* = |\nabla \ell_n^m(\theta^*)|_\infty$ and*

$$D(\hat{\theta}, \theta) = (\hat{\theta} - \theta)' \left[\nabla \ell_n^m(\hat{\theta}) - \nabla \ell_n^m(\theta) \right].$$

Then

$$(\lambda - z^*)|\tilde{\theta}_{T^c}|_1 \leq D(\hat{\theta}, \theta^*) + (\lambda - z^*)|\tilde{\theta}_{T^c}|_1 \leq (\lambda + z^*)|\tilde{\theta}_T|_1. \quad (20)$$

Besides, for arbitrary $\xi > 1$ on the event

$$\Omega_1 = \left\{ |\nabla \ell_n^m(\theta^*)|_\infty \leq \frac{\xi - 1}{\xi + 1} \lambda \right\} \quad (21)$$

the random vector $\tilde{\theta}$ belongs to the cone $\mathcal{C}(\xi, T)$.

Proof The proof is the same as the proof of Huang et al. (2013, Lemma 3.1). It is quoted here to make the paper complete.

Convexity of the MCMC approximation $\ell_n^m(\theta)$ easily implies the first inequality in (20). The same property combined with convexity of the Lasso penalty gives us that zero has to belong to the subgradient of (7) at the minimizer $\hat{\theta}$, i.e.

$$\begin{cases} \nabla_{rs} \ell_n^m(\hat{\theta}) = -\lambda \text{sign}(\hat{\theta}_{rs}), & \text{if } \hat{\theta}_{rs} \neq 0 \\ |\nabla_{rs} \ell_n^m(\hat{\theta})| \leq \lambda, & \text{if } \hat{\theta}_{rs} = 0, \end{cases} \quad (22)$$

where we use $\nabla \ell_n^m(\theta) = \left(\nabla_{rs} \ell_n^m(\hat{\theta}) \right)_{r < s}$ and $\text{sign}(t) = 1$ for $t > 0$, $\text{sign}(t) = -1$ for $t < 0$, $\text{sign}(t) = 0$ for $t = 0$. Using (22) and properties of the l_1 -norm we obtain that

$$\begin{aligned} D(\hat{\theta}, \theta^*) &= \sum_{(r,s) \in T} \tilde{\theta}_{rs} \nabla_{rs} \ell_n^m(\theta^* + \tilde{\theta}) + \sum_{(r,s) \in T^c} \hat{\theta}_{rs} \nabla_{rs} \ell_n^m(\theta^* + \tilde{\theta}) - \tilde{\theta}' \nabla \ell_n^m(\theta^*) \\ &\leq \lambda \sum_{(r,s) \in T} |\tilde{\theta}_{rs}| - \lambda \sum_{(r,s) \in T^c} |\hat{\theta}_{rs}| + |\tilde{\theta}|_1 z^* \\ &\leq \lambda |\tilde{\theta}_T|_1 - \lambda |\tilde{\theta}_{T^c}|_1 + z^* |\tilde{\theta}_T|_1 + z^* |\tilde{\theta}_{T^c}|_1 \\ &= (\lambda + z^*) |\tilde{\theta}_T|_1 + (z^* - \lambda) |\tilde{\theta}_{T^c}|_1. \end{aligned}$$

Thus, the second inequality in (20) is also established. To prove the last claim of the lemma notice that on the event Ω_1 we obtain from (20)

$$|\tilde{\theta}_{T^c}|_1 \leq \frac{\lambda + z^*}{\lambda - z^*} |\tilde{\theta}_T|_1 \leq \xi |\tilde{\theta}_T|_1.$$

■

The second lemma is an adaptation of Huang et al. (2013, Theorem 3.1) to our problem.

Lemma 6 *Let $\xi > 1$. Moreover, let us denote $\tau = \frac{(\xi+1)\bar{d}_0\lambda}{F(\xi, T)}$ and an event*

$$\Omega_2 = \{\tau < e^{-1}\}. \quad (23)$$

Then $\Omega_1 \cap \Omega_2 \subset A$, where

$$A = \left\{ |\hat{\theta} - \theta^*|_\infty \leq \frac{2\xi e^\eta \lambda}{(\xi + 1)\bar{F}(\xi, T)} \right\}, \quad (24)$$

where $\eta < 1$ is the smaller solution of the equation $\eta e^{-\eta} = \tau$.

Proof Suppose we are on the event $\Omega_1 \cap \Omega_2$. Denote again $\tilde{\theta} = \hat{\theta} - \theta^*$ and notice that $\theta = \frac{\tilde{\theta}}{|\tilde{\theta}|_1} \in \mathcal{C}(\xi, T)$ by Lemma 5. Consider the function

$$g(t) = \theta' \nabla \ell_n^m(\theta^* + t\theta) - \theta' \nabla \ell_n^m(\theta^*)$$

for each $t \geq 0$. This function is nondecreasing, because $\ell_n^m(\cdot)$ is convex. Thus, we obtain $g(t) \leq g(|\tilde{\theta}|_1)$ for every $t \in (0, |\tilde{\theta}|_1)$. On the event Ω_1 and from Lemma 5 we have that

$$\theta' [\nabla \ell_n^m(\theta^* + t\theta) - \nabla \ell_n^m(\theta^*)] + \frac{2\lambda}{\xi + 1} |\theta_{T^c}|_1 \leq \frac{2\lambda\xi}{\xi + 1} |\theta_T|_1. \quad (25)$$

In further argumentation we consider all nonnegative t satisfying (25), that is an interval $[0, \tilde{t}]$ for some $\tilde{t} > 0$. Proceeding similarly to the proof of Huang et al. (2013, Lemma 3.2) we obtain

$$t\theta' [\nabla \ell_n^m(\theta^* + t\theta) - \nabla \ell_n^m(\theta^*)] \geq t^2 \exp(-\gamma_{t\theta}) \theta' \nabla^2 \ell_n^m(\theta^*) \theta, \quad (26)$$

where $\gamma_{t\theta} = t \max_{k,l} |\theta' J(Y^k) - \theta' J(Y^l)| \leq 2t$, because $J(Y^k) = (Y^k(r)Y^k(s))_{rs}$ and $|\theta|_1 = 1$. Therefore, the right-hand side in (26) can be lower bounded by

$$t^2 \exp(-2t) \theta' \nabla^2 \ell_n^m(\theta^*) \theta. \quad (27)$$

Using the definition of $\bar{F}(\xi, T)$, the fact that $\theta \in \mathcal{C}(\xi, T)$, the bound (27) and (25) we obtain

$$\begin{aligned} t \exp(-2t) \frac{\bar{F}(\xi, T) |\theta_T|_1^2}{\bar{d}_0} &\leq t \exp(-2t) \theta' \nabla^2 \ell_n^m(\theta^*) \theta \\ &\leq \theta' [\nabla \ell_n^m(\theta^* + t\theta) - \nabla \ell_n^m(\theta^*)] \\ &\leq \frac{2\lambda\xi}{\xi + 1} |\theta_T|_1 - \frac{2\lambda}{\xi + 1} |\theta_{T^c}|_1 \\ &\leq \lambda(\xi + 1) |\theta_T|_1^2 / 2. \end{aligned}$$

So, every t satisfying (25) has to fulfill the inequality $2t \exp(-2t) \leq \tau$. In particular, $2\tilde{t} \exp(-2\tilde{t}) \leq \tau$. We are on Ω_2 , so it implies that $2\tilde{t} \leq \eta$, where η is the smaller solution of the equation $\eta \exp(-\eta) = \tau$. We know also that $|\tilde{\theta}|_1 \leq \tilde{t}$, so

$$\begin{aligned} |\tilde{\theta}|_1 \exp(-\eta) &\leq \tilde{t} \exp(-2\tilde{t}) \leq \frac{\tilde{t} \exp(-2\tilde{t}) \theta' \nabla^2 \ell_n^m(\theta^*) \theta}{\bar{F}(\xi, T) |\theta_T|_1 |\theta|_\infty} \\ &\leq \frac{\theta' [\nabla \ell_n^m(\theta^* + \tilde{t}\theta) - \nabla \ell_n^m(\theta^*)]}{\bar{F}(\xi, T) |\theta_T|_1 |\theta|_\infty} \\ &\leq \frac{2\lambda\xi}{(\xi + 1) \bar{F}(\xi, T) |\theta|_\infty}, \end{aligned}$$

where we have used bounds (27) and (25). Using the equality $|\theta|_\infty = \frac{|\tilde{\theta}|_\infty}{|\tilde{\theta}|_1}$, we finish the proof. \blacksquare

Lemma 7 *For every natural n, m and positive t*

$$P(|\nabla \ell_n^m(\theta^*)|_\infty \leq t) \geq 1 - 2\bar{d} \exp(-nt^2/8) - \beta_1 \exp\left(-\frac{m\beta_2}{4M^2}\right) - 2\bar{d}\beta_1 \exp\left(-\frac{t^2 m\beta_2}{64M^2}\right).$$

Proof We can rewrite $\nabla \ell_n^m(\theta^*)$ as

$$\nabla \ell_n^m(\theta^*) = - \left[\frac{1}{n} \sum_{i=1}^n J(Y_i) - \frac{\nabla C(\theta^*)}{C(\theta^*)} \right] + \frac{\frac{1}{m} \sum_{k=1}^m w_k(\theta^*) \left[J(Y^k) - \frac{\nabla C(\theta^*)}{C(\theta^*)} \right]}{\frac{1}{m} \sum_{k=1}^m w_k(\theta^*)} \quad (28)$$

Notice that the first term in (28) depends only on the initial sample Y_1, \dots, Y_n and is an average of i.i.d random variables. The second term depends only on the MCMC sample Y^1, \dots, Y^m . We start the analysis with the former one. Using Hoeffding's inequality we obtain for each natural n , positive t and a pair of indices $r < s$

$$P\left(\left| \frac{1}{n} \sum_{i=1}^n J_{rs}(Y_i) - \frac{\nabla_{rs} C(\theta^*)}{C(\theta^*)} \right| > t/2\right) \leq 2 \exp(-nt^2/8).$$

Therefore, by the union bound we have

$$P\left(\left| \frac{1}{n} \sum_{i=1}^n J(Y_i) - \frac{\nabla C(\theta^*)}{C(\theta^*)} \right|_\infty > t/2\right) \leq 2\bar{d} \exp(-nt^2/8). \quad (29)$$

Next, we investigate the second expression in (28). Its denominator is an average that depends on the Markov chain. To handle it we can use Theorem 4 in section 3. Notice that $\frac{\exp[(\theta^*)' J(y)]}{h(y)} \leq MC(\theta^*)$ for every y and $\mathbb{E}_{Y \sim h} \frac{\exp[(\theta^*)' J(Y)]}{h(Y)} = C(\theta^*)$. Therefore, for every n, m

$$P\left(\frac{1}{m} \sum_{k=1}^m w_k(\theta^*) \geq C(\theta^*)/2\right) \geq 1 - \beta_1 \exp\left(-\frac{m\beta_2}{4M^2}\right). \quad (30)$$

Finally, we bound the l_∞ -norm of the numerator of the second term in (28). We fix a pair of indices $r < s$. It is not difficult to calculate that

$$\mathbb{E}_{Y \sim h} \left[\frac{\exp[(\theta^*)'J(Y)]}{h(Y)} \left(J_{rs}(Y) - \frac{\nabla_{rs}C(\theta^*)}{C(\theta^*)} \right) \right] = 0$$

and for every y

$$\frac{\exp((\theta^*)'J(y))}{h(y)} \left| J_{rs}(y) - \frac{\nabla_{rs}C(\theta^*)}{C(\theta^*)} \right| \leq 2MC(\theta^*).$$

From Theorem 4 we obtain for every positive t

$$\left| \frac{1}{m} \sum_{k=1}^m w_k(\theta^*) \left[J_{rs}(Y^k) - \frac{\nabla_{rs}C(\theta^*)}{C(\theta^*)} \right] \right| \geq tC(\theta^*)/4$$

with probability at least $1 - 2\beta_1 \exp\left(-\frac{t^2 m \beta_2}{64M^2}\right)$. Using union bounds we estimate the numerator of the second expression in (28). This fact and (30) imply that for every positive t and natural n, m with probability at least $1 - \beta_1 \exp\left(-\frac{m \beta_2}{4M^2}\right) - 2\bar{d}\beta_1 \exp\left(-\frac{t^2 m \beta_2}{64M^2}\right)$ we have

$$\left| \frac{\frac{1}{m} \sum_{k=1}^m w_k(\theta^*) \left[J(Y^k) - \frac{\nabla C(\theta^*)}{C(\theta^*)} \right]}{\frac{1}{m} \sum_{k=1}^m w_k(\theta^*)} \right|_\infty \leq t/2. \quad (31)$$

Taking (29) and (31) together we finish the proof. ■

Corollary 8 *Let $\varepsilon > 0, \xi > 1$ and*

$$\lambda = \frac{\xi + 1}{\xi - 1} \max \left(2\sqrt{\frac{2 \log(2\bar{d}/\varepsilon)}{n}}, 8M\sqrt{\frac{\log[(2\bar{d} + 1)\beta_1/\varepsilon]}{m\beta_2}} \right).$$

Conditions (13) and (14) imply

$$P(\Omega_1) \geq 1 - 2\varepsilon.$$

Proof We take $t = \frac{\xi-1}{\xi+1}\lambda$ in Lemma 7. ■

Lemma 9 *For every n, m and positive t*

$$\begin{aligned} P(|\nabla^2 \ell_n^m(\theta^*) - \nabla^2 \log C(\theta^*)|_\infty \leq t) &\geq 1 - 2\beta_1 \exp\left(-\frac{m\beta_2}{4M^2}\right) \\ &\quad - 2\bar{d}(\bar{d} + 1)\beta_1 \exp\left(-\frac{t^2 m \beta_2}{256M^2}\right). \end{aligned} \quad (32)$$

Proof To estimate the l_∞ -norm of the matrix difference in (32) we bound l_∞ -norms of two matrices:

$$\frac{\sum_{k=1}^m w_k(\theta^*) J(Y^k)^{\otimes 2}}{\sum_{k=1}^m w_k(\theta^*)} - \frac{\nabla^2 C(\theta^*)}{C(\theta^*)} = \frac{\frac{1}{m} \sum_{k=1}^m w_k(\theta^*) \left[J(Y^k)^{\otimes 2} - \frac{\nabla^2 C(\theta^*)}{C(\theta^*)} \right]}{\frac{1}{m} \sum_{k=1}^m w_k(\theta^*)} \quad (33)$$

and

$$\left[\frac{\frac{1}{m} \sum_{k=1}^m w_k(\theta) J(Y^k)}{\frac{1}{m} \sum_{k=1}^m w_k(\theta)} \right]^{\otimes 2} - \left[\frac{\nabla C(\theta^*)}{C(\theta^*)} \right]^{\otimes 2}. \quad (34)$$

The denominator of the right-hand side of (33) has been estimated in the proof of Lemma 7, so we bound the numerator. We can calculate that

$$\mathbb{E}_{Y \sim h} \left[\frac{\exp[(\theta^*)' J(Y)]}{h(Y)} \left(J(Y)^{\otimes 2} - \frac{\nabla^2 C(\theta^*)}{C(\theta^*)} \right) \right] = 0$$

and for every y and two pairs of indices $r < s, r' < s'$ we have

$$\frac{\exp((\theta^*)' J(y))}{h(y)} \left| J_{rs}(y) J_{r's'}(y) - \frac{\nabla_{rs,r's'}^2 C(\theta^*)}{C(\theta^*)} \right| \leq 2MC(\theta^*).$$

Using the union bound and (30) we upper-bound the l_∞ -norm of the right-hand side of (33) by $t/2$ with probability at least

$$1 - \beta_1 \exp\left(-\frac{m\beta_2}{4M^2}\right) - 2\bar{d}^2 \beta_1 \exp\left(-\frac{t^2 m \beta_2}{64M^2}\right).$$

The last step of the proof is handling with the l_∞ -norm of (34). This expression can be upper-bounded by

$$\left| \frac{\sum_{k=1}^m w_k(\theta) J(Y^k)}{\sum_{k=1}^m w_k(\theta)} - \frac{\nabla C(\theta^*)}{C(\theta^*)} \right|_\infty \left(\left| \frac{\sum_{k=1}^m w_k(\theta) J(Y^k)}{\sum_{k=1}^m w_k(\theta)} \right|_\infty + \left| \frac{\nabla C(\theta^*)}{C(\theta^*)} \right|_\infty \right). \quad (35)$$

The first term in (35) has been bounded with high probability in the proof of Lemma 7. The remaining two can be easily estimated by one. Therefore, for every positive t the l_∞ -norm of (34) is not greater than $t/2$ with probability at least

$$1 - \beta_1 \exp\left(-\frac{m\beta_2}{4M^2}\right) - 2\bar{d} \beta_1 \exp\left(-\frac{t^2 m \beta_2}{256M^2}\right).$$

Putting together this fact and the bound of (33) we finish the proof. \blacksquare

Corollary 10 *If (14), then for every $\varepsilon > 0$ the following inequality*

$$\bar{F}(\xi, T) \geq F(\xi, T) - 16\bar{d}_0(1 + \xi)^2 M \sqrt{\frac{\log [2\bar{d}(\bar{d} + 1)\beta_1/\varepsilon]}{m\beta_2}}$$

has probability at least $1 - 2\varepsilon$.

Proof We take

$$t = 16M \sqrt{\frac{\log [2\bar{d}(\bar{d} + 1)\beta_1/\varepsilon]}{m\beta_2}}$$

in Lemma 9 and use Huang et al. (2013, Lemma 4.1 (ii)). ■

Appendix B. Proofs of main results

Proof [Proof of Theorem 2] Fix $\varepsilon > 0, \xi > 1$ and denote $\gamma = \gamma(\xi) = \frac{\alpha(\xi)-2}{\alpha(\xi)} \in (0, 1)$. First, from Corollary 8 we know that $P(\Omega_1) \geq 1 - 2\varepsilon$. Using the condition (14) we obtain that

$$F(\xi, T) - 16\bar{d}_0(1 + \xi)^2 M \sqrt{\frac{\log [2\bar{d}(\bar{d} + 1)\beta_1/\varepsilon]}{m\beta_2}} \geq \gamma F(\xi, T).$$

Therefore, from Corollary 10 we have that $P(\bar{F}(\xi, T) \geq \gamma F(\xi, T)) \geq 1 - 2\varepsilon$. It is not difficult to calculate that

$$\frac{(1 + \xi)\bar{d}_0\lambda}{\gamma F(\xi, T)} \leq e^{-1},$$

so we have also $P(\Omega_2) \geq 1 - 2\varepsilon$. To finish the proof we use Lemma 6 (with $\eta = 1$ for simplicity) and again bound $\bar{F}(\xi, T)$ from below by $\gamma F(\xi, T)$ in the event A defined in (24). ■

Proof [Proof of Corollary 3] The proof is a simple consequence of the uniform bound (15) obtained in Theorem 2. Indeed, for an arbitrary pair of indices $(r, s) \notin T$ we obtain

$$|\hat{\theta}_{rs}| = |\hat{\theta}_{rs} - \theta_{rs}^*| \leq R_n^m,$$

so $(r, s) \notin \hat{T}$. Analogously, if we take a pair $(r, s) \in T$, then

$$|\hat{\theta}_{rs}| \geq |\theta_{rs}^*| - |\hat{\theta}_{rs} - \theta_{rs}^*| > 2\delta - R_n^m \geq \delta.$$

■

References

- Yves F. Atchadé, Gersende Fort, and Eric Moulines. On perturbed proximal gradient algorithms. *Journal of Machine Learning Research*, 18:310–342, 2017.
- Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, 2008.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2:183–202, 2009.
- Julian Besag. Spacial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B*, 36:192–236, 1974.
- Peter J Bickel, Ya’acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37:1705–1732, 2009.
- Guy Bresler, Elchanan Mossel, and Allan Sly. Reconstruction of Markov random fields from samples: Some observations and algorithms. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 343–356. Springer, 2008.
- Peter Bühlmann and Sara van de Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Series in Statistics, New York: Springer, 2011.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, 2001.
- Charles J Geyer. On the convergence of Monte Carlo maximum likelihood calculations. *Journal of the Royal Statistical Society. Series B (Methodological)*, 56:261–274, 1994.
- Charles J Geyer and Elizabeth A Thompson. Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 54:657–699, 1992.
- Jian Guo, Elizaveta Levina, George Michailidis, and Ji Zhu. Joint structure estimation for categorical Markov networks. *Technical report*, 2010.
- Holger Höfling and Robert Tibshirani. Estimation of sparse binary pairwise Markov networks using pseudolikelihoods. *Journal of Machine Learning Research*, 10:883–906, 2009.
- Jean Honorio. Convergence Rates of Biased Stochastic Optimization for Learning Sparse Ising Models. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1099–1106, 2012.
- Jian Huang and Cun-Hui Zhang. Estimation and selection via absolute penalized convex minimization and its multistage adaptive applications. *Journal of Machine Learning Research*, 13:1839–1864, 2012.
- Jian Huang, Tingni Sun, Zhiliang Ying, Yi Yu, and Cun-Hui Zhang. Oracle inequalities for the lasso in the Cox model. *Annals of Statistics*, 41(3):1142–1165, 2013.

- Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925.
- Ali Jalali, Christopher C Johnson, and Pradeep K Ravikumar. On learning discrete graphical models using greedy methods. In *Advances in Neural Information Processing Systems*, pages 1935–1943, 2011.
- Ioannis Kontoyiannis and Sean P Meyn. Geometric ergodicity and the spectral gap of non-reversible markov chains. *Probability Theory and Related Fields*, 154:327–339, 2012.
- Su-In Lee, Varun Ganapathi, and Daphne Koller. Efficient Structure Learning of Markov Networks using l_1 -Regularization. In *Advances in Neural Information Processing Systems*, pages 817–824, 2006.
- Błażej Miasojedow. Hoeffding’s inequalities for geometrically ergodic markov chains on general state space. *Statistics & Probability Letters*, 87:115–120, 2014.
- Elchanan Mossel and Allan Sly. Exact thresholds for Ising–Gibbs samplers on general graphs. *The Annals of Probability*, 41:294–328, 2013.
- Sahand Negahban, Bin Yu, Martin J Wainwright, and Pradeep K Ravikumar. A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers. In *Advances in Neural Information Processing Systems*, pages 1348–1356, 2009.
- Piotr Pokarowski and Jan Mielniczuk. Combined l_1 and greedy l_0 penalized least squares for linear model selection. *Journal of Machine Learning Research*, 16:961–992, 2015.
- Pradeep Ravikumar, Martin J Wainwright, and John Lafferty. High-dimensional Ising model selection using l_1 -regularized logistic regression. *The Annals of Statistics*, 38:1287–1319, 2010.
- Gareth O Roberts and Jeffrey S Rosenthal. Geometric ergodicity and hybrid Markov chains. *Electronic Communications in Probability*, 2:13–25, 1997.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58:267–288, 1996.
- Douglas Turnbull, Luke Barrington, David Torres, and Gert Lanckriet. Semantic annotation and retrieval of music and sound effects. *IEEE Transactions on Audio, Speech, and Language Processing*, 16:467–476, 2008.
- Claudia D. van Borkulo, Denny Borsboom, Sacha Epskamp, Tessa F. Blanken, Lynn Boschloo, Robert A. Schoevers, and Lourens J. Waldorp. A new method for constructing networks from binary data. *Scientific Reports*, 4, 2014.
- Sara van de Geer. High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, 36:614–645, 2008.
- Vivian Viallon, Onureena Banerjee, Eric Jouglu, Grégoire Rey, and Joel Coste. Empirical comparison study of approximate methods for structure selection in binary graphical models. *Biometrical Journal*, 56:307–331, 2013.

- Martin J. Wainwright and Michael I. Jordan. Log-determinant relaxation for approximate inference in discrete markov random fields. *IEEE Transactions on Signal Processing*, 54:2099–2109, 2006.
- Lingzhou Xue, Hui Zou, and Tianxi Cai. Nonconcave penalized composite conditional likelihood estimation of sparse Ising models. *The Annals of Statistics*, 40:1403–1429, 2012.
- Fei Ye and Cun-Hui Zhang. Rate Minimavity of the Lasso and Dantzig Selector for the l_q loss in l_r Balls. *Journal of Machine Learning Research*, 11:3519–3540, 2010.
- Tong Zhang. On the consistency of feature selection using greedy least squares regression. *Journal of Machine Learning Research*, 10:555–568, 2009.
- Peng Zhao and Bin Yu. On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2563, 2006.
- Shuheng Zhou. Thresholding procedures for high dimensional variable selection and statistical estimation. In *Advances in Neural Information Processing Systems*, pages 2304–2312, 2009.