



UvA-DARE (Digital Academic Repository)

Bayesian structure learning for Markov Random Fields with a spike and slab prior

Chen, Y.; Welling, M.

Publication date

2012

Document Version

Final published version

Published in

Uncertainty in Artificial Intelligence: proceedings of the Twenty-Eight conference (2012): August 15-17, 2012, Catalina Island, CA

[Link to publication](#)

Citation for published version (APA):

Chen, Y., & Welling, M. (2012). Bayesian structure learning for Markov Random Fields with a spike and slab prior. In N. de Freitas, & K. Murphy (Eds.), *Uncertainty in Artificial Intelligence: proceedings of the Twenty-Eight conference (2012): August 15-17, 2012, Catalina Island, CA* (pp. 174-184). AUAI Press. <http://www.auai.org/uai2012/proceedings.pdf>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Bayesian Structure Learning for Markov Random Fields with a Spike and Slab Prior

Yutian Chen

Department of Computer Science
University of California, Irvine
Irvine, CA 92697

Max Welling

Department of Computer Science
University of California, Irvine
Irvine, CA 92697

Abstract

In recent years a number of methods have been developed for automatically learning the (sparse) connectivity structure of Markov Random Fields. These methods are mostly based on L_1 -regularized optimization which has a number of disadvantages such as the inability to assess model uncertainty and expensive cross-validation to find the optimal regularization parameter. Moreover, the model's predictive performance may degrade dramatically with a sub-optimal value of the regularization parameter (which is sometimes desirable to induce sparseness). We propose a fully Bayesian approach based on a “spike and slab” prior (similar to L_0 regularization) that does not suffer from these shortcomings. We develop an approximate MCMC method combining Langevin dynamics and reversible jump MCMC to conduct inference in this model. Experiments show that the proposed model learns a good combination of the structure and parameter values without the need for separate hyper-parameter tuning. Moreover, the model's predictive performance is much more robust than L_1 -based methods with hyper-parameter settings that induce highly sparse model structures.

1 Introduction

Undirected probabilistic graphical models, also known as Markov Random Fields (MRFs), have been widely used in a large variety of domains including computer vision (Li, 2009), natural language processing (Sha and Pereira, 2003), and social networks (Robins et al., 2007). The structure of the model is defined through a set of features defined on subsets of random variables. Automated methods to select relevant features are becoming increasingly important in a time where the proliferation of sensors make it possi-

ble to measure a multitude of data-attributes. There is also an increasing interest in *sparse* model structures because they help against overfitting and are computationally more tractable than dense model structures.

In this paper we focus on a particular type of MRF, called a log-linear model, where structure learning or feature selection is integrated with parameter estimation. Although structure learning has been extensively studied for directed graphical models, it is typically more difficult for undirected models due to the intractable normalization term of the probability distribution, known as the partition function. Traditional algorithms apply only to restricted types of structures with low tree-width (Andrew and Gao, 2007; Tsuruoka et al., 2009; Hu et al., 2009) or special models such as Gaussian graphical models (Jones et al., 2005) so that accurate inference can be conducted efficiently.

For an arbitrary structure, various methods have been proposed in the literature, generally categorized into two approaches. One approach is based on separate tests on an edge or the neighbourhood of a node so that there is no need to compute the joint distribution (Wainwright et al., 2007; Bresler et al., 2008; Ravikumar et al., 2010). The other approach is based on maximum likelihood estimation (MLE) with a sparsity inducing criterion. These methods require approximate inference algorithms in order to estimate the log-likelihood such as Gibbs sampling (Della Pietra et al., 1997), loopy belief propagation (Lee et al., 2006; Zhu et al., 2010), or pseudo-likelihood (Höfling and Tibshirani, 2009). A popular choice of such a criterion is L_1 regularization (Riezler and Vasserman, 2004; Dudik et al., 2004) which enjoys several good properties such as a convex objective function and a consistency guarantee. However, L_1 -regularized MLE is usually sensitive to the choice the regularization strength, and these optimization-based methods cannot provide a credible interval for the learned structure. Also, in order to learn a sparse structure, a strong penalty has to be imposed on all the edges which usually results in suboptimal parameter values.

We will follow a third approach to MRF structure learning in a fully Bayesian framework which has not been ex-

plored yet. The Bayesian approach considers the structure of a graphical model as random. Inference in a Bayesian model provides inherent regularization, and offers a fully probabilistic characterization of the underlying structure. It was shown in Park and Casella (2008) that Bayesian models with a Gaussian or Laplace prior distribution (corresponding to L_2 or L_1 regularization) do not exhibit sparse structure. Mohamed et al. (2011) proposes to use a “spike and slab” prior for learning *directed* graphical models which corresponds to the ideal L_0 regularization. This model exhibits better robustness against over-fitting than the related L_1 approaches. Unlike the Laplace/Gaussian prior, the posterior distribution over parameters for a “spike and slab” prior is no longer guaranteed to be unimodal. However, approximate inference methods have been successfully applied in the context of directed models using MCMC (Mohamed et al., 2011) and expectation propagation (Hernández-Lobato et al., 2010).

Unfortunately, Bayesian inference for MRFs is much harder than for directed networks due to the partition function. This feature renders even MCMC sampling intractable which caused some people to dub these problems “double intractability” (Murray et al., 2006). Nevertheless, variational methods (Parise and Welling, 2006; Qi et al., 2005) and MCMC methods (Murray and Ghahramani, 2004) have been successfully explored for approximate inference when the model structure is fixed.

We propose a Bayesian structure learning method with a spike and slab prior for MRFs and devise an approximate MCMC method to draw samples of both the model structure and the model parameters by combining a modified Langevin sampler with a reversible jump MCMC method. Experiments show that the posterior distribution estimated by our inference method matches the actual distribution very well. Moreover, our method offers better robustness to both under-fitting and over-fitting than L_1 -regularized methods. A related but different application of the spike and slab distribution in MRFs is shown in Courville et al. (2011) for modelling hidden random variables.

This paper is organized as follows: we first introduce a hierarchical Bayesian model for MRF structure learning in section 2 and then describe an approximate MCMC method in section 3, 4, and 5 to draw samples for the model parameters, structure, and other hyper-parameters respectively. Experiments are conducted in section 6 on two simulated data sets and a real-world dataset, followed by a discussion section.

2 Learning the Structure of MRFs as Bayesian Inference

The probability distribution of a MRF is defined by a set of potential functions. Consider a widely used family of

MRFs with log-linear parametrization:

$$P(\mathbf{x}|\boldsymbol{\theta}) = \frac{1}{Z(\boldsymbol{\theta})} \exp\left(\sum_{\alpha} \theta_{\alpha} f_{\alpha}(\mathbf{x}_{\alpha})\right) \quad (1)$$

where each potential function is defined as the exponential of the product between a feature function f_{α} of a set of variables \mathbf{x}_{α} and an associated parameter θ_{α} . Z is called the partition function. All the variables in the scope of a potential function form a clique in their graphical representation. When a parameter θ_{α} has a value of zero, we could equivalently remove feature f_{α} and all the edges between variables in \mathbf{x}_{α} (if these variables are not also in the scope of other features) without changing the distribution of \mathbf{x} . Therefore, by learning the parameters of this MRF model we can simultaneously learn the structure of a model if we allow some parameters to go to zero.

The Bayesian learning approach to graphical models considers parameters as a random variable subject to a prior. Given observed data, we can infer the posterior distribution of the parameters and their connectivity structure. Two commonly used priors, the Laplace and the Gaussian distribution, correspond to the L_1 and L_2 penalties respectively in the associated optimization-based approach. Although a model learned by L_1 -penalized MLE is able to obtain a sparse structure, the full Bayesian treatment usually results in a fully connected model with many weak edges as observed in Park and Casella (2008), without special approximate assumptions like the ones in Lin and Lee (2006). We propose to use the “spike and slab” prior to learn a sparse structure for MRFs in a fully Bayesian approach. The spike and slab prior (Mitchell and Beauchamp, 1988; Ishwaran and Rao, 2005) is a mixture distribution which consists of a point mass at zero (spike) and a widely spread distribution (slab):

$$P(\theta_{\alpha}) = (1 - p_0)\delta(\theta_{\alpha}) + p_0\mathcal{N}(\theta_{\alpha}; 0, \sigma_0^2) \quad (2)$$

where $p_0 \in [0, 1]$, δ is the Dirac delta function, and σ_0 is usually large enough to be uninformative. The spike component controls the sparsity of the structure in the posterior distribution while the slab component usually applies a mild shrinkage effect on the parameters of the existing edges even in a highly sparse model. This type of *selective* shrinkage is different from the global shrinkage imposed by L_1/L_2 regularization, and enjoys benefits in parameter estimation as demonstrated in the experiment section.

The Bayesian MRF with the spike and slab prior is formulated as follows:

$$\begin{aligned} P(\mathbf{x}|\boldsymbol{\theta}) &= \frac{1}{Z(\boldsymbol{\theta})} \exp\left(\sum_{\alpha} \theta_{\alpha} f_{\alpha}(\mathbf{x}_{\alpha})\right) \\ \theta_{\alpha} &= Y_{\alpha} A_{\alpha} \\ Y_{\alpha} &\sim \text{Bern}(p_0) \quad p_0 \sim \text{Beta}(a, b) \\ A_{\alpha} &\sim \mathcal{N}(0, \sigma_0^2) \quad \sigma_0^{-2} \sim \Gamma(c, d) \end{aligned} \quad (3)$$




Figure 1: Illustration of the MCMC for A_α and Y_α .

where \mathbf{x} is a set of state variables and a , b , c , and d are hyper-parameters. In the experiments we will use pairwise features in which case $\alpha = (i, j)$ plus bias terms given by $\sum_i \theta_i f_i(x_i)$ in the expression for the log-probability. We will use a normal prior $\theta_i \sim \mathcal{N}(0, \sigma_b^2)$ for these bias terms. σ_b is chosen to be large enough to act as an uninformative prior. Y_α is a binary random variable representing the existence of the edges in the clique \mathbf{x}_α , and A_α is the actual value of the parameter θ_α when the edges are instantiated. It is easy to observe that given p_0 and σ_0 , θ_α has the same distribution as in equation 2. We use a hierarchical model for θ so that the inference will be insensitive to the choice of the hyper-parameters. In fact, experiments show that with a simple setting of the hyper-parameters, proper values of the sparsity parameter p_0 and the variance σ_0 are learned automatically by our model for all the data sets without the necessity of cross-validation.

Unlike the optimization-based methods which estimate a single structure, the Bayesian approach expresses uncertainty about the existence of edges through its posterior distribution, $P(\mathbf{Y}|\mathcal{D})$. We have applied a simple thresholding on $P(Y_\alpha|\mathcal{D})$ for edge detection in the experiments although more sophisticated methods can conceivably give better results.

Standard approaches to posterior inference do not work for Bayesian MRFs because it is intractable to compute the probability of a state \mathbf{x} (due to the intractability of the partition function). We devised an approximate MCMC algorithm for inference, where we draw samples of the continuous variable A_α by a modified Langevin dynamics algorithm, and samples of the discrete variable Y_α jointly with A_α by a reversible jump MCMC method, as illustrated in Figure 1 and explained in the following sections.

3 Sampling Parameter Values by Langevin Dynamics

Given \mathbf{Y} , p_0 , σ_0 , and an observed data set $\mathcal{D} = \{\mathbf{x}^{(m)}\}, m = 1 \dots N$, the conditional distribution of parameters $\{A_\alpha : Y_\alpha = 1\}$ is the posterior distribution of an MRF with a fixed edge set induced by $\{\alpha : Y_\alpha = 1\}$ and an independent Gaussian prior $\mathcal{N}(0, \sigma_0^2)$. We consider draw-

ing samples of A_α with fixed \mathbf{Y} in this section and will use θ_α and A_α interchangeably to refer to a nonzero parameter. Even for an MRF model with a fixed structure, MCMC is still intractable. Approximate MCMC methods have been discussed in Murray and Ghahramani (2004) among which Langevin Monte Carlo (LMC) with “brief sampling” to compute the required expectations in the gradients, shows good performance. (Welling and Teh, 2011) further shows that LMC with a noisy gradient can draw samples from the exact posterior distribution when the step size approaches zero.

Langevin dynamics is described as the hybrid Monte Carlo (HMC) method with one leapfrog step in section 5.5.2 of Neal (2010):

$$\begin{aligned} \mathbf{p}_{t+\varepsilon/2} &= \mathbf{p}_t + \frac{\varepsilon C}{2} \mathbf{g}(\boldsymbol{\theta}_t) \\ \boldsymbol{\theta}_{t+\varepsilon} &= \boldsymbol{\theta}_t + \varepsilon C \mathbf{p}_{t+\varepsilon/2} \\ \mathbf{p}_{t+\varepsilon} &= \mathbf{p}_{t+\varepsilon/2} + \frac{\varepsilon C}{2} \mathbf{g}(\boldsymbol{\theta}_{t+\varepsilon}) \end{aligned} \quad (4)$$

where \mathbf{p} is the auxiliary momentum, ε is the step size, C is a positive definite preconditioning matrix, and \mathbf{g} is the gradient of the log-posterior probability $\log P(\boldsymbol{\theta}|\mathcal{D})$ ¹. A new value of \mathbf{p} is drawn at every iteration from an isotropic Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbb{I})$ and then discarded after $\boldsymbol{\theta}$ is updated. The leapfrog step is usually followed by a Metropolis-Hastings accept/reject step to ensure detailed balance. But since the rejection rate decays as ε^3 , that step can be skipped for small step sizes without incurring much error.

In a MRF, the gradient term \mathbf{g} involves computing an expectation over exponentially many states as

$$g_\alpha(\theta) = \sum_{m=1}^N f_\alpha(\mathbf{x}_\alpha^{(m)}) - \mathbb{N} \mathbb{E}_{P(\mathbf{x}|\boldsymbol{\theta})} f_\alpha(\mathbf{x}_\alpha) - \frac{\theta_\alpha}{\sigma_0^2} \quad (5)$$

The expectation is estimated by a set of state samples $\{\tilde{\mathbf{x}}^{(s)}\}$ in the “brief Langevin” algorithm of Murray and Ghahramani (2004), where these samples are drawn by running a few steps of Gibbs sampling initialized from a subset of the training data. We adopt the “brief Langevin” algorithm with three modifications for faster mixing.

3.1 Persistent Gibbs Sampling

We maintain a set of persistent Markov chains for the state samples $\{\tilde{\mathbf{x}}^{(s)}\}$ by initializing Gibbs sampling at the last states of the previous iteration instead of the data. This is motivated by the persistent contrastive divergence algorithm of Tieleman (2008). When $\boldsymbol{\theta}$ changes slowly enough, the Gibbs sampler will approximately sample from the stationary distribution, even when allowed a few steps at every iteration.

¹We omit all the other random variables that $\boldsymbol{\theta}$ is conditioned on in this section for ease of notation.

3.2 Preconditioning

When the posterior distribution of $\{\theta_\alpha\}$ has different scales along different variables, the original LMC with a common step size for all θ_α 's will traverse the parameter space slowly. We adopt a preconditioning matrix C to speed up the mixing, where C satisfies $CC^T = H$ with H is the Hessian matrix of $\log P(\theta_{\text{MAP}}|\mathcal{D})$, computed as:

$$H(\theta_{\text{MAP}}) = \text{Cov}_{P(\mathbf{x}|\theta_{\text{MAP}})} \mathbf{f}(\mathbf{x}) + \sigma_0^{-2} \quad (6)$$

This is reminiscent of the observed Fisher information matrix in Girolami and Calderhead (2011) except that we use the MAP estimate with the prior. We approximate $H(\theta_{\text{MAP}})$ by averaging over $H(\theta_t)$ during a burn-in period and estimate $\text{Cov}_{P(\mathbf{x}|\theta_t)} \mathbf{f}$ with the set of state samples from the persistent Markov chains. The adoption of a preconditioning matrix also helps us pick a common step size parameter ε suitable for different training sets.

3.3 Partial Momentum Refreshment

The momentum term \mathbf{p} in the leapfrog step represents the update direction of the parameter. Langevin dynamics is known to explore the parameter space through inefficient random walk behavior because it draws an independent sample for \mathbf{p} at every iteration. We can achieve a better mixing rate with the partial momentum refreshment method proposed in Horowitz (1991). When \mathbf{p} is updated at every step by:

$$\mathbf{p}_t \leftarrow \alpha \mathbf{p}_t + \beta \mathbf{n}_t \quad (7)$$

where $n_t \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$, and α, β satisfy $\alpha^2 + \beta^2 = 1$, the momentum is partially preserved from the previous iteration and thereby suppresses the random-walk behavior in a similar fashion as HMC with multiple leapfrog steps.

α controls how much momentum to be carried over from the previous iteration. With a large value of α , LMC reduces the auto-correlation between samples significantly relative to LMC without partial momentum refreshment. The improved mixing rate is illustrated in Figure 2. We also show that the mean and standard deviation of the posterior distribution does not change. However, caution should be exercised especially when the step size η is large because a value of α that is too large would increase the error in the update equation which we do not correct with a Metropolis-Hastings step because that is intractable.

4 Sampling Edges by Reversible Jump MCMC

Langevin dynamics handles the continuous change in the parameter value A_α . As for discrete changes in the model structure, Y_α , we propose an approximate reversible jump MCMC (RJMCMC) step (Green, 1995) to




Figure 2: Comparison of Langevin dynamics on the block model in section 6.1 with partial momentum refreshment $\alpha = 0.9$ against $\alpha = 0$. Step size $\eta = 10^{-3}$. Top: the auto-correlation of two typical parameters. Bottom: the posterior mean and standard deviation of all parameters estimated with 10K samples.

sample Y_α and A_α jointly from the conditional distribution $P(\mathbf{A}, \mathbf{Y}|p_0, \sigma_0, \mathcal{D})$. The proposed Markov chain adds/deletes one clique (or simply one edge when $\alpha = (i, j)$) at a time. When an edge does not exist, i.e., $Y_\alpha = 0$, the variable A_α can be excluded from the model, and therefore we consider the jump between a full model with $Y_\alpha \neq 0$, $A_\alpha = a$ and a degenerate model with $Y_\alpha = 0$.

The proposed RJMCMC is as follows: when $Y_\alpha = 0$, propose adding an edge with probability P_{add} and sample $A_\alpha = a$ from a proposal distribution $q(A)$ with support on $[-\Delta_\alpha, \Delta_\alpha]$; when $Y_\alpha = 1$ and $|A_\alpha| \leq \Delta_\alpha$, propose deleting an edge with P_{del} . The reason of restricting the proposed move within $[-\Delta_\alpha, \Delta_\alpha]$ will be explained later. It is easy to see that the Jacobian is 1. The jump is then accepted by the Metropolis-Hastings algorithm with a probability:

$$Q_{add} = \min\{1, Q^*(a)\}, \quad Q_{del} = \min\{1, 1/Q^*(a)\}$$

$$Q^*(a) = \exp\left(a \sum_m f_\alpha(\mathbf{x}_\alpha^{(m)})\right) \left(\frac{Z(Y_\alpha = 0)}{Z(Y_\alpha = 1, A_\alpha = a)} \right)^N$$

$$\frac{p_0 \mathcal{N}(A_\alpha = a | 0, \sigma_0^2)}{(1 - p_0)} \frac{P_{del}}{P_{add} q(A_\alpha = a)} \quad (8)$$

The factors in the first line of Q^* represent the ratio of the model likelihoods while the other two are respectively the ratio of the prior distributions and the ratio of the proposal distributions.

4.1 Unbiased Estimate to Q^* and $1/Q^*$

Computing the partition functions in equation 8 is generally intractable. However, noticing that $Z(Y_\alpha = 1, A_\alpha = a) \rightarrow Z(Y_\alpha = 0)$, as $a \rightarrow 0$, the log-ratio of the two partition functions should be well approximated by a quadratic approximation at the origin when a is small. In this way

we reduce the problem of estimating the partition function to computing the first and second order derivatives of the log-partition function. We employ a second order Taylor expansion for the ratio of partition functions in Q^* as follows:

$$\begin{aligned} & \left(\frac{Z(Y_\alpha = 0)}{Z(Y_\alpha = 1, A_\alpha = a)} \right)^N \stackrel{\text{def}}{=} R_{add} \\ & \approx \exp \left(-Na \frac{\partial \log(Z)}{\partial A_\alpha} \Big|_{A_\alpha=0} - \frac{Na^2}{2} \frac{\partial^2 \log(Z)}{\partial A_\alpha^2} \Big|_{A_\alpha=0} \right) \end{aligned} \quad (9)$$

We know that the k th order derivatives of the log-partition function of an MRF correspond to the k th order centralized moments (or cumulants) of the features, that is,

$$\frac{\partial \log(Z)}{\partial A_\alpha} = \mathbb{E}f_\alpha, \quad \frac{\partial^2 \log(Z)}{\partial A_\alpha^2} = \text{Var}f_\alpha \quad (10)$$

Given a set of n state samples $\tilde{\mathbf{x}}^{(s)} \sim P(\mathbf{x}|Y_\alpha = 0)$ from the persistent Markov chains, we can compute an unbiased estimate of the mean and variance of f_α as the sample mean \bar{f}_α and sample variance $S_\alpha^2 = \sum_s (f_\alpha(\tilde{\mathbf{x}}^{(s)}) - \bar{f}_\alpha)^2 / (n-1)$ respectively. Consequently we obtain an estimate of R_{add} by plugging \bar{f}_α and S_α^2 into equation 9, which is unbiased in the logarithmic domain, denoted as \tilde{R}_{add} .

An unbiased estimate in the logarithmic domain, unfortunately, is no longer unbiased once transformed to the linear domain. To correct the bias induced by the transformation, we take another Taylor expansion of \tilde{R}_{add}/R_{add} around $a = 0$. After some derivation, we obtain an unbiased estimate of R_{add} upto the second order of a given by

$$\tilde{R}_{add}(a) = \exp \left(-Na\bar{f}_\alpha - \frac{Na^2}{2} S_\alpha^2 \right) \left(1 + \frac{N^2 a^2}{2n} S_\alpha^2 \right)^{-1}$$

with variance:

$$\text{Var}(\tilde{R}_{add}/R_{add}) = \frac{N^2 a^2}{n} \text{Var}(f_\alpha) + o(a^3) \quad (11)$$

Similarly, we can also obtain an unbiased estimate, \tilde{R}_{del} , in $1/Q^*$ when considering deleting an edge, with the same formula as \tilde{R}_{add} except that a is replaced by $-a$ and the sample mean and variance are now estimated with respect to $P(\mathbf{x}|Y_\alpha = 1, A_\alpha = a)$. If we plug in \tilde{R}_{add} (or \tilde{R}_{del}) into Q_{add} (or Q_{del}) we get an unbiased estimate of the acceptance probability except when Q^* (or $1/Q^*$) is close to 1 in which case the min operation causes extra bias. Since the variance can be computed as a function of a in equation 11, we can estimate how large a jump range Δ can be used in order to keep the error in the acceptance probability negligible. A larger data set requires a smaller jump range or alternatively a larger sample set that grows quadratically with the size of the data set.

4.2 Optimal Proposal Distribution q

After plugging equations 9, 10 and 5 into equation 8, we obtain the acceptance probability as a function of a :

$$\begin{aligned} Q_{add} &= \min \left\{ 1, \frac{h(a)}{q(A_\alpha = a)} \text{const} \right\} \\ h(a) &= \exp \left(-\frac{a^2}{2} \left(\frac{1}{\sigma_0^2} + \text{Var}f_\alpha \right) + aNg_\alpha \right) \end{aligned} \quad (12)$$

where g_α and $\text{Var}f_\alpha$ are defined at $\theta_\alpha = 0$. Clearly, the optimal proposal distribution in terms of minimal variance is given by the following truncated Gaussian distribution

$$q_{opt}(A_\alpha = a|\theta) \propto h(a), \quad a \in [-\Delta_\alpha, \Delta_\alpha] \quad (13)$$

When adding an edge, we have state samples $\tilde{\mathbf{x}}^{(s)} \sim P(\mathbf{x}|Y_\alpha = 0)$. The expectation $\mathbb{E}f_\alpha$ in g_α can be estimated by $f_\alpha(\{\tilde{\mathbf{x}}^{(s)}\})$ and $\text{Var}f_\alpha$ by $S_\alpha^2(\{\tilde{\mathbf{x}}^{(s)}\})$. When deleting an edge, the samples are from $P(\mathbf{x}|Y_\alpha = 1, A_\alpha = a)$. We apply a quadratic approximation for $\log Z(\theta_\alpha)$, use equation 10, and estimate $\text{Var}f_\alpha \approx S_\alpha^2$ and $\mathbb{E}f_\alpha \approx \bar{f}_\alpha - aS_\alpha^2$.

4.3 Parallel Proposals

Since the most computationally demanding step is to obtain a set of state samples $\{\tilde{\mathbf{x}}^{(s)}\}$, we want to reuse the samples whenever possible. Given that Δ_α is small enough, the parameter value does not change much when we accept an ‘‘add or delete edge’’ move. We can thus assume that the distribution of A_α on an edge is not affected much by an accepted move on other edges. As a result we do not have to rerun the Gibbs sampler after every edge operation, and we can propose jumps for all $\{\alpha : |A_\alpha| < \Delta_\alpha\}$ in parallel, using the same set of samples. This reduces the computation time significantly.

5 Sampling for Hyper-parameters

Given \mathbf{A} and \mathbf{Y} , the hyper-parameters are easy to sample when using conjugate priors:

$$\begin{aligned} p_0|\mathbf{Y} &\sim \text{Beta}(a + \sum_\alpha I(Y_\alpha = 1), b + \sum_\alpha I(Y_\alpha = 0)) \\ \sigma_0^{-2}|\mathbf{Y}, \mathbf{A} &\sim \Gamma(c + \frac{1}{2} \sum_\alpha I(Y_\alpha = 1), d + \frac{1}{2} \sum_{\alpha: Y_\alpha=1} A_\alpha^2) \end{aligned} \quad (14)$$

where I is the indicator function.

The whole inference algorithm is summarized in Alg 1.

6 Experiments

6.1 Datasets

We assess the performance of our proposed method on two simulated data sets and one real-world data set. For the

Algorithm 1 MCMC for Bayesian MRFs with Langevin Dynamics and RJMCMC

(Parameters: number of iterations $ITER$, number of Gibbs sampling steps N_{Gibbs} , sample size n , step size ε , partial momentum refreshment α , RJMCMC proposal width Δ_α .)

Initialize \mathbf{A} , \mathbf{Y} , and momentum \mathbf{p} randomly

for $iter = 1 \rightarrow ITER$ **do**

 Sample p_0 and σ_0 given \mathbf{A} and \mathbf{Y}

 Run Gibbs sampling for N_{Gibbs} steps to draw $\{\mathbf{x}^{(s)}\}_{s=1}^n$.

 Run LMC to draw $\{A_\alpha : Y_\alpha \neq 0\}$.

 Run RJMCMC to propose adding an edge for $\{\alpha : Y_\alpha = 0\}$, and deleting an edge for $\{\alpha : Y_\alpha = 1, |A_\alpha| < \Delta_\alpha\}$

end for

simulated data, we randomly generate sparse Ising models with binary ± 1 states and with parameters sampled from a Gaussian distribution $\mathcal{N}(0, \sigma^2)$ where $\sigma = 0.5$ for edges and 0.1 for node biases. These models are then converted to their equivalent Boltzmann machines with binary $\{0, 1\}$ states from which we draw exact samples. Two structures are considered: (1) a block model with 12 nodes equally divided into 3 groups. Edges are added randomly within a group with a probability of 0.8, and across groups with 0.1. Edges within a group are strong and positively coupled. There are 66 candidate edges with 20 edges in the ground truth model. (2) a 10×10 lattice with 4950 candidate edges and with 180 edges in the ground truth model.

For the real data, we use the MNIST digits image data. We convert the gray scale pixel values to binary values by thresholding at a value of 50, resize the images to a 14×14 scale, and then pick a 9×12 patch centered in each image where the average value of each pixel is in the range of $[10^{-4}, 1 - 10^{-4}]$. The last step is necessary because the other competing models do not have regularization on their biases, which will result in divergent parameter estimates for pixels that are always 0 or 1.

6.2 Model Specification

We compare our Bayesian structure learning algorithm with two other approaches. One is proposed in Wainwright et al. (2007) which recovers the neighbourhood of nodes by training separate L_1 -regularized logistic regressions on each variable. While its goal is edge detection, we can also use it as a parameter estimation method with two output models, “Wain Max” and “Wain Min”, as defined in Höfling and Tibshirani (2009). We implement the “Wain Max/Min” methods with the Lasso regularized generalized linear model package of Friedman et al. (2010)². The other

approach is one of the several variants of L_1 -regularized MLE methods which use a pseudo-likelihood approximation (Höfling and Tibshirani, 2009), denoted as “MLE”³.

For our Bayesian model, we consider two schemes to specify a model for prediction. One is the fully Bayesian approach, referred as “Bayes”, in which we random pick 100 model samples in the Markov chain and approximate the Bayesian model by a mixture model of these 100 components. The other one is to obtain a single model by applying a threshold of 0.5 to $P(Y_{i,j}|\mathcal{D})$ and estimate the posterior mean of the included edges, referred to as “Bayes PM”.

The performance of the Bayesian model is insensitive to the choice of hyper-parameters. We simply set $a = b = c = d = 5$ for p_0 and σ_0 , and $\sigma_b = 10$ across all experiments. For the parameters of the MCMC method, we also use a common setting. We use a diagonal approximation to the feature covariance $\text{Cov}\mathbf{f}$ and thereby the preconditioning matrix C . We set the sample size $n = 100$, number of Gibbs sampling steps $N_{Gibbs} = 1$, LMC step size $\varepsilon = 10^{-3}$, and momentum refreshment rate $\alpha = 0.9$. The RJMCMC proposal width is set as $\Delta_\alpha = 0.01/\sqrt{N\text{Var}\mathbf{f}}$ to achieve a small estimation error as in equation 11. For each experiment, we run the MCMC algorithm to collect $10K$ samples with a subsampling interval of 1000. Since the exact partition function can be computed on the small block model, we also run an exact MCMC, “Bayes Exact”, with an exact gradient and accept/reject decision as well as larger values in ε , α , and Δ than the approximate MCMC.

Different levels of sparsity have to be considered in L_1 -based methods for an optimal regularization strength. For the Bayesian method, we learn a single sparsity level. However, for the sake of comparison, we also consider a method with p_0 as a parameter and vary it between $(0, 1)$ to induce different sparsity, referred to with a suffix “ p_0 ”.

6.3 Accuracy of Inference

We first evaluate the validity of the proposed MCMC method on the block data where exact inference can be carried out. The marginal posterior distribution of an edge parameter, $\theta_{i,j}$, is a mixture of a point mass at zero and a continuous component with a single mode. Figure 3 shows the histogram of samples in the continuous component of four randomly picked $\theta_{i,j}$ ’s. The title above each plot is the posterior probability of the edge (i, j) being present in the model, i.e. $\theta_{i,j} \neq 0$ or $Y_{i,j} = 1$. In this figure, the marginal distribution estimated from the approximate MCMC method matches the distribution from “Bayes Exact” very well. For a more comprehensive comparison, we run “Bayes p_0 ” and “Bayes Exact p_0 ” methods, and vary the value of p_0 from 10^{-4} to 0.99 to achieve different levels of sparsity. We estimate and collect across different p_0

²Code provided at <http://www-stat.stanford.edu/~tibs/glmnet-matlab>

³Code provided at <http://holgerhoefling.com>




Figure 3: Histogram of the parameter samples of four randomly picked edges from “Bayes” and “Bayes Exact” when $Y_{i,j} = 1$. The training data is from the block model with $N = 100$.




Figure 4: Comparison of “Bayes p_0 ” and “Bayes Exact p_0 ” on posterior probability of $Y_{i,j} = 1$, and the posterior mean and standard deviation of $\theta_{i,j}$ when $Y_{i,j} = 1$. The training data is from the block model with $N = 100$.

values the posterior probability of $\theta_{i,j} \neq 0$, posterior mean and standard deviation of $\theta_{i,j}$ in the continuous component, as shown in Figure 4. Each point represents a parameter under some value of p_0 . We find the approximate MCMC procedure produces about the same values for these three statistics as the exact MCMC method.

6.4 Simulated Data

We then compare the performance of various methods on simulated data sets for two tasks: structure recovery and model estimation. The accuracy of recovering the true structure is measured by the precision and recall of the true edges. The quality of the estimated models is evaluated by the predictive performance on a held-out validation set. Since computing the log-likelihood of a MRF is in general intractable, we use the conditional log-likelihood (CLL) on a group of variables instead, which is a generalization of the conditional marginal log-likelihood in Lee et al. (2006). For each data case in the validation set, we randomly choose a group of variables, which is all the variables in the block model and a 3×3 grid in the lattice and MNIST models, and compute $\log P(\{x_i\}_{i \in \text{group}} | \{x_j\}_{j \notin \text{group}})$.

We train models on the simulated data ranging between 50 and 1000 items. For the Bayesian models, we remove an edge if the posterior probability $P(Y_{i,j} = 1 | \mathcal{D}) < 0.5$.




Figure 5: Precision-Recall Curves for the lattice data with 100 data cases.

Figure 5 shows typical precision-recall curves for different methods. It turns out that all the models with a sparsity tuning parameter perform similarly across all the training sets. The “Bayes” model tends to find a structure with high precision.

We then consider the average CLL as a function of the edge density of a model. The edge density is defined as the percentage of edges present in a model, i.e., $1 - \text{sparsity}$. For the fully Bayesian model, we measure the average density in the Markov chain. In fact, the variance of the edge density is usually small, suggesting that most model samples have about the same number of edges. We show the average CLL of different models trained with 100 data cases in Figure 6. The results on other data sizes are omitted because they have the same tendency. All the Bayesian models give robust prediction performance in the sparse and dense model ranges and “Bayes p_0 ” outperforms all the other methods with a tunable sparsity parameter for almost all settings of p_0 . Moreover, the curve of “Bayes p_0 ” peaks at the true density level. In contrast, L_1 methods would underfit to the data for sparse models or overfit for dense models. “MLE” performs better than “Wain” models. This makes sense as the “Wain” models were not designed for MRF parameter estimation. Figure 7 compares the Bayesian models with exact and approximate inference. Again, the approximate MCMC method generates about the same posterior distribution as the “exact” method.

The difference between Bayesian models and L_1 -based models could be partially explained by the different prior/regularization. As shown in Figure 8, to achieve a sparse structure, we have to use strong regularization in the L_1 models which causes global shrinkage for all parameters, resulting in under-fitting. On the other hand, to obtain a dense structure, weak regularization must be applied globally which leads to over-fitting. In contrast, with a spike and slab prior, the parameter value of existing edges is not affected directly by p_0 . Instead, their variance is controlled by another random variable σ_0 which fits the data automatically with a weak hierarchical prior. The behavior of selective shrinkage in the spike and slab prior is also discussed in Mohamed et al. (2011); Ishwaran and Rao (2005).




Figure 6: Mean and standard deviation of average CLL at different density levels for block (left) and lattice (right) model with 100 training data cases. “Wain Min” is not plotted as it is always inferior to “Wain Max”. Vertical line: true edge density.




Figure 7: Average CLL at different density levels for the block model with 100 training cases.




Figure 8: Absolute value of edge parameters in the true model, “Wain Max”, and “Bayes PM” for lattice data with 100 samples. Parameters are plotted only if they are not zero in at least one model and sorted along the x-axis by the absolute value in the true model. The parameters in “Wain Max” are too small in sparse models (top, edge density $\approx 3\%$) and too large in dense models (bottom, 65%). Inset: zoom-in at the right-hand side. “Wain Min” and “MLE” are similar to “Wain Max”.

The Bayesian models, “Bayes p_0 ” and “Bayes PM p_0 ” do however show an interesting “under-fitting” phenomenon at a large density levels. This is due to a misspecified value for p_0 . Since the standard deviation σ_0 is shared by all the $A_{i,j}$ ’s whose $Y_{i,j} = 1$, when we fix p_0 at an improperly large value, it forces a lot of non-existing edges to be included in the model, which consequently brings down the posterior distribution of σ_0 . This results in too small values on real edges as shown in the inset of Figure 8 and thereby a decrease in the model predictive accuracy.

However, this “misbehavior” in return just suggests the ability of our Bayesian model to learn the true structure. Once we release p_0 through a hierarchical prior, the model will abandon these improper values in p_0 and automatically find a good structure and parameters. The vertical line in Figure 6 indicates the sparsity of the true model. Both “Bayes” and “Bayes PM” find sparsity levels very close to the true value, while L_1 -based methods are under-fitted at that same level as shown in the upper panel of Figure 8. We show the joint performance of edge detection and param-




Figure 9: CLL vs F1 score for block and lattice model with 100 data cases. A good model is at the upper right corner.

ter learning in Figure 9 where the performance of edge detection is summarized by the F1 score (the harmonic mean of the precision and recall). The Bayesian models with a hierarchical p_0 prior achieve both a high F-1 score and CLL value near the upper right corner.

6.5 MNIST Data

Since there does not exist ground truth in the model structure of the MIST data set, we evaluate how well we can learn a sparse model for prediction. Figure 10 shows the average CLL on 10K test images with a model trained on




Figure 10: Mean and standard deviation of average CLL versus edge density on MNIST with 100 and 1000 data cases.

100 and 1000 images respectively. In the sparse and dense model ranges, we observe again a better performance of “Bayes” than L_1 -based methods. “Bayes PM” also shows robustness to under/over-fitting although it seems that simply computing the posterior mean does not provide sufficiently good model parameters in the median density range.

To get a more intuitive comparison about the quality of learned sparse models, we train models on 1000 images by different methods with a density of 0.2 and then run Gibbs sampling to draw 36 samples from each model. The images are shown in Figure 11. While it is hard to get good reconstruction using a model without hidden variables, the Bayesian methods produce qualitatively better images than competing methods, even though “Bayes PM” does not have higher CLL than “MLE” at this level.

A common limitation of learning Bayesian models with the MCMC inference method is that it is much slower than training a model with a point estimate. However, as shown in the experiments, the Bayesian methods are able to learn a good combination of parameters and a structure without the need to tune hyper-parameters through cross validation. Also, the Bayesian methods learn sparser models than L_1 -based methods without sacrificing predictive performance significantly. Because the computational complexity of inference grows exponentially with the maximum clique size of MRFs, L_1 -based models at their optimal (not so



Figure 11: Samples from learned models at an edge density level of 0.2.

sparse) regularization level can in fact become significantly more computationally expensive than their Bayesian counterparts at prediction time. Turning up the regularization will result in sparser models but at the cost of under-fitting the data and thus sacrificing predictive accuracy.

7 Discussion

We propose Bayesian structure learning for MRFs with a spike and slab prior. An approximate MCMC method is proposed to achieve effective inference based on Langevin dynamics and reversible jump MCMC. As far as we know this is the first attempt to learn MRF structures in the fully Bayesian approach using spike and slab priors. Related work was presented in Parise and Welling (2006) with a variational method for Bayesian MRF model selection. However this method can only compare a given list of candidate models instead of searching in the exponentially large structure space.

The proposed MCMC method is shown to provide accurate posterior distributions at small step sizes. The selective shrinkage property of the spike and slab prior enables us to learn an MRF at different sparsity levels without noticeably suffering from under-fitting or over-fitting even for a small data set. Experiments with simulated data and real-world data show that the Bayesian method can learn both an accurate structure and a set of parameter values with strong predictive performance. In contrast the L_1 -based methods could fail to accomplish both tasks with a single choice of the regularization strength. Also, the performance of our Bayesian model is largely insensitive to the choice of hyper-parameters. It provides an automated way to choose a proper sparsity level, while L_1 methods usually rely on cross-validation to find their optimal regularization setting.

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. 0914783, 0928427, 1018433.

References

- G. Andrew and J. Gao. Scalable training of L1-regularized log-linear models. In *Proceedings of the 24th international conference on Machine learning*, pages 33–40. ACM, 2007.
- G. Bresler, E. Mossel, and A. Sly. Reconstruction of Markov random fields from samples: Some observations and algorithms. *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 343–356, 2008.
- A. Courville, J. Bergstra, and Y. Bengio. A spike and slab restricted Boltzmann machine. *Journal of Machine Learning Research, W&CP*, 15, 2011.
- S. Della Pietra, V. Della Pietra, and J. Lafferty. Inducing features of random fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 19(4):380–393, 1997.
- M. Dudik, S. Phillips, and R. Schapire. Performance guarantees for regularized maximum entropy density estimation. *Learning Theory*, pages 472–486, 2004.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010. URL <http://www.jstatsoft.org/v33/i01/>.
- M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- P.J. Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732, 1995.
- D. Hernández-Lobato, J. Hernández-Lobato, T. Helleputte, and P. Dupont. Expectation propagation for Bayesian multi-task feature selection. *Machine Learning and Knowledge Discovery in Databases*, pages 522–537, 2010.
- H. Höfling and R. Tibshirani. Estimation of sparse binary pairwise Markov networks using pseudo-likelihoods. *The Journal of Machine Learning Research*, 10:883–906, 2009.
- A.M. Horowitz. A generalized guided Monte Carlo algorithms. *Physics Letters B*, 268(2):247–252, 1991.
- J. Hu, A. Joshi, and V.E. Johnson. Log-linear models for gene association. *Journal of the American Statistical Association*, 104(486):597–607, 2009.
- H. Ishwaran and J.S. Rao. Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics*, 33(2):730–773, 2005.
- B. Jones, C. Carvalho, A. Dobra, C. Hans, C. Carter, and M. West. Experiments in stochastic computation for high-dimensional graphical models. *Statistical Science*, 20(4):388–400, 2005.
- S.I. Lee, V. Ganapathi, and D. Koller. Efficient structure learning of Markov networks using L1 regularization. In *In NIPS*. Citeseer, 2006.
- S.Z. Li. *Markov random field modeling in image analysis*. Springer-Verlag New York Inc, 2009.
- Y. Lin and D.D. Lee. Bayesian L1-norm sparse learning. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, volume 5, pages V–V. IEEE, 2006.
- T.J. Mitchell and J.J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, pages 1023–1032, 1988.
- S. Mohamed, K. Heller, and Z. Ghahramani. Bayesian and L1 approaches to sparse unsupervised learning. *Arxiv preprint arXiv:1106.1157*, 2011.
- I. Murray and Z. Ghahramani. Bayesian learning in undirected graphical models: approximate MCMC algorithms. In *Proceedings of the 14th Annual Conference on Uncertainty in AI*, pages 392–399, 2004.
- I. Murray, Z. Ghahramani, and D.J.C. MacKay. MCMC for doubly-intractable distributions. In *Proceedings of the 14th Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, Pittsburgh, PA, 2006.
- R.M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo: Methods and Applications*, page 113, 2010.
- S. Parise and M. Welling. Structure learning in Markov random fields. *Advances in Neural Information Processing Systems*, 2006.
- T. Park and G. Casella. The Bayesian lasso. *Journal of the American Statistical Association*, 103(482):681–686, 2008.
- Y. Qi, M. Szummer, and T.P. Minka. Bayesian conditional random fields. In *Artificial Intelligence and Statistics*, 2005.
- P. Ravikumar, M.J. Wainwright, and J.D. Lafferty. High-dimensional ising model selection using L1-regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- S. Riezler and A. Vasserman. Incremental feature selection and L1 regularization for relaxed maximum-entropy modeling. In *Proceedings of EMNLP*, volume 4, 2004.
- G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph models for social networks. *Social networks*, 29(2):173–191, 2007.
- F. Sha and F. Pereira. Shallow parsing with conditional random fields. In *Proceedings of the 2003 Conference of*

the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1, pages 134–141. Association for Computational Linguistics, 2003.

- T. Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradients. In *Proceedings of the International Conference on Machine Learning*, volume 25, pages 1064–1071, 2008.
- Y. Tsuruoka, J. Tsujii, and S. Ananiadou. Stochastic gradient descent training for L1-regularized log-linear models with cumulative penalty. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-Volume 1*, pages 477–485. Association for Computational Linguistics, 2009.
- M.J. Wainwright, P. Ravikumar, and J.D. Lafferty. High-dimensional graphical model selection using L1-regularized logistic regression. *Advances in neural information processing systems*, 19:1465, 2007.
- M. Welling and Y.W. Teh. Bayesian learning via stochastic gradient langevin dynamics. 2011.
- J. Zhu, N. Lao, and E.P. Xing. Grafting-light: fast, incremental feature selection and structure learning of Markov random fields. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 303–312. ACM, 2010.