



HAL
open science

Hidden Gibbs random fields model selection using Block Likelihood Information Criterion

Julien Stoehr, Jean-Michel Marin, Pierre Pudlo

► **To cite this version:**

Julien Stoehr, Jean-Michel Marin, Pierre Pudlo. Hidden Gibbs random fields model selection using Block Likelihood Information Criterion. 2024. hal-01330202

HAL Id: hal-01330202

<https://hal.science/hal-01330202>

Preprint submitted on 29 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hidden Gibbs random fields model selection using Block Likelihood Information Criterion

Julien Stoehr¹, Jean-Michel Marin^{*2}, and Pierre Pudlo³

¹School of Mathematical Sciences & Insight Centre for Data Analytics, University College Dublin, Ireland

²Institut Montpellierain Alexander Grothendieck, Université de Montpellier, France

³Institut de Mathématiques de Marseille, Aix-Marseille Université, France

March 13, 2022

Abstract

Performing model selection between Gibbs random fields is a very challenging task. Indeed, due to the Markovian dependence structure, the normalizing constant of the fields cannot be computed using standard analytical or numerical methods. Furthermore, such unobserved fields cannot be integrated out and the likelihood evaluation is a doubly intractable problem. This forms a central issue to pick the model that best fits an observed data. We introduce a new approximate version of the *Bayesian Information Criterion*. We partition the lattice into continuous rectangular blocks and we approximate the probability measure of the hidden Gibbs field by the product of some Gibbs distributions over the blocks. On that basis, we estimate the likelihood and derive the *Block Likelihood Information Criterion* (BLIC) that answers model choice questions such as the selection of the dependency structure or the number of latent states. We study the performances of BLIC for those questions. In addition, we present a comparison with ABC algorithms to point out that the novel criterion offers a better trade-off between time efficiency and reliable results.

Keywords: Hidden Markov random fields; model selection; Bayesian Information Criterion

1 Introduction

Gibbs or discrete Markov random fields have appeared as convenient statistical model to analyse different types of spatially correlated data. Notable examples are the autologistic model (Besag, 1974) and its extension the Potts model used to describe the spatial dependency of discrete random variables (*e.g.*, shades of grey or colors) on the vertices of an undirected graph (*e.g.*, a regular grid of pixels). In particular, hidden Markov random fields offer an appropriate representation for practical settings where the true state is unknown. The general framework can be described as an observed data \mathbf{y} which is a noisy or incomplete version of an unobserved discrete latent process \mathbf{x} . Shaped by the development of Geman and Geman (1984) and Besag (1986), these models have enjoyed

*jean-michel.marin@umontpellier.fr

great success in image analysis – see for example [Alfö et al. \(2008\)](#) and [Moore et al. \(2014\)](#) who performed image segmentation with the help of this modelling – but also in other applications including disease mapping (*e.g.*, [Green and Richardson, 2002](#)) and genetic analysis (*e.g.*, [François et al., 2006](#), [Friel et al., 2009](#)) to name a few. Despite their popularity, Gibbs random fields suffer from major computational difficulties since their normalizing constant is intractable. This forms a central issue in statistical analysis as the computation of the likelihood is an integral part of the procedure for both parameter inference (*e.g.*, [Celeux et al., 2003](#), [Friel et al., 2009](#), [McGrory et al., 2009](#), [Everitt, 2012](#)) and model selection (*e.g.*, [Grelaud et al., 2009](#), [Friel, 2013](#), [Cucala and Marin, 2013](#), [Stoehr et al., 2015](#)). Remark the exception of small lattices on which we can apply the recursive algorithm of [Reeves and Pettitt \(2004\)](#), [Friel and Rue \(2007\)](#) and obtain an exact computation of the normalizing constant. However, the complexity in time of the above algorithm grows exponentially and is thus helpless on large lattices.

The present paper cares about the problem of selecting the number of latent states as well as the dependency structure of hidden Potts model and explores the opportunity of using the Bayesian Information Criterion (BIC, [Schwarz, 1978](#)) to answer the question. If the problem of recovering the number of hidden states is common in image segmentation, the problem of selecting a dependency structure has received little attention in the literature. [Stoehr et al. \(2015\)](#) have proposed to use approximate Bayesian computation (ABC) model choice (*e.g.*, [Marin et al., 2012](#)) based on geometric summary statistics to tackle the choice of an underlying graph but their approach is restricted to the latter. While our work is motivated by a more general issue, it offers a way to overcome the computational burden of ABC algorithms.

Model choice is a problem of probabilistic model comparison. The standard approach to compare one model against another is based on the Bayes factor ([Kass and Raftery, 1995](#)) that involves the ratio of the evidence of each model. However the evidence can usually not be computed with standard procedure due to a high-dimensional integral. Various approximations have been proposed but a commonly used one, if only for its simplicity, is BIC that is an asymptotic estimate of the evidence based on the Laplace method for integrals. The criterion is a simple penalized function of the maximized log-likelihood which, in the context of hidden Gibbs random fields, cannot be computed since it requires to integrate the intractable Gibbs distribution over the latent space configurations. As regards the simpler case of observed Markov random field solutions have been brought by penalized pseudolikelihood ([Ji and Seymour, 1996](#)) or MCMC approximation of BIC ([Seymour and Ji, 1996](#)). To circumvent the computational difficulties in the hidden case, little has been done before the work of [Stanford and Raftery \(2002\)](#) and [Forbes and Peyrard \(2003\)](#). Both propose approximations that consist in replacing the intractable likelihood with a product distribution on system of independent variables to make the computation tractable. Our main contribution is to show that larger collections of variables, namely blocks of the lattice, can be considered by taking advantage of the exact recursion of [Reeves and Pettitt \(2004\)](#) and leads to an efficient criterion : the Block Likelihood Information Criterion (BLIC). In particular, we will show that a reasonable approximation of the Gibbs distribution is a product of Gibbs distributions on each independent block. Such ideas have occurred in the context of composite likelihood but the use of non-genuine probability distribution results in misspecified model (*e.g.*, [Okabayashi et al., 2011](#), [Friel, 2012](#), [Stoehr and Friel, 2015](#)) that we have decided to avoid.

The paper is organized as follows: Section 2 presents hidden Gibbs random fields. In section 3, after recalling the basis of BIC, we introduced our Block Likelihood Information Criterion (BLIC). In Section 4, to assess the performances of the novel criterion, it is compared to pre-existing

criteria on simulated data sets. We fill in our study with a comparison between BLIC and the ABC algorithm of [Stoehr et al. \(2015\)](#).

2 Hidden Gibbs random fields

A discrete random field \mathbf{X} is a collection of random variables X_i indexed by a finite set $\mathcal{S} = \{1, \dots, n\}$, whose elements are called sites, and taking values in a finite state space $\mathcal{X} := \{0, \dots, K-1\}$, interpreted as colors. For a given subset $A \subset \mathcal{S}$, \mathbf{X}_A and \mathbf{x}_A respectively define the random process on A , *i.e.*, $\{X_i, i \in A\}$, and a realisation of \mathbf{X}_A . Denotes $\mathcal{S} \setminus A = -A$ the complement of A in \mathcal{S} . When modeling a digital image, the sites are lying on a regular 2D-grid of pixels, and their dependency is given by an undirected graph \mathcal{G} which induces a topology on \mathcal{S} : by definition, sites i and j are adjacent or neighbor if and only if i and j are linked by an edge in \mathcal{G} . A random field \mathbf{X} is a Markov random field with respect to \mathcal{G} , if for all configuration \mathbf{x} and for all sites i

$$\mathbf{P}(X_i = x_i \mid \mathbf{X}_{-i} = \mathbf{x}_{-i}) = \mathbf{P}(X_i = x_i \mid \mathbf{X}_{\mathcal{N}(i)} = \mathbf{x}_{\mathcal{N}(i)}), \quad (1)$$

where $\mathcal{N}(i)$ denotes the set of all the adjacent sites to i in \mathcal{G} . The Hammersley-Clifford theorem states that if the distribution of a Markov random field with respect to a graph \mathcal{G} is positive for all configuration \mathbf{x} then it admits a Gibbs representation for the same topology (see for example [Grimmett \(1973\)](#), [Besag \(1974\)](#) and for a historical perspective [Clifford \(1990\)](#)), namely a probability measure π on \mathcal{X}^n given by

$$\pi(\mathbf{x} \mid \psi, \mathcal{G}) = \frac{1}{Z(\psi, \mathcal{G})} \exp\{-H(\mathbf{x} \mid \psi, \mathcal{G})\}, \quad (2)$$

where $\psi = (\psi_1, \dots, \psi_d)$ is a vector of parameters, H denotes the energy function or Hamiltonian. The present paper solely focuses on models whose Hamiltonian linearly depends on the parameter ψ , that is

$$H(\mathbf{x} \mid \psi, \mathcal{G}) = -\psi^T \mathbf{S}(\mathbf{x}).$$

where $\mathbf{S}(\mathbf{x}) = (s_1(\mathbf{x}), \dots, s_d(\mathbf{x}))$ is a vector of sufficient statistics. The inherent difficulty of all these models that arises from the intractable normalizing constant, called the partition function, defined by

$$Z(\psi, \mathcal{G}) = \sum_{\mathbf{x} \in \mathcal{X}^n} \exp\{\psi^T \mathbf{S}(\mathbf{x})\}$$

The latter is a summation over the numerous possible realizations of the random field \mathbf{X} , that cannot be computed directly (except for small grids and small number of colors K).

In hidden Markov random fields, the latent process is observed indirectly through another field; this permits the modelling of noise that may happen upon many concrete situations. The aim is to infer some properties of a latent state \mathbf{x} given an observation \mathbf{y} . Precisely, given the realization \mathbf{x} of the latent, the observation \mathbf{y} is a family of random variables indexed by the set of sites \mathcal{S} , and taking values in a set \mathcal{Y} , *i.e.*, $\mathbf{y} = \{y_i; i \in \mathcal{S}\}$, and are commonly assumed as independent draws that form a noisy version of the hidden field. Consequently, we set the conditional distribution of \mathbf{Y} knowing $\mathbf{X} = \mathbf{x}$, also called emission distribution, as the product

$$\pi(\mathbf{y} \mid \mathbf{x}, \phi) = \prod_{i \in \mathcal{S}} \pi(y_i \mid x_i, \phi),$$

where $\pi(y_i | x_i, \phi)$ is the marginal noise distribution parametrized by ϕ , that is given for any site i . Those marginal distributions are for instance discrete distributions (Everitt, 2012), Gaussian (e.g., Besag et al., 1991, Qian and Titterton, 1991, Forbes and Peyrard, 2003, Cucala and Marin, 2013) or Poisson distributions (e.g., Besag et al., 1991). Model of noise that takes into account information of the nearest neighbors have also been explored (Besag, 1986). Hence the likelihood of the hidden Gibbs random field with parameter ψ on the graph \mathcal{G} and emission distribution $\pi(\cdot | \mathbf{x}, \phi)$ is given by

$$\pi(\mathbf{y} | \phi, \psi) = \sum_{\mathcal{X}^n} \pi(\mathbf{y} | \mathbf{x}, \phi) \pi(\mathbf{x} | \psi, \mathcal{G}). \quad (3)$$

The latter faces a double intractable issue as neither the likelihood of the latent field, nor the above sum can be computed directly: the cardinality of the range of the sum is of combinatorial complexity.

3 Block Likelihood Information Criterion

The Bayesian Information Criterion offers a mean arising from Bayesian viewpoint to select a statistical model. In what follows, we provide solely the foundation that motivates our contribution and we refer the reader for instance to Raftery (1995) for a more detailed presentation.

3.1 Background on Bayesian Information Criterion

We are given n independent and identically distributed observations $\mathbf{y} = \{y_1, \dots, y_n\}$ from an unknown statistical model to estimate. The Bayesian approach to model selection is based on posterior model probabilities. Consider a finite set of models $\{m : 1, \dots, M\}$ where each one is defined by a probability density function π_m related to a parameter space Θ_m . The model that best fits an observation \mathbf{y} is the model with the highest posterior probability

$$\pi(m | \mathbf{y}) = \frac{\pi(m)e(\mathbf{y} | m)}{\sum_{m'} \pi(m')e(\mathbf{y} | m')},$$

where $e(\mathbf{y} | m)$ denotes the evidence of m , that is the joint distribution of (\mathbf{y}, θ_m) integrated over space parameter Θ_m

$$e(\mathbf{y} | m) = \int \pi_m(\mathbf{y} | \theta_m) \pi_m(\theta_m) d\theta_m.$$

Under the assumption of model being equally likely *a priori*, it is equivalent to choose the model with the largest evidence. From the Laplace method for integrals, under regularity conditions, the evidence of model m can be written as

$$\log e(\mathbf{y} | m) = \log \pi_m(\mathbf{y} | \hat{\theta}_{\text{MLE}}) - d_m \log(n) + R_m(\hat{\theta}_{\text{MLE}}) + \mathcal{O}(n^{-\frac{1}{2}}), \quad (4)$$

where where $\hat{\theta}_{\text{MLE}}$ is the maximum likelihood estimator of π_m , d_m is the number of free parameters for model m and R_m is bounded as the sample size grows to infinity (e.g., Schwarz, 1978, Tierney and Kadane, 1986).

BIC is an asymptotical estimate of the evidence defined by

$$-2 \log e(\mathbf{y} | m) \simeq \text{BIC}(m) = -2 \log \pi_m(\mathbf{y} | \hat{\theta}_{\text{MLE}}) + d_m \log(n). \quad (5)$$

The $d_m \log(n)$ term corresponds to a penalty term which increases with the complexity of the model. Thus selecting the model with the largest evidence is equivalent to choose the model which minimizes BIC. Regardless of the prior on parameter, the error in (5) is, in general, solely bounded and does not go to zero even with an infinite amount of data. The approximation may hence seem somewhat crude. However as observed by Kass and Raftery (1995) the criterion does not appear to be qualitatively misleading as long as the sample size n is much larger than the number d_m of free parameters in the model. In addition, a reasonable choice of the prior can lead to much smaller error. Indeed, Kass and Wasserman (1995) have found that the error is $\mathcal{O}(n^{-1/2})$ for a well chosen multivariate normal prior distribution.

BIC can be defined beside the special case of independent random variables. In the latter case the number of free parameter is, in general, not equal to the dimension of the parameter space as for the independent case. The consistency of BIC has been proven in various situations such as independent and identically distributed processes from the exponential families (Haughton, 1988), mixture models (Keribin, 2000), Markov chains (Csiszár et al., 2000, Gassiat, 2002). When dealing with observed Markov random fields, aside from the problem of intractable likelihoods the number of free parameters in the penalty term has no simple formula. In the context of selecting a neighborhood system, Csiszár and Talata (2006) proposed to replace the likelihood by the pseudolikelihood (Besag, 1975) and modify the penalty term as the number of all possible configurations for the neighboring sites. The resulting criterion is shown to be consistent as regards this model choice. Up to our knowledge such a result has not been yet derived for hidden Markov random field. The problem of approximating BIC could be termed a triple intractable problem since neither the maximum likelihood estimate $\hat{\theta}_{\text{MLE}}$ nor the incomplete likelihood $\pi_m(\cdot | \theta)$ can be computed with standard methods since they require to integrate over the latent configuration space and no simple definition of d_m is available.

3.2 Gibbs distribution approximations

A convenient way to circumvent the issues of computing BIC is to replace the Gibbs distribution by tractable surrogates since it avoids the use of time consuming simulations methods. As for the pseudolikelihood (Besag, 1975) and more generally composite likelihood (Lindsay, 1988), the main idea consists in replacing the original Markov distribution by a product of easily normalized distribution. But while composite likelihoods are not a genuine probability distribution for Gibbs random field, the focus hereafter is solely on valid probability function by considering system of independent variables. This choice is motivated by the observations that at finite sample size, when dealing with composite likelihood, misspecification of the model has to be taken into account (*e.g.*, Friel, 2012, Stoehr and Friel, 2015), so that constant terms may appear in the remainder R_m in (4).

Finding good approximations of the Gibbs distribution has long standing antecedents in statistical mechanics when one aims at predicting the response to the system to a change in the Hamiltonian. One important technique is based on a variational approach as the minimizer of the free energy, sometimes referred to as variational or Gibbs free energy and defined with the Kullback-Leibler divergence between \mathbf{P} and the target distribution $\pi(\cdot | \psi, \mathcal{G})$ as

$$F(\mathbf{P}) = -\log Z(\psi, \mathcal{G}) + \text{KL}(\mathbf{P}, \pi(\cdot | \psi, \mathcal{G})). \quad (6)$$

The Kullback-Leibler divergence being non-negative and zero if and only if $\mathbf{P} = \pi(\cdot | \psi, \mathcal{G})$, the free energy has an optimal lower bound achieved for $\mathbf{P} = \pi(\cdot | \psi, \mathcal{G})$. Minimizing the free energy

with respect to the set of probability distribution on \mathcal{X}^n allows to recover the Gibbs distribution but presents the same computational intractability. A solution is to minimize the Kullback-Leibler divergence over a restricted class of tractable probability distribution on \mathcal{X}^n . This is the basis of mean field approaches that aim at minimizing the Kullback-Leibler divergence over the set of probability functions that factorize on sites of the lattice. The minimization of (6) over this set leads to fixed point equations for each marginal of \mathbf{P} (see for example [Jordan et al., 1999](#)). The resulting solution motivates the mean field-like approximations of [Celeux et al. \(2003\)](#) for which the neighbors of a site i are set to well chosen constant independently of the value at the given site, namely

$$\mathbf{P}^{\text{MF-like}}(\mathbf{x} \mid \psi, \mathcal{G}) = \prod_{i \in \mathcal{S}} \pi(x_i \mid \mathbf{X}_{\mathcal{N}(i)} = \tilde{\mathbf{x}}_{\mathcal{N}(i)}, \psi, \mathcal{G}). \quad (7)$$

Instead of considering distributions that completely factorize on single sites, we are hereafter interested in tractable approximations that factorize over larger sets of nodes, namely blocks of the lattice. Consider a partition of \mathcal{S} into contiguous rectangular blocks, namely

$$\mathcal{S} = \bigsqcup_{\ell=1}^C A(\ell),$$

and denote \tilde{D} the class of independent probability distributions \mathbf{P} that factorize with respect to this partition, that is if $\mathcal{X}_{A(\ell)}^n$ stands for the configuration space of the block $A(\ell)$, for all \mathbf{x} in \mathcal{X}^n

$$\mathbf{P}(\mathbf{x}) = \prod_{\ell=1}^C \mathbf{P}_\ell(x_{A(\ell)}), \text{ where } \mathbf{P}_\ell \in \mathcal{M}_1^+(\mathcal{X}_{A(\ell)}^n) \text{ and } \mathbf{P} \in \mathcal{M}_1^+(\mathcal{X}^n).$$

To take over from the Gibbs likelihood, we propose to explore the opportunity of probability distributions in \tilde{D} of the form

$$\mathbf{P}(\mathbf{x} \mid \tilde{\mathbf{x}}, A(1), \dots, A(C), \psi) = \prod_{\ell=1}^C \pi(\mathbf{x}_{A(\ell)} \mid \mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}, \psi, \mathcal{G}), \quad (8)$$

where $\tilde{\mathbf{x}}$ is a constant field in \mathcal{X}^n to specify and $B(\ell)$ is either the border of $A(\ell)$, *i.e.*, elements of the absolute complement of $A(\ell)$ that are connected to elements of $A(\ell)$ in \mathcal{G} , or the empty set. In the latter case, we are cancelling the edges in \mathcal{G} that link elements of $A(\ell)$ to elements of any other subset of \mathcal{S} such that the factorization is independent of $\tilde{\mathbf{x}}$. The Gibbs distribution is then simply replaced by the product of the likelihood restricted to $A(\ell)$. For instance a Potts model on \mathcal{X}^n is replaced with a product of Potts models on $\mathcal{X}_{A(\ell)}^n$. To underline that point, $\tilde{\mathbf{x}}$ is omitted in what follows when $B(\ell) = \emptyset$. Note that composite likelihoods differs from (8) in most cases since blocks are not allowed to overlap and contrary to conditional composite likelihoods, neighbors are set to constants. The only example of composite likelihoods that lies in \tilde{D} is marginal composite likelihoods for non overlapping blocks.

The assumption of independent blocks leads to tractable BIC approximations. Indeed, plugging

the probability distribution (8) in place of the Gibbs distribution in (3) yields

$$\begin{aligned}
\mathbf{P}_m(\mathbf{y} \mid \tilde{\mathbf{x}}, \theta) &= \sum_{\mathbf{x} \in \mathcal{X}^n} \pi(\mathbf{y} \mid \mathbf{x}, \phi) \mathbf{P}(\mathbf{x} \mid \tilde{\mathbf{x}}, A(1), \dots, A(C), \psi) \\
&= \prod_{\ell=1}^C \sum_{\mathbf{x}_{A(\ell)}} \left\{ \prod_{i \in A(\ell)} \pi(y_i \mid x_i, \phi) \right\} \pi(\mathbf{x}_{A(\ell)} \mid \mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}, \psi, \mathcal{G}) \\
&= \prod_{\ell=1}^C \sum_{\mathbf{x}_{A(\ell)}} \pi(\mathbf{y}_{A(\ell)} \mid \mathbf{x}_{A(\ell)}, \phi) \pi(\mathbf{x}_{A(\ell)} \mid \mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}, \psi, \mathcal{G}). \tag{9}
\end{aligned}$$

This estimate of the incomplete likelihood $\pi_m(\cdot \mid \theta)$ leads to the following BIC approximations

$$\text{BIC}(m) \approx -2 \log \mathbf{P}_m(\mathbf{y} \mid \tilde{\mathbf{x}}, \theta^*) + d_m \log(|\mathcal{S}|) := \text{BLIC}^{\tilde{\mathbf{x}}}(m \mid \theta^*), \tag{10}$$

where $\theta^* = (\phi^*, \psi^*)$ is a parameter value to specify. We refer to these approximations as Block Likelihood Information Criterion (BLIC). In the first instance, the number of free parameters d_m is set to the dimension of Θ_m , that is we are neglecting the interaction between variables within a block in the penalty term.

Our proposal relies on that each term of the product (9) can be computed using the recursion of [Friel and Rue \(2007\)](#) as long as the blocks are small enough. Indeed for models whose potential linearly depends on the parameter, the probability distribution on $A(\ell)$ can be written as a Gibbs distribution on the block conditioned on the fixed border $\tilde{\mathbf{x}}_{B(\ell)}$, namely

$$\pi(\mathbf{x}_{A(\ell)} \mid \mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}, \psi, \mathcal{G}) = \frac{1}{Z(\psi, \mathcal{G}, \tilde{\mathbf{x}}_{B(\ell)})} \exp\{\psi^T \mathbf{S}(\mathbf{x}_{A(\ell)} \mid \tilde{\mathbf{x}})\},$$

where $\mathbf{S}(\mathbf{x}_{A(\ell)} \mid \tilde{\mathbf{x}})$ is the restriction of \mathbf{S} to the subgraph defined on the set $A(\ell)$ and conditioned on the fixed border $\tilde{\mathbf{x}}_{B(\ell)}$, and $Z(\psi, \mathcal{G}, \tilde{\mathbf{x}}_{B(\ell)})$ is the corresponding normalizing constant. Assuming that all the marginals of the emission distribution are positive, it follows

$$\begin{aligned}
&\sum_{\mathbf{x}_{A(\ell)}} \pi(\mathbf{y}_{A(\ell)} \mid \mathbf{x}_{A(\ell)}, \phi) \pi(\mathbf{x}_{A(\ell)} \mid \mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}, \psi, \mathcal{G}) \\
&= \frac{1}{Z(\psi, \mathcal{G}, \tilde{\mathbf{x}}_{B(\ell)})} \underbrace{\sum_{\mathbf{x}_{A(\ell)}} \exp\{\log \pi(\mathbf{y}_{A(\ell)} \mid \mathbf{x}_{A(\ell)}, \phi) + \psi^T \mathbf{S}(\mathbf{x}_{A(\ell)} \mid \tilde{\mathbf{x}})\}}_{=Z(\theta, \mathcal{G}, \mathbf{y}_{A(\ell)}, \tilde{\mathbf{x}}_{B(\ell)})}.
\end{aligned}$$

The term $Z(\theta, \mathcal{G}, \mathbf{y}_{A(\ell)}, \tilde{\mathbf{x}}_{B(\ell)})$ corresponds to the normalizing constant of the conditional random field $\mathbf{X}_{A(\ell)}$ knowing $\mathbf{Y}_{A(\ell)} = \mathbf{y}_{A(\ell)}$ and $\mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}$, that is the initial model with an extra potential on singletons. Then the algebraic simplification at the core of the algorithm of [Friel and Rue \(2007\)](#) applies for both normalizing constants, such that we can exactly compute the Block Likelihood Information Criterion, namely

$$\text{BLIC}^{\tilde{\mathbf{x}}}(m \mid \theta^*) = -2 \sum_{\ell=1}^C \left\{ \log Z(\theta^*, \mathcal{G}, \mathbf{y}_{A(\ell)}, \tilde{\mathbf{x}}_{B(\ell)}) - \log Z(\psi^*, \mathcal{G}, \tilde{\mathbf{x}}_{B(\ell)}) \right\} + d_m \log(|\mathcal{S}|). \tag{11}$$

3.3 Related model choice criteria

This approach encompasses the Pseudolikelihood Information Criterion (PLIC, [Stanford and Raftery, 2002](#)) as well as the mean field-like approximations $\text{BIC}^{\text{MF-like}}$ proposed by [Forbes and Peyrard \(2003\)](#). When one considers the finest partition of \mathcal{S} , that is distributions that factorize on sites, they have already proposed ingenious solutions for choosing $\tilde{\mathbf{x}}$ and estimating $\hat{\theta}_*$ in (10). Indeed, [Stanford and Raftery \(2002\)](#) suggest to set $(\tilde{\mathbf{x}}, \hat{\theta}_*)$ to the final estimates $(\hat{\theta}^{\text{ICM}}, \tilde{\mathbf{x}}^{\text{ICM}})$ of the unsupervised Iterated Conditional Modes (ICM, [Besag, 1986](#)) algorithm, while [Forbes and Peyrard \(2003\)](#) put forward the use of the output $(\hat{\theta}^{\text{MF-like}}, \tilde{\mathbf{x}}^{\text{MF-like}})$ of the simulated field algorithm of [Celeux et al. \(2003\)](#). To make this statement clear, we could note

$$\begin{aligned} \text{PLIC}(m) &= \text{BLIC}_{\tilde{\mathbf{x}}^{\text{ICM}}} \left(m \mid \hat{\theta}^{\text{ICM}} \right), \\ \text{BIC}^{\text{MF-like}}(m) &= \text{BLIC}_{\tilde{\mathbf{x}}^{\text{MF-like}}} \left(m \mid \hat{\theta}^{\text{MF-like}} \right). \end{aligned}$$

Whilst PLIC shows good result as regards the selection of the number of components of the hidden state, ICM performs poorly for the parameter estimation in comparison with the EM-like algorithm of [Celeux et al. \(2003\)](#). Hence we advocate in favour of the latter in what follows to get estimates of $\hat{\theta}_{\text{MLE}}$ and to fix a segmented random field $\tilde{\mathbf{x}}$.

We shall also remark that for a factorization over the graph nodes when $B(\ell) = \emptyset$ we retrieve a mixture model. Indeed, turning off all the edges in \mathcal{G} leads to approximate the Gibbs distribution by a multinomial distribution with event probabilities depending on the potential on singletons. Hence if marginal emission distribution are Gaussian random variables depending on the component on the latent site associated, we would deal with a classical Gaussian mixture model.

4 Comparison of BIC approximations

Our primary intent with the BIC approximations was to choose the number of latent states as well as the dependency structure of a hidden Markov random fields. The following numerical experiments illustrate the performances as regards these questions for realizations of a hidden Potts model.

4.1 Hidden Potts models

This numerical part of the paper focuses on observations for which the hidden field is modelled by a K -states Potts model. While being widely used in practice (*e.g.*, [Hurn et al., 2003](#), [Alfò et al., 2008](#), [François et al., 2006](#), [Moores et al., 2014](#)), the model is representative of the computational difficulties of hidden Gibbs random field. The model sets a probability distribution on $\mathcal{X}^n = \{1, \dots, K\}^n$ parametrized by a scalar ψ that adjusts the level of dependency between adjacent sites and whose Hamiltonian is given by

$$H(\mathbf{x} \mid \psi, \mathcal{G}) = -\psi \sum_{i \overset{\mathcal{G}}{\sim} j} \mathbf{1}\{x_i = x_j\}.$$

The above sum $i \overset{\mathcal{G}}{\sim} j$ ranges the set of edges of the graph \mathcal{G} . In the statistical physic literature, ψ is interpreted as the inverse of a temperature, and when the temperature drops below a fixed threshold, values x_i of a typical realization of the field are almost all equal (the model then exhibits strong dependency between all sites). These peculiarities of Potts models are called phase transitions.




Figure 1: Neighborhood graphs \mathcal{G} of hidden Potts model. (a) The four closest neighbour graph \mathcal{G}_4 defining model $\text{HPM}(\mathcal{G}_4, \theta, K)$. (b) The eight closest neighbour graph \mathcal{G}_8 defining model $\text{HPM}(\mathcal{G}_8, \theta, K)$.

We set the emission distribution such that the marginal distribution are Gaussian distribution entered at the value of the related nodes, namely

$$y_i \mid x_i = k \sim \mathcal{N}(k, \sigma_k^2) \quad k \in \{0, \dots, K-1\},$$

where σ_k is the standard deviation for sites belonging to class k . Even though the noise model is homoscedastic, we still index the standard deviation by k since we do not use the assumption of a constant variance in the estimation procedure, such that the number of parameters estimated is $d_m = 2 \times k + 1$. The parameter to be estimated with the ICM or simulated field algorithms is then $\theta = (\phi, \psi)$, with $\phi = \{(k, \sigma_k) : k = 0, \dots, K-1\}$. We denote $\text{HPM}(\mathcal{G}, \theta, K)$, the hidden K -states Potts model defined above.

The common point of our examples is to select the hidden Potts model that better fits a given observation \mathbf{y}^{obs} composed of $n = 100 \times 100$ pixels among a collection

$$\mathcal{M} = \{\text{HPM}(\mathcal{G}, \theta, K) : K = K_{\min}, \dots, K_{\max} ; \mathcal{G} \in \{\mathcal{G}_4, \mathcal{G}_8\}\},$$

where K is the number of colors of the corresponding model and \mathcal{G} is one of the two possible neighborhood systems: \mathcal{G}_4 and \mathcal{G}_8 , see Figure 1. For each model $\text{HPM}(\mathcal{G}, \theta, K)$, the estimate $\hat{\theta}_{\text{MLE}}$ and the segmented field $\tilde{\mathbf{x}}$ were computed using SpaCEM³ (see the Documentation on <http://spacem3.gforge.inria.fr>). The software allows the implementation of the unsupervised ICM algorithm as well as the simulated field algorithm and provides computation of PLIC, the mean field-like approximations $\text{BIC}^{\text{MF-like}}$ and BIC^{GBF} . The ICM and the EM-like algorithms were both initialized with a simple K -means procedure. The stopping criterion is then settled to a number of 200 iterations that is enough to ensure the convergence of the procedure.

In what follows, we restrict each $A(\ell)$ to be of the same dimension and in particular square block of dimension $b \times b$. For the sake of clarity the Block Likelihood Criterion is indexed by the dimension of the blocks, namely for a partition of square blocks of size $b \times b$ for which $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}^{\text{MF-like}}$ and $\hat{\theta}_{\text{MLE}} = \hat{\theta}^{\text{MF-like}}$, we note it $\text{BLIC}_{b \times b}^{\text{MF-like}}$. As already mentioned, we then have $\text{BIC}^{\text{MF-like}} = \text{BLIC}_{1 \times 1}^{\text{MF-like}}$. We recall that when $B(\ell) = \emptyset$, $\tilde{\mathbf{x}}$ is omitted in the previous notations, that is for a square blocks partition we note our criterion $\text{BLIC}_{b \times b}$. Then $\text{BLIC}_{1 \times 1}$ is the BIC approximations corresponding to a finite independent mixture model. All criterion were tested on simulated images obtained using the Swendsen-Wang algorithm. We describe below the different experiments settings we have considered and the results we got.

4.2 First experiment: selection of the number of colors

In this experiment the dependency structure is assumed to be known and the aim is to recover the number K of colors of the latent configuration. We considered realizations from hidden Potts models with $K_T = 4$ colors and $\sigma_k = 0.5$. The interaction parameter ψ was set close to the phase transition, namely $\psi = 1$ and $\psi = 0.4$ for \mathcal{G}_4 and \mathcal{G}_8 respectively. These values of the parameter ensure the images present homogeneous regions and then the observations exhibit some spatial structure. Such settings illustrate the advantage of taking into account spatial information of the model. Obviously, for values of ψ where the interaction is weaker, the benefit of the criterion that include the dependency structure of the model is not clear. The latter could even be misleading in comparison with BIC approximations for independent mixture models when ψ is close to zero. On the other side, when ψ is above the phase transition, the distribution on \mathcal{X}^n becomes heavily multi-modal and there is almost solely one class represented in the image regardless the number of colors of the model. We carried out 100 simulations from the first order neighborhood structure \mathcal{G}_4 and 100 simulations from the second order neighborhood structure \mathcal{G}_8 .

Table 1: Selected K in the first experiment for 100 realizations from $\text{HPM}(\mathcal{G}_4, \theta, 4)$ and 100 realizations from $\text{HPM}(\mathcal{G}_8, \theta, 4)$ using Pseudolikelihood Information Criterion (PLIC), mean field-like approximations ($\text{BIC}^{\text{MF-like}}$, BIC^{GBF}) and Block Likelihood Information Criterion (BLIC) for various sizes of blocks and border conditions.

HPM($\mathcal{G}_4, \theta, 4$)							HPM($\mathcal{G}_8, \theta, 4$)						
K	2	3	4	5	6	7	K	2	3	4	5	6	7
PLIC	0	9	91	0	0	0	PLIC	0	7	93	0	0	0
$\text{BIC}^{\text{MF-like}}$	0	0	39	23	16	22	$\text{BIC}^{\text{MF-like}}$	0	0	43	18	19	20
BIC^{GBF}	0	0	39	25	18	18	BIC^{GBF}	0	0	52	20	19	9
$\text{BLIC}_{2 \times 2}^{\text{MF-like}}$	0	0	58	18	8	16	$\text{BLIC}_{2 \times 2}^{\text{MF-like}}$	0	0	52	14	17	17
$\text{BLIC}_{1 \times 1}$	0	0	97	1	2	0	$\text{BLIC}_{1 \times 1}$	0	3	90	1	4	2
$\text{BLIC}_{2 \times 2}$	0	0	100	0	0	0	$\text{BLIC}_{2 \times 2}$	0	1	99	0	0	0
							$\text{BLIC}_{4 \times 4}$	0	0	100	0	0	0

The results obtained for the different criterion are reported in Table 1. For $b \geq 2$, $\text{BLIC}_{b \times b}$ outperform the different criterion even though PLIC and $\text{BLIC}_{1 \times 1}$ provide good results. By contrast approximations based on mean field-like approximations, that is $\text{BIC}^{\text{MF-like}}$, BIC^{GBF} and $\text{BLIC}_{2 \times 2}^{\text{MF-like}}$, perform poorly. These conclusions need nonetheless to be put into perspective. Figure 2(a) shows that the main issue encountered by these criterion is their inability to discriminate between the more complex models. Indeed these BIC approximations reach a plateau from $K = 4$, a problem that other criterion do not face. As an example, Figure 2(b) and Figure 2(c) represent boxplots of the difference between BIC values for $\text{HPM}(\mathcal{G}_4, \theta, K)$ as K is increasing for the 100 realizations, namely

$$\Delta(K \rightarrow K + 1) = \text{BIC} \left(\text{HPM}(\mathcal{G}, \hat{\theta}_{\text{MLE}}, K + 1) \right) - \text{BIC} \left(\text{HPM}(\mathcal{G}, \hat{\theta}_{\text{MLE}}, K) \right),$$

for $K = K_{\min}, \dots, K_{\max}$. Hence, BIC approximations grow with K if $\Delta(K \rightarrow K + 1) \geq 0$ and decrease otherwise. It appears that $\text{BLIC}_{2 \times 2}$ increases systematically from $K = 4$ whereas $\text{BIC}^{\text{MF-like}}$

tend to be constant, or even decreases, so that none minimum can be clearly identified. We do not provide the boxplots for $\text{BIC}^{\text{MF-like}}$ and BIC^{GBF} because they are significantly the same.

Finally these results illustrate in particular the importance of a well chosen segmented field $\tilde{\mathbf{x}}$. Indeed PLIC and $\text{BIC}^{\text{MF-like}}$ are both criterion of type $\text{BLIC}_{1 \times 1}^{\tilde{\mathbf{x}}}$ but their performances greatly differ on this example. As regards the selection of K , $\text{BLIC}_{b \times b}$ circumvent this question whilst performing better.

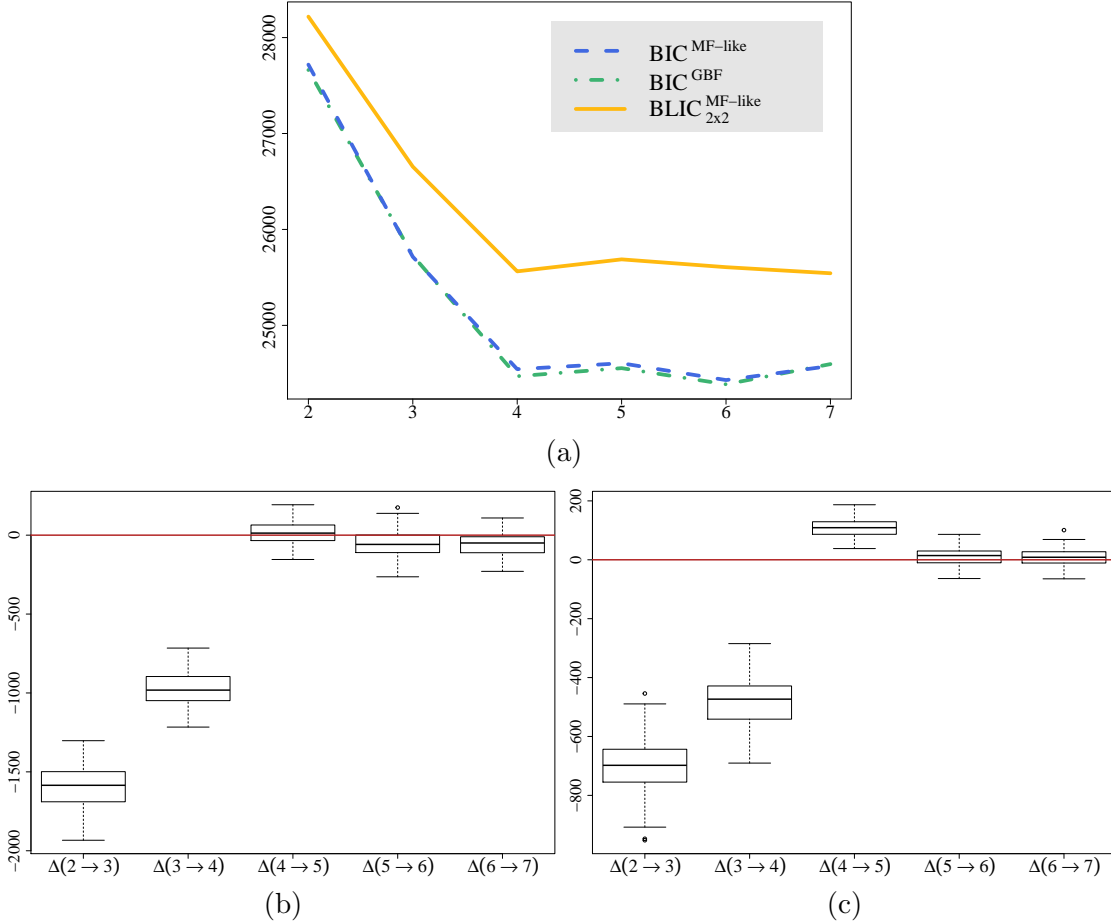


Figure 2: First experiment results. (a) $\text{BIC}^{\text{MF-like}}$, BIC^{GBF} and $\text{BLIC}_{2 \times 2}^{\text{MF-like}}$ values for one realization of a first order hidden Potts model $\text{HPM}(\mathcal{G}, \theta, 4)$. (b) Difference between $\text{BLIC}_{2 \times 2}^{\text{MF-like}}$ values for 100 realization of a first order hidden Potts model $\text{HPM}(\mathcal{G}_4, \theta, 4)$ as K is increasing. (c) Difference between $\text{BLIC}_{2 \times 2}$ values for 100 realization of a first order hidden Potts model $\text{HPM}(\mathcal{G}_4, \theta, 4)$ as K is increasing

4.3 Second experiment: selection of the dependency structure

For this second experiments the setting was exactly the same than for the first experiment. The only difference is that as first instance the number of colors K_T is assumed to be known while the neighborhood system has to be chosen. To answer such a question it is obvious that we can not use criterion $\text{BLIC}_{1 \times 1}$ based on independent mixture model.

As regards this question, all but two criterion perform very well, see Table 2. In the first place,

H

Table 2: Selected \mathcal{G} in the second experiment for 100 realizations from $\text{HPM}(\mathcal{G}_4, \theta, 4)$ and 100 realizations from $\text{HPM}(\mathcal{G}_8, \theta, 4)$ using Pseudolikelihood Information Criterion (PLIC), mean field-like approximations ($\text{BIC}^{\text{MF-like}}$, BIC^{GBF}) and Block Likelihood Information Criterion (BLIC) for various sizes of blocks and border conditions.

HPM($\mathcal{G}_4, \theta, 4$)			HPM($\mathcal{G}_8, \theta, 4$)		
	\mathcal{G}_4	\mathcal{G}_8		\mathcal{G}_4	\mathcal{G}_8
PLIC	53	47	PLIC	0	100
$\text{BIC}^{\text{MF-like}}$	100	0	$\text{BIC}^{\text{MF-like}}$	0	100
BIC^{GBF}	100	0	BIC^{GBF}	0	100
$\text{BLIC}_{2 \times 2}^{\text{MF-like}}$	100	0	$\text{BLIC}_{2 \times 2}^{\text{MF-like}}$	0	100
$\text{BLIC}_{2 \times 2}$	100	0	$\text{BLIC}_{2 \times 2}$	59	41
			$\text{BLIC}_{4 \times 4}$	0	100

PLIC faces trouble to select the correct \mathcal{G}_4 . This illustrates the importance of the estimation of the interaction parameter ψ . We have observed that the ICM algorithm whilst providing good segmented field, produces poorer estimates of the parameter than the simulated field algorithm. This has an impact quite important since ψ sets the strength of interaction between neighboring nodes of the graph \mathcal{G} and is most representative of the spatial correlation. On the other hand, $\text{BLIC}_{2 \times 2}$ fails to select the neighborhood system for second order hidden Potts model $\text{HPM}(\mathcal{G}_8, \theta, 4)$. This conclusion can be simply explained by the fact that the block does not include enough spatial information to discriminate between the competing models. When the primary purpose is the selection of a dependency structure, we should use block large enough to be informative regarding the different neighborhood systems in competition.

Aside the two above exceptions, the good performances of all criteria can be surprising. The same experiment has been done for stronger noise with $\sigma_k = 0.75$ and $\sigma_k = 1$. The conclusion remains the same. It appears that for a conditionally independent noise process, neighborhood systems are readily distinguished close to the phase transition. This is not true for any parameter value as illustrated in the third experiment.

In the second instance, we supposed that K_T and \mathcal{G} were unknown, so that we were interested in the joint selection of the number of colors and of the dependency graph. For this example, the results remain the same than in Table 1 with the exception of PLIC. Indeed, the different criteria manage to differentiate the model in terms of the graph \mathcal{G} so that their performances are directly related to their ability to choose the correct number of colors.

4.4 Third experiment: BLIC versus ABC

This third experiment is the occasion to compare BLIC with the ABC procedures proposed by [Stoehr et al. \(2015\)](#). We return to the problem of solely selecting the dependency graph when the number of colors is known. We still consider a homoscedastic Gaussian noise but over bicolor Potts models ($K=2$). The standard deviation $\sigma_k = 0.39$, $k \in \{0, 1\}$, was set so that the probability of a wrong prediction of the latent color with a marginal MAP rule on the Gaussian model is about 10% in the thresholding step of the ABC procedure. Regarding the dependency parameter ψ , we set prior distributions below the phase transition which occurs at different levels depending on the

H

Table 3: Evaluation of the prior error rate of ABC procedures and of the error rate for the model choice criterion in the third experiment.

Train size	5,000	100,000	Criterion	Error rate
2D statistics	14.2%	13.8%	PLIC	19.8%
4D statistics	10.8%	9.8%	BIC ^{MF-like}	7.6%
6D statistics	8.6%	6.9%	BIC ^{GBF}	7.1%
Adaptive ABC	8.2%	6.7%	BLIC _{4×4}	7.7%

neighborhood structure. Precisely we used a uniform distribution over $(0; 1)$ when the adjacency is given by \mathcal{G}_4 and a uniform distribution over $(0; 0.35)$ with \mathcal{G}_8 . In order to examine the performance of model choice criteria in comparison of ABC, we carried out 1000 realizations from $\text{HPM}(\mathcal{G}_4, \theta, 2)$ and 1000 realizations from $\text{HPM}(\mathcal{G}_8, \theta, 2)$ with parameters from the priors. The results are presented in Table 3

The novel ABC procedure introduced by [Stoehr et al. \(2015\)](#) appears to provide the best performances but for a training reference table of size 100 000. This reinforces the idea that for unlimited computation possibilities, ABC can efficiently address situations where the likelihood is intractable. However, Table 3 suggest that for a much lower computational cost it is possible to get equivalent, or even better, error rate by using model choice criterion BIC^{MF-like}, BIC^{GBF} or BLIC_{b×b}, while PLIC seems not to be overtaken. In this example, BIC^{GBF} slightly supersede BIC^{MF-like} and BLIC_{b×b}. This can be explained by the fact that for parameter from the prior close to zero, the assumption of independence between the sites is almost true. In the latter case, estimating BIC using the first order approximations of the partition function of Gibbs distribution ([Forbes and Peyrard, 2003](#)) may be preferable than using normalizing constants defined on blocks.

5 Conclusion and perspective

In the present article, we considered BIC to perform model selection when dealing with hidden Markov random fields. To avoid time consuming simulation methods like MCMC or ABC algorithms, we proposed to move towards variational methods and in particular to use valid probability distributions over non-overlapping blocks of the lattice in place of the intractable likelihood (Section 3.2). Consequently, we derived Block Likelihood Information Criterion to discriminate between hidden Markov random fields.

The numerical results (Section 4) highlighted that the approximations of BIC based on independent blocks without fixed border provide better performances comparing to pre-existing criteria as regards the inference of the number of latent colors. This conclusion has to be brought into perspective for the selection of the dependency structure as the size of the blocks should be wide enough unless BLIC can be misleading. According to the numerical results (Section 4.4), the opportunity we have explored appears to be a satisfactory alternative to ABC model choice algorithms which besides their computational cost are delicate to calibrate (*e.g.*, [Stoehr et al., 2015](#)). Our approach offers thus an appealing trade-off between efficient computation and reliable results.

While the numerical part of the paper assess its efficiency, the novel criterion makes in its current version two major approximations that are worth exploring. First mention, the choice of a particular

substitute is lead by any optimality conditions. From that viewpoint, the construction of an optimal approximations regarding the variational free energy over the set of probability distributions that factorize on blocks is yet to be studied. The second level of approximations concerns the penalty term. The next step of our work cannot be reduced to the sole aim of improving the quality of the approximations. Through Section 4.2, we have seen that an optimal solution with respect to the Kullback- Leibler divergence is not sufficient to ensure a good behaviour of model choice criteria, especially if the more complex model are not enough penalized. The penalty term used is solely valid for independent variable. We have neglected the interaction within a block, an assumption that slightly modified the number of free parameter. The impact of dependence variables on the penalty term is a logical follow-up to our work.

References

- Alfö, M., Nieddu, L., and Vicari, D. (2008). A finite mixture model for image segmentation. *Statistics and Computing*, 18(2):137–150.
- Besag, J. E. (1974). Spatial Interaction and the Statistical Analysis of Lattice Systems (with Discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236.
- Besag, J. E. (1975). Statistical Analysis of Non-Lattice Data. *The Statistician*, 24:179–195.
- Besag, J. E. (1986). On the Statistical Analysis of Dirty Pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):259–302.
- Besag, J. E., York, J., and Mollié, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the institute of statistical mathematics*, 43(1):1–20.
- Celeux, G., Forbes, F., and Peyrard, N. (2003). EM procedures using mean field-like approximations for Markov model-based image segmentation. *Pattern Recognition*, 36(1):131–144.
- Clifford, P. (1990). Markov random fields in statistics. *Disorder in physical systems: A volume in honour of John M. Hammersley*, pages 19–32.
- Csiszár, I., Shields, P. C., et al. (2000). The consistency of the BIC Markov order estimator. *The Annals of Statistics*, 28(6):1601–1619.
- Csiszár, I. and Talata, Z. (2006). Consistent Estimation of the Basic Neighborhood of Markov Random Fields. *The Annals of Statistics*, 34(1):123–145.
- Cucala, L. and Marin, J.-M. (2013). Bayesian Inference on a Mixture Model With Spatial Dependence. *Journal of Computational and Graphical Statistics*, 22(3):584–597.
- Everitt, R. G. (2012). Bayesian Parameter Estimation for Latent Markov Random Fields and Social Networks. *Journal of Computational and Graphical Statistics*, 21(4):940–960.
- Forbes, F. and Peyrard, N. (2003). Hidden Markov random field model selection criteria based on mean field-like approximations. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(9):1089–1101.

- François, O., Ancelet, S., and Guillot, G. (2006). Bayesian Clustering Using Hidden Markov Random Fields in Spatial Population Genetics. *Genetics*, 174(2):805–816.
- Friel, N. (2012). Bayesian Inference for Gibbs Random Fields Using Composite Likelihoods. In *Proceedings of the Winter Simulation Conference*, number 28 in WSC '12, pages 1–8. Winter Simulation Conference.
- Friel, N. (2013). Evidence and Bayes Factor Estimation for Gibbs Random Fields. *Journal of Computational and Graphical Statistics*, 22(3):518–532.
- Friel, N., Pettitt, A. N., Reeves, R., and Wit, E. (2009). Bayesian Inference in Hidden Markov Random Fields for Binary Data Defined on Large Lattices. *Journal of Computational and Graphical Statistics*, 18(2):243–261.
- Friel, N. and Rue, H. (2007). Recursive computing and simulation-free inference for general factorizable models. *Biometrika*, 94(3):661–672.
- Gassiat, E. (2002). Likelihood ratio inequalities with applications to various mixtures. *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, 38(6):897–906.
- Geman, S. and Geman, D. (1984). Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741.
- Green, P. J. and Richardson, S. (2002). Hidden Markov Models and Disease Mapping. *Journal of the American Statistical Association*, 97(460):1055–1070.
- Grelaud, A., Robert, C. P., Marin, J.-M., Rodolphe, F., and Taly, J.-F. (2009). ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 4(2):317–336.
- Grimmett, G. R. (1973). A theorem about random fields. *Bulletin of the London Mathematical Society*, 5(1):81–84.
- Haughton, D. M. A. (1988). On the Choice of a Model to Fit Data from an Exponential Family. *The Annals of Statistics*, 16(1):342–355.
- Hurn, M. A., Husby, O. K., and Rue, H. (2003). A Tutorial on Image Analysis. In *Spatial Statistics and Computational Methods*, volume 173 of *Lecture Notes in Statistics*, pages 87–141. Springer New York.
- Ji, C. and Seymour, L. (1996). A consistent model selection procedure for Markov random fields based on penalized pseudolikelihood. *The annals of applied probability*, pages 423–443.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. (1999). An Introduction to Variational Methods for Graphical Models. *Machine learning*, 37(2):183–233.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, 90(430):773–795.
- Kass, R. E. and Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the schwarz criterion. *Journal of the American Statistical Association*, 90(431):928–934.

- Keribin, C. (2000). Consistent Estimation of the Order of Mixture Models. *Sankhy: The Indian Journal of Statistics, Series A (1961-2002)*, 62(1):49–66.
- Lindsay, B. G. (1988). Composite likelihood methods. *Contemporary Mathematics*, 80(1):221–39.
- Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). Approximate Bayesian Computational methods. *Statistics and Computing*, 22(6):1167–1180.
- McGrory, C. A., Titterton, D. M., Reeves, R., and Pettitt, A. N. (2009). Variational Bayes for estimating the parameters of a hidden Potts model. *Statistics and Computing*, 19(3):329–340.
- Moore, M. T., Hargrave, C. E., Harden, F., and Mengersen, K. (2014). Segmentation of cone-beam CT using a hidden Markov random field with informative priors. *Journal of Physics : Conference Series*, 489.
- Okabayashi, S., Johnson, L., and Geyer, C. J. (2011). Extending pseudo-likelihood for Potts models. *Statistica Sinica*, 21(1):331.
- Qian, W. and Titterton, D. (1991). Estimation of parameters in hidden Markov models. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 337(1647):407–428.
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological methodology*, 25:111–164.
- Reeves, R. and Pettitt, A. N. (2004). Efficient recursions for general factorisable models. *Biometrika*, 91(3):751–757.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464.
- Seymour, L. and Ji, C. (1996). Approximate Bayes model selection procedures for Gibbs-Markov random fields. *Journal of Statistical Planning and Inference*, 51(1):75–97.
- Stanford, D. C. and Raftery, A. E. (2002). Approximate Bayes factors for image segmentation: The pseudolikelihood information criterion (PLIC). *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(11):1517–1520.
- Stoehr, J. and Friel, N. (2015). Calibration of conditional composite likelihood for bayesian inference on gibbs random fields. In *JMLR W&CP: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38, pages 921–929.
- Stoehr, J., Pudlo, P., and Cucala, L. (2015). Adaptive ABC model choice and geometric summary statistics for hidden Gibbs random fields. *Statistics and Computing*, 25(1):129–141.
- Tierney, L. and Kadane, J. B. (1986). Accurate Approximations for Posterior Moments and Marginal Densities. *Journal of the American Statistical Association*, 81(393):82–86.