



HAL
open science

A review on statistical inference methods for discrete Markov random fields

Julien Stoehr, Richard Everitt, Matthew T. Moores

► **To cite this version:**

Julien Stoehr, Richard Everitt, Matthew T. Moores. A review on statistical inference methods for discrete Markov random fields. 2017. hal-01462078v2

HAL Id: hal-01462078

<https://hal.science/hal-01462078v2>

Preprint submitted on 11 Apr 2017

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.

A review on statistical inference methods for discrete Markov random fields

Julien Stoehr¹

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

Abstract

Developing satisfactory methodology for the analysis of Markov random field 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. This forms a central issue for any statistical approach as the likelihood is an integral part of the procedure. Furthermore, such unobserved fields cannot be integrated out and the likelihood evaluation becomes a doubly intractable problem. This report gives an overview of some of the methods used in the literature to analyse such observed or unobserved random fields.

Keywords: statistics; Markov random fields; parameter estimation; model selection.

1 Introduction

The problem of developing satisfactory methodology for the analysis of spatially correlated data has been of a constant interest for more than half a century now. Constructing a joint probability distribution to describe the global properties of such data is somewhat complicated but the difficulty can be bypassed by specifying the local characteristics via conditional probability instead. This proposition has become feasible with the introduction of Markov random fields (or Gibbs distribution) as a family of flexible parametric models for spatial data (*the Hammersley-Clifford theorem*, Besag, 1974). Markov random fields are spatial processes related to lattice structure, the conditional probability at each nodes of the lattice being dependent only upon its neighbours, that is useful in a wide range of applications. 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} .

Gibbs distributions originally appears in statistical physics to describe equilibrium state of a physical systems which consists of a very large number of interacting particles such as ferromagnet ideal gases (Lanford and Ruelle, 1969). But they have since been useful in many other modelling areas, surged by the development in the statistical community since the 1970's. Indeed, they 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. Shaped by the development of Geman and Geman (1984) and Besag (1986), – see for example Alfò et al. (2008) and Moores et al. (2014) who performed image segmentation with the help of this modelling – and 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. The exponential random graph model or p^* model (Wasserman and Pattison, 1996) is another prominent example (Frank and Strauss, 1986) and arguably the most popular statistical model for social network analysis (*e.g.*, Robins et al., 2007).

Interests in these models is not so much about Markov laws that may govern data but rather the flexible and stabilizing properties they offer in modelling. Whilst the Gibbs-Markov equivalence provides an explicit form of the joint distribution and thus a global description of the model, this is marred by a considerable computational curse. Conditional probabilities can be easily computed, but the joint and the marginal distribution are meanwhile unavailable since the normalising constant is of combinatorial complexity and generally can not be evaluated with standard analytical or numerical methods. 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 and model selection. 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 aforementioned algorithm grows exponentially and is thus helpless on large lattices. Many deterministic or stochastic approximations have been proposed for circumventing this difficulty and developing methods that are computationally efficient and accurate is still an area of active research. Solutions to deal with the intractable likelihood are of two kinds. On one hand, one can rely on pseudo-model as surrogates for the likelihood. Such solutions typically stems from composite likelihood or variational approaches. On the other hand Monte Carlo methods have played a major role to estimate the intractable likelihood in both frequentist and Bayesian paradigm.

The present survey paper cares about the problem of carrying out statistical inference (mostly in a Bayesian framework) for Markov random fields. When dealing with hidden random fields, the focus is solely on hidden data represented by discrete models such as the Ising or the Potts models. Both are widely used examples and representative of the general level of difficulty. Aims may be to infer on parameters of the model or on the latent state \mathbf{x} . The paper is organised as follows: it begins by introducing the existence of Markov random fields with some specific examples (Section 2). The difficulties inherent to the analysis of such a stochastic model are especially pointed out in Section 3. The special case of small regular lattices is described in Section 4. As befits a survey paper, Section 5 focuses on solutions based on pseudo-models while Section 6 is dedicated to a state of the art concerning sampling method.

2 Markov random field and Gibbs distribution

2.1 Gibbs-Markov equivalence

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 $\mathcal{X} := \{0, \dots, K - 1\}$. 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 realization of \mathbf{X}_A . Denote by $\mathcal{S} \setminus A = -A$ the complement of A in \mathcal{S} . When modelling local interactions, the sites are lying on an undirected graph \mathcal{G} which

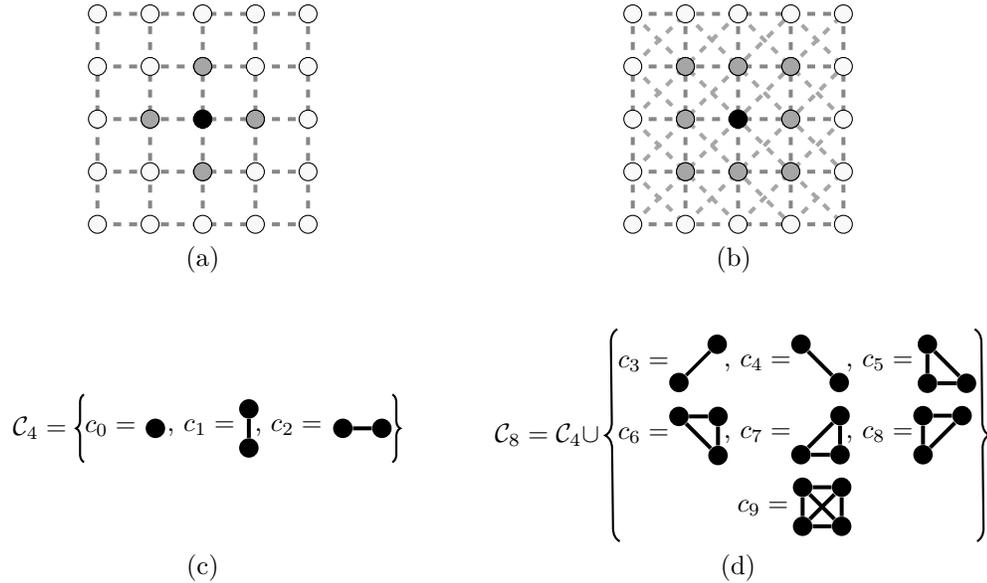


Figure 1: First and second order neighbourhood graphs \mathcal{G} with corresponding cliques. (a) The four closest neighbours graph \mathcal{G}_4 . neighbours of the vertex in black are represented by vertices in gray. (b) The eight closest neighbours graph \mathcal{G}_8 . neighbours of the vertex in black are represented by vertices in gray. (c) Cliques of graph \mathcal{G}_4 . (d) Cliques of graph \mathcal{G}_8 .

induces a topology on \mathcal{S} : by definition, sites i and j are adjacent or neighbour 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 it satisfies the following Markov property

$$\mathbf{P}(X_i = x_i \mid \mathbf{X}_{-i} = \mathbf{x}_{-i}) = \mathbf{P}(X_i = x_i \mid \mathbf{X}_{\mathcal{N}(i)} = \mathbf{x}_{\mathcal{N}(i)}), \tag{2.1}$$

where $\mathcal{N}(i)$ denotes the set of all the adjacent sites to i in \mathcal{G} . It is worth noting that any random field is a Markov random field with respect to the trivial topology, that is the cliques of \mathcal{G} are either the empty set or the entire set of sites \mathcal{S} . Recall a clique c in an undirected graph \mathcal{G} is any single vertex or a subset of vertices such that every two vertices in c are connected by an edge in \mathcal{G} . As an example, when modelling a digital image, the lattice is interpreted as a regular 2D-grid of pixels and the random variables states as shades of grey or colors. Two widely used adjacency structures are the graph \mathcal{G}_4 (first order lattice), respectively \mathcal{G}_8 (second order lattice), for which the neighbourhood of a site is composed of the four, respectively eight, closest sites on a two-dimensional regular lattice, except on the boundaries of the lattice, see Figure 1.

The difficulty with the Markov formulation is that one defines a set of conditional distributions which does not guarantee the existence of a joint distribution. 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 *e.g.*, [Grimmett \(1973\)](#), [Besag \(1974\)](#) and for a historical perspective [Clifford \(1990\)](#)), namely a density function on \mathcal{X} parametrised by $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^d$ and given with respect to the counting measure by

$$f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) = \frac{1}{Z(\boldsymbol{\theta}, \mathcal{G})} \exp\{-V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x})\} := \frac{1}{Z(\boldsymbol{\theta}, \mathcal{G})} q(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}), \tag{2.2}$$

where $V_{\boldsymbol{\theta}, \mathcal{G}}$ denotes the energy function which can be written as a sum over the set \mathcal{C} of all cliques of the graph, namely $V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x}) = \sum_{c \in \mathcal{C}} V_{\boldsymbol{\theta}, c}(\mathbf{x})$ for all configuration $\mathbf{x} \in \mathcal{X}$. The inherent difficulty of all these models arises from the intractable normalizing constant, sometimes called the partition function, defined by

$$Z(\boldsymbol{\theta}, \mathcal{G}) = \sum_{\mathbf{x} \in \mathcal{X}} \exp \{-V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x})\}.$$

The latter is a summation over the numerous possible realisations of the random field \mathbf{X} , that is of combinatorial complexity and cannot be computed directly (except for small grids and small number of states K). For binary variables X_i , the number of possible configurations reaches 2^n .

2.2 Autologistic model and related distributions

The formulation in terms of potential allows the local dependency of the Markov field to be specified and leads to a class of flexible parametric models for spatial data. In most cases, cliques of size one (singleton c_0) and two (doubletons $\mathcal{C} = \{c_1, c_2, c_3, c_4\}$) are assumed to be satisfactory to model the spatial dependency and potential functions related to larger cliques are set to zero. We present below popular examples of such pairwise models broadly used in the literature.

Autologistic model–Ising model The autologistic model first proposed by [Besag \(1972\)](#) is a pairwise-interaction Markov random field for binary (zero-one) spatial process. Denote $\boldsymbol{\theta} = (\alpha, \beta)$. The joint distribution is given by the following energy function

$$V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x}) = -\alpha \sum_{i=1}^n x_i - \beta \sum_{\substack{i \sim j \\ \mathcal{G}}} x_i x_j, \quad (2.3)$$

where the above sum $\sum_{i \sim j}^{\mathcal{G}}$ ranges the set of edges of the graph \mathcal{G} . The full-conditional probability thus writes like a logistic regression where the explanatory variables are the neighbours and themselves observations. The parameter α controls the level of 0 – 1 whereas the parameters $\{\beta\}$ model the dependency between two neighbouring sites i and j . One usually prefers to consider variables taking values in $\{-1, 1\}$ instead of $\{0, 1\}$ since it offers a more parsimonious parametrisation and avoids non-invariance issues when one switches states 0 and 1 as mentioned by [Pettitt et al. \(2003\)](#). A well known example is the general Ising model of ferromagnetism ([Ising, 1925](#)) that consists of discrete variables representing spins of atoms. The Gibbs distribution (2.3) is sometimes referred to as the Boltzmann distribution in statistical physics. The potential on singletons describes local contributions from external fields to the total energy. Spins most likely line up in the same direction of α , that is, in the positive, respectively negative, direction if $\alpha > 0$, respectively $\alpha < 0$. Putting differently α adjusts non-equal abundances of the two state values. The parameters $\{\beta\}$ represent the interaction strength between neighbours i and j . When $\beta > 0$ the interaction is called ferromagnetic and adjacent spins tend to be aligned, that is neighbouring sites with same sign have higher probability. When $\beta < 0$ the interaction is called anti-ferromagnetic and adjacent spins tend to have opposite signs.

Potts model The Potts model (Potts, 1952) originally appears in statistical mechanics to model interacting spins but has been used in other modelling area since then. It is a pairwise Markov random field that extends the Ising model to K possible states. The model sets a probability distribution on parametrized by $\boldsymbol{\theta} = (\alpha_0, \dots, \alpha_{K-1}, \beta)$, namely

$$\mathbf{V}_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x}) = - \sum_{i=1}^n \sum_{k=0}^{K-1} \alpha_k \mathbf{1}\{x_i = k\} - \beta \sum_{i \sim_j^{\mathcal{G}}} \mathbf{1}\{x_i = x_j\}, \quad (2.4)$$

where $\mathbf{1}\{A\}$ is the indicator function equal to 1 if A is true and 0 otherwise. Note that a potential function can be defined up to an additive constant. To ensure that potential functions on singletons are uniquely determined, one usually imposes the constraint $\sum_{k=0}^{K-1} \alpha_k = 0$. The 2-states Potts model is equivalent to the Ising model up to a constant for interaction parameter β , that is $\beta_{\text{Potts}} = 2\beta_{\text{Ising}}$.

2.3 Phase transition

One major peculiarity of Markov random field in the absence of an external field (*i.e.*, $\alpha = 0$) is a symmetry breaking for large values of parameter β due to a discontinuity of the partition function when the number of sites n tends to infinity. When the parameter β is zero, the random field is a system of independent uniform variables and all configurations are equally distributed. Increasing β favours the variable X_i to be equal to the dominant state among its neighbours and leads to patches of like-valued variables in the graph, such that once β tends to infinity values x_i are all equal. The distribution thus becomes multi-modal. In physics this is known as phase transition. This transition phenomenon has been widely study in both physics and probability, see for example Georgii (2011) for further details. The two dimensional Ising model is known to have a phase transition at a critical value β_c . Onsager (1944) obtained an exact value of β_c for the Ising model on the first order square lattice, namely

$$\beta_c = \frac{1}{2} \log \left\{ 1 + \sqrt{2} \right\} \approx 0.44.$$

The latter extends to a K -states Potts model on the first order lattice

$$\beta_c = \log \left\{ 1 + \sqrt{K} \right\},$$

see for instance Matveev and Shrock (1996) for specific results to Potts model on the square lattice and Wu (1982) for a broader overview.

The transition is more rapid than the number of neighbours increases. To illustrate this point, Figure 2 gives the average proportion of homogeneous pairs of neighbours, and the corresponding variance, for 2-states Potts model on the first and second order lattices of size 100×100 . Indeed, phase transition corresponds to discontinuity at β_c of $\beta \mapsto \lim_{n \rightarrow \infty} \frac{1}{n} \nabla \log Z(\beta, \mathcal{G})$. Since $\nabla^2 \log Z(\beta, \mathcal{G}) = \mathbf{Var} \{ \mathbf{S}(\mathbf{X}) \}$, where $\mathbf{S}(\mathbf{X}) = \sum_{i \sim_j^{\mathcal{G}}} \mathbf{1}\{X_i = X_j\}$ is the number of homogeneous pairs of a Potts random field \mathbf{X} , the discontinuity condition can thus be written as

$$\lim_{\beta \rightarrow \beta_c} \lim_{n \rightarrow \infty} \mathbf{Var} \{ \mathbf{S}(\mathbf{X}) \} = \infty.$$

Mention this is all theoretical asymptotic considerations and the discontinuity does not show itself on finite lattice realizations but the variance becomes increasingly sharper as the size grows.

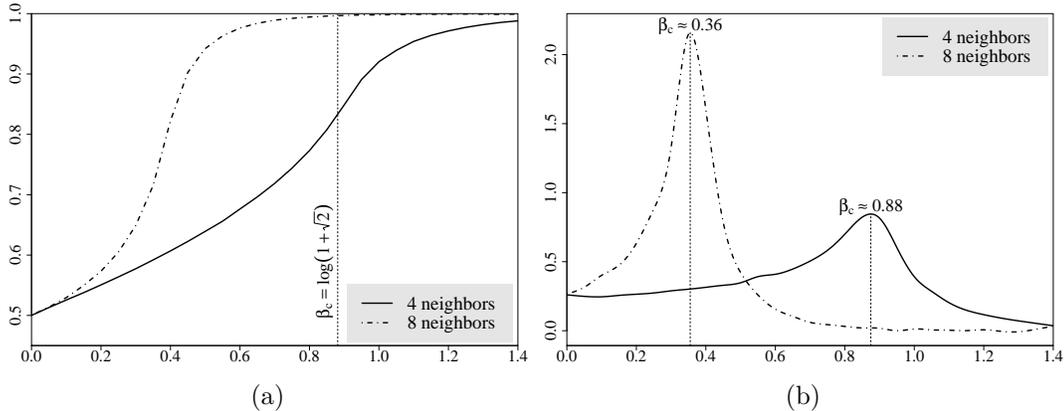


Figure 2: Phase transition for a 2-states Potts model with respect to the first order and second order 100×100 regular square lattices. (a) Average proportion of homogeneous pairs of neighbours. (b) Variance of the number of homogeneous pairs of neighbours.

2.4 Hidden Gibbs random field

Hidden Markov random fields has encountered a large interest over the past decade. It offer an appropriate representation for practical settings where the true state is unknown and observed indirectly through another field; this permits the modelling of noise that may happen upon many concrete situations: image analysis, (*e.g.*, Besag, 1986, Forbes and Peyrard, 2003, Alfö et al., 2008, Moores et al., 2014), disease mapping (*e.g.*, Green and Richardson, 2002), genetic analysis (François et al., 2006), to name but a few. The unobserved data is modelled as a discrete Markov random field \mathbf{X} associated to an energy function $V_{\theta, \mathcal{G}}$, as defined in (2.2). 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

$$g_{\phi}(\mathbf{y} | \mathbf{x}) = \prod_{i \in \mathcal{S}} g_i(y_i | x_i, \phi),$$

where g_i 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, 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 neighbours have also been explored (Besag, 1986).

Assuming that all the marginal distributions g_i are positive, one may write the joint distribution of (\mathbf{X}, \mathbf{Y}) , also called the complete likelihood, as

$$\mathbf{P}(\mathbf{x}, \mathbf{y} | \theta, \phi, \mathcal{G}) = \frac{1}{Z(\theta, \mathcal{G})} \exp \left\{ -V_{\theta, \mathcal{G}}(\mathbf{x}) + \sum_{i \in \mathcal{S}} \log g_i(y_i | x_i, \phi) \right\}.$$

The conditional field \mathbf{X} given $\mathbf{Y} = \mathbf{y}$ is thus a Markov random field and the noise can be interpreted as a non homogeneous external potential on singleton which is a bond to the unobserved data.

3 Statistical analysis issues

The intractable normalising constant $Z(\boldsymbol{\theta}, \mathcal{G})$ forms a central issue for both parameter and model selection problems as the likelihood is an integral part of the statistical procedure. Below, we introduce some of the classical problems of the literature

Maximum likelihood estimator Under the statistical model $f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})$, computing the maximum likelihood estimator, namely

$$\hat{\boldsymbol{\theta}}_{\text{MLE}} = \arg \max_{\boldsymbol{\theta}} \log f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}), \quad (3.1)$$

is challenging. Indeed, closed-form gradients are typically out of reach. Furthermore, one cannot rely on differentiation techniques such as automatic differentiation (*e.g.*, [Neidinger, 2010](#)) since point-wise estimation is impossible due to the intractability issue.

Computations of posterior distributions Consider a Bayesian posterior distribution expressed as

$$\pi(\boldsymbol{\theta} \mid \mathbf{x}) \propto f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) \pi(\boldsymbol{\theta}), \quad (3.2)$$

where $f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})$ denoted the likelihood of the observed data $\mathbf{x} \in \mathcal{X}$ and $\pi(\boldsymbol{\theta})$ denotes a prior density on the parameter space Θ with respect to a reference measure (often the Lebesgue measure of the Euclidean space). Here we are concerned with the situation where the unnormalised posterior distribution, the right-hand-side of (3.2) is intractable. This complication results in what is often termed a doubly-intractable posterior distribution, since the posterior distribution itself is normalised by the evidence (or marginal likelihood) which is typically also intractable.

Bayesian model choice Model choice is a problem of probabilistic model comparison. Assume we are given a set $\mathcal{M} = \{m : 1, \dots, M\}$ of stochastic models with respective parameter spaces Ψ_m embedded into Euclidean spaces of various dimensions. Bayesian approach to model selection considers the model itself as an unknown parameter of interest. The joint Bayesian distribution sets: a prior π on the model space, a prior density π_m on each parameter space Ψ_m with respect to a reference measure (often the Lebesgue measure of the Euclidean space), the likelihood f_m of the data \mathbf{y} within each model. On the extended parameter space $\Psi = \bigcup_{m=1}^M \{m\} \times \Psi_m$, the Bayesian analysis targets posterior model probabilities, that is the marginal in \mathcal{M} of the posterior distribution for $(m, \boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_M)$ given $\mathbf{Y} = \mathbf{y}$,

$$\pi(m \mid \mathbf{y}) = \frac{e(\mathbf{y} \mid m) \pi(m)}{\sum_{m'=1}^M e(\mathbf{y} \mid m') \pi(m')},$$

where $e(\mathbf{y} | m)$ denotes the evidence (or integrated likelihood) of model m defined as

$$e(\mathbf{y} | m) = \int_{\Psi_m} f_m(\mathbf{y} | \boldsymbol{\psi}_m) \pi_m(\boldsymbol{\psi}_m) d\boldsymbol{\psi}_m. \quad (3.3)$$

When the goal of the Bayesian analysis is the selection of the model that best fits the observed data \mathbf{y} , it is performed through the maximum *a posteriori* (MAP) defined by

$$\hat{m}_{\text{MAP}}(\mathbf{y}) = \arg \max_m \pi(m | \mathbf{y}). \quad (3.4)$$

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 (3.3) of each model. Under the assumption of model being equally likely *a priori*, the MAP rule (3.4) is equivalent to choose the model with the largest evidence (3.3) and in place of a fully Bayesian approach, model choice criterion can be used. However, the evidence can usually not be computed with standard procedure because of a high-dimensional integral. Thus much of the research in model selection area focuses on evaluating it by numerical methods.

Selecting the model among a collection of Markov random fields is a daunting task as none of (3.3) and (3.4) are analytically available. The model selection problem for hidden Markov random fields is even more complicated and could be termed as a triply-intractable problem. Indeed in addition to the integral on Ψ_m which is typically intractable, the stochastic model for \mathbf{Y} is based on the latent process \mathbf{X} in \mathcal{X} , that is

$$f_m(\mathbf{y} | \boldsymbol{\psi}_m = (\boldsymbol{\theta}_m, \boldsymbol{\phi}_m)) = \int_{\mathcal{X}} g_{\boldsymbol{\phi}_m}(\mathbf{y} | \mathbf{x}) f(\mathbf{x} | \boldsymbol{\theta}_m, \mathcal{G}_m) \mu(d\mathbf{x}), \quad (3.5)$$

with μ the counting measure (discrete case). Both the integral and the Gibbs distribution are intractable and consequently so is the posterior distribution.

4 Recursive algorithm for small Markov random field

When the Markov random field is defined on a small enough lattices, it is possible to answer the difficulty of computing the normalising constant by relying on generalised recursions for general factorisable models Reeves and Pettitt (2004). This method is based on an algebraic simplification due to the reduction in dependence arising from the Markov property. It applies to unnormalized likelihoods that can be expressed as a product of factors, each of which is dependent on only a subset of the lattice sites. More specifically the unnormalized version of a Gibbs distribution can be write as

$$q(\mathbf{x} | \boldsymbol{\theta}, \mathcal{G}) = \prod_{i=1}^{n-r} q_i(\mathbf{x}_{i:i+r} | \boldsymbol{\theta}, \mathcal{G}),$$

where each factor q_i depends on a subset $\mathbf{x}_{i:r} = \{x_i, \dots, x_{i+r}\}$ of \mathbf{x} , where r is defined to be the *lag* of the model.

As a result of this factorisation, the summation for the normalizing constant can be represented as

$$Z(\boldsymbol{\theta}, \mathcal{G}) = \sum_{\mathbf{x}_{n-r:n}} q_{n-r}(\mathbf{x}_{n-r:n} \mid \boldsymbol{\theta}, \mathcal{G}) \cdots \sum_{\mathbf{x}_{1:1+r}} q_1(\mathbf{x}_{1:1+r} \mid \boldsymbol{\theta}, \mathcal{G}).$$

The latter can be computed much more efficiently than the straightforward summation over the K^n possible lattice realisations using the following steps

$$\begin{aligned} Z_1(\mathbf{x}_{2:1+r}) &= \sum_{x_1} q_1(\mathbf{x}_{1:1+r}), \\ Z_i(\mathbf{x}_{i+1:i+r}) &= \sum_{x_i} q_i(\mathbf{x}_{i:i+r}) Z_{i-1}(\mathbf{x}_{i:i+r-1}), \text{ for all } i \in \{2, \dots, n-r\}, \\ Z(\boldsymbol{\psi}, \mathcal{G}) &= \sum_{\mathbf{x}_{n-r+1:n}} Z_{n-r}(\mathbf{x}_{n-r+1:n}). \end{aligned}$$

The complexity of the troublesome summation is significantly cut down since the forward algorithm solely relies on K^r possible configurations. Consider a rectangular lattice $h \times w = n$, where h stands for the height and w for the width of the lattice, with a first order neighbourhood system \mathcal{G}_4 (see Figure 1.(a)). The minimum *lag* representation for a pairwise model define on such a lattice occurs for r given by the smaller of the number of rows or columns in the lattice. Without the loss of generality, assume $h \leq w$ and lattice points are ordered from top to bottom in each column and columns from left to right. The complexity in time of the algorithm is then exponential in the number of rows and linear in the number of column. The algorithm of [Reeves and Pettitt \(2004\)](#) was extended in [Friel and Rue \(2007\)](#) to also allow exact draws from $\pi(\mathbf{x} \mid \boldsymbol{\psi}, \mathcal{G})$ for small enough lattices. The reader can find below an example of implementation for the general Potts model.

Following the work of [Friel and Rue \(2007\)](#), the R-package GiRaF available on CRAN proposes amongst other tools an ingenious implementation of those recursions for the autologistic, the Ising and the Potts model. Indeed, a naive implementation of the aforementioned recursions for such models can substantially increase the cost in time and memory allocation. For instance, the unnormalized Potts distribution associated to energy function (2.4) writes as

$$q(\mathbf{x} \mid \boldsymbol{\psi}, \mathcal{G}_4) = \prod_{i=1}^{n-h} q_i(\mathbf{x}_{i:i+h} \mid \boldsymbol{\psi}, \mathcal{G}_4),$$

where

- for all lattice point i except the ones on the last row or last column

$$q_i(\mathbf{x}_{i:i+h} \mid \boldsymbol{\psi}, \mathcal{G}_4) = \exp \left(\sum_{k=0}^{K-1} \alpha_k \mathbf{1}\{x_i = k\} + \beta_0 \mathbf{1}\{x_i = x_{i+1}\} + \beta_1 \mathbf{1}\{x_i = x_{i+h}\} \right). \quad (4.1)$$

- When lattice point i is on the last row x_{i+1} drops out of (4.1), that is

$$q_i(\mathbf{x}_{i:i+h} \mid \boldsymbol{\psi}, \mathcal{G}_4) = \exp \left(\sum_{k=0}^{K-1} \alpha_k \mathbf{1}\{x_i = k\} + \beta_1 \mathbf{1}\{x_i = x_{i+h}\} \right). \quad (4.2)$$

- The last factor takes into account all potentials within the last column

$$q_{n-h}(\mathbf{x}_{n-h:n} \mid \psi, \mathcal{G}_4) = \exp \left(\sum_{i=n-h}^n \sum_{k=0}^{K-1} \alpha_k \mathbf{1}\{x_i = k\} + \beta_1 \mathbf{1}\{x_{n-h} = x_n\} + \beta_0 \sum_{i=n-h+1}^n \mathbf{1}\{x_i = x_{i+1}\} \right).$$

One shall remark that for a homogeneous random field (*i.e.*, parameters are independent of the location of the sites), factors (4.1) and (4.2) only depend on the value of the random variables $\mathbf{X}_{i:i+h}$ but not on the actual position of the sites. Hence the number of factors to be computed is $2K^h$ instead of $h(w-1)K^h$. In term of implementation that also means factors can be computed for the different possible configurations once upstream the recursion. Furthermore with a first order neighbourhood, factor at a site merely involves its neighbour below and on its right, thereby reducing the number of possible factor to $K^3 + K^2$. Finally, mention it is straightforward to extend this algorithm to hidden Markov random field since as already mention in Section 2.4 the noise corresponds to a non homogeneous potential on singleton and hence the model still writes as a general factorisable model.

5 Pseudo-model and variational approaches

A point of view to overcome model’s bottleneck is to replace the true model with another pseudo-model selected among a collection of much simpler probability distribution, like for variational Bayes (Jaakkola and Jordan, 2000), or with an easily-normalised full conditional distributions (Lindsay, 1988). Both options have been explored in the literature and we discuss some of the solutions below.

5.1 Composite likelihood

A composite likelihood (Lindsay, 1988) approximates the joint distribution as the product of easy normalised full-conditional distributions

$$f_{\text{CL}}(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) = \prod_{i=1}^C f(\mathbf{x}_{A(i)} \mid \mathbf{x}_{B(i)}, \boldsymbol{\theta}, \mathcal{G})^{w_i}, \tag{5.1}$$

where $\{A(i)\}_{i=1}^C \subseteq \mathcal{P}(\mathcal{S})$ and $\{B(i)\}_{i=1}^C \subseteq \mathcal{P}(\mathcal{S})$ denote sets of subset of \mathcal{S} . It has encountered considerable interests in the statistics literature and the reader may refer to Varin et al. (2011) for a comprehensive overview. One of the earliest approach using composite likelihood is the pseudolikelihood (Besag, 1975) which approximates the joint distribution of \mathbf{x} as the product of full-conditional distributions for each site i ,

$$f_{\text{pseudo}}(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) = \prod_{i=1}^n f(x_i \mid \mathbf{x}_{-i}, \boldsymbol{\theta}, \mathcal{G}). \tag{5.2}$$

The pseudolikelihood (5.2) is not a genuine probability distribution, except if the random variables X_i are independent. The Markov property ensures that each term in the product only involves nearest neighbours, and so the normalising constant of each full-conditional is straightforward to compute. It is straightforward to prove that a unique maximum exists and it is easy to compute.

Geman and Graffigne (1986) demonstrate the consistency of the maximum pseudolikelihood estimator

$$\hat{\boldsymbol{\theta}}_{\text{MPLE}} = \arg \max_{\boldsymbol{\theta}} \log f_{\text{pseudo}}(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}).$$

when the lattice size tends to infinity for discrete Markov random field but maximum pseudolikelihood estimator has generally larger asymptotic variance than maximum likelihood estimator and does not achieve the Cramer-Rao lower bound (Lindsay, 1988). Practically, this approximation has been shown to lead to unreliable estimates of $\boldsymbol{\theta}$ (e.g., Rydén and Titterton, 1998, Friel and Pettitt, 2004, Cucala et al., 2009, Friel et al., 2009).

5.2 Variational approaches and parameter estimation

Variational methods refer to a class of deterministic approaches. They consist in introducing a variational function as an approximation to the likelihood in order to solve a simplified version of an optimization problem. It 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 a probability distribution \mathbf{P} and the target distribution $f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})$ as

$$F(\mathbf{P}) = -\log Z(\boldsymbol{\theta}, \mathcal{G}) + \text{KL}(\mathbf{P} \parallel f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})). \quad (5.3)$$

Although the Kullback–Leibler divergence is not a true metric, it has the non-negative property with divergence zero if and only if distributions are equal almost everywhere. The free energy has then an optimal lower bound achieved for $\mathbf{P} = f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})$. Minimizing the free energy with respect to the set of probability distribution on \mathcal{X} 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} . 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. namely for all \mathbf{x} in $\mathcal{X} = \prod_{i \in \mathcal{S}} \mathcal{X}_i$,

$$\mathbf{P}(\mathbf{x}) = \prod_{i \in \mathcal{S}} \mathbf{P}_i(x_i), \text{ where } \mathbf{P}_i \in \mathcal{M}_1^+(\mathcal{X}_i) \text{ and } \mathbf{P} \in \mathcal{M}_1^+(\mathcal{X}).$$

The minimization of (5.3) over this set leads to fixed point equations for each marginal of \mathbf{P} and the optimal solution is the so called mean field approximation (see for example Jordan et al., 1999).

Variational EM algorithm

In what follows we focus on the issue of estimating the parameters of a hidden Markov random fields. Let consider \mathbf{y} a noisy version of a Markov random field \mathbf{x} . One can write the log-likelihood as follows

$$\log f(\mathbf{y} | \boldsymbol{\psi}) = \underbrace{\sum_{\mathbf{x} \in \mathcal{X}} \log \left\{ \frac{f(\mathbf{x}, \mathbf{y} | \boldsymbol{\psi}, \mathcal{G})}{\mathbf{P}(\mathbf{x})} \right\} \mathbf{P}(\mathbf{x})}_{=F(\mathbf{P}, \boldsymbol{\psi})} + \underbrace{\sum_{\mathbf{x} \in \mathcal{X}} \log \left\{ \frac{\mathbf{P}(\mathbf{x})}{f(\mathbf{x} | \mathbf{y}, \boldsymbol{\psi}, \mathcal{G})} \right\} \mathbf{P}(\mathbf{x})}_{=\text{KL}(\mathbf{P} \| f(\cdot | \mathbf{y}, \boldsymbol{\psi}, \mathcal{G}))}. \quad (5.4)$$

Solutions based on variational approaches focus on minimising the KL-term or equivalently maximising the F . For instance, this relaxation of the original issue has shown good performances for approximating the maximum likelihood estimate (Celeux et al., 2003), as well as for Bayesian inference on hidden Markov random fields (McGrory et al., 2009). We discuss it below.

Celeux et al. (2003) explore the opportunity to use the Expectation-Maximization (EM) algorithm (Dempster et al., 1977). To address the problem of finding the maximum likelihood estimator for the statistical model $f(\mathbf{y} | \boldsymbol{\psi})$, the EM algorithm iterates between two steps:

1. **E step:** one computes the conditional expectation of the log-joint distribution with respect to the distribution of \mathbf{X} given $\mathbf{Y} = \mathbf{y}$ at the current parameter value $\boldsymbol{\psi}^{(t)}$

$$\begin{aligned} Q(\boldsymbol{\psi} | \boldsymbol{\psi}^{(t)}) &= \mathbf{E}_{\boldsymbol{\psi}^{(t)}} \{ \log f(\mathbf{X}, \mathbf{y} | \boldsymbol{\psi}, \mathcal{G}) | \mathbf{Y} = \mathbf{y} \} \\ &= \sum_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x} | \mathbf{y}, \boldsymbol{\psi}^{(t)}, \mathcal{G}) \log f(\mathbf{x}, \mathbf{y} | \boldsymbol{\psi}, \mathcal{G}). \end{aligned}$$

2. **M step:** one maximises Q with respect to $\boldsymbol{\psi}$,

$$\boldsymbol{\psi}^{(t+1)} = \arg \max_{\boldsymbol{\psi}} Q(\boldsymbol{\psi} | \boldsymbol{\psi}^{(t)}).$$

We refer the reader to Wu (1983) for convergence results. The EM scheme cannot be applied directly to hidden Markov random fields as it yields analytically intractable updates. The function Q can be written as

$$Q(\boldsymbol{\psi} | \boldsymbol{\psi}^{(t)}) = \mathbf{E}_{\boldsymbol{\psi}^{(t)}} \{ \log g_{\phi}(\mathbf{y} | \mathbf{X}) | \mathbf{Y} = \mathbf{y}, \} + \mathbf{E}_{\boldsymbol{\psi}^{(t)}} \{ \log f(\mathbf{X} | \boldsymbol{\theta}, \mathcal{G}) | \mathbf{Y} = \mathbf{y} \}.$$

and hence requires evaluating $f(\cdot | \boldsymbol{\theta}, \mathcal{G})$ and $f(\cdot | \mathbf{y}, \boldsymbol{\psi}, \mathcal{G})$ (for computing $\mathbf{E}_{\boldsymbol{\psi}}$) which are both unavailable. Many stochastic or deterministic schemes have been proposed to handle intractability in the EM steps, such as the Gibbsian-EM (Chalmond, 1989), the Monte-Carlo EM (Wei and Tanner, 1990) or the Restoration-Maximization algorithm (Qian and Titterton, 1991). The aim of the variational EM (VEM) is to maximize the function F instead of Q in order to get a tractable version of the EM algorithm. This shift in the formulation leads to an alternating optimization procedure which can be described as follows: let denote \mathcal{D} a set of probability distributions on the latent space, given a current value $(\mathbf{P}^{(t)}, \boldsymbol{\theta}^{(t)})$ in $\mathcal{D} \times \Theta$, updates with

$$\mathbf{P}^{(t+1)} = \arg \max_{\mathbf{P} \in \mathcal{D}} F(\mathbf{P}, \boldsymbol{\psi}^{(t)}) = \arg \min_{\mathbf{P} \in \mathcal{D}} \text{KL}(\mathbf{P} \| f(\cdot | \mathbf{y}, \boldsymbol{\psi}^{(t)}, \mathcal{G})), \quad (5.5)$$

$$\boldsymbol{\psi}^{(t+1)} = \arg \max_{\boldsymbol{\psi}} F(\mathbf{P}^{(t+1)}, \boldsymbol{\psi}) = \arg \max_{\boldsymbol{\theta}} \sum_{\mathbf{x} \in \mathcal{X}} \mathbf{P}^{(t+1)}(\mathbf{x}) \log \pi(\mathbf{x}, \mathbf{y} | \boldsymbol{\psi}, \mathcal{G}). \quad (5.6)$$

Optimising F over the class of independent probability distributions \mathbf{P} that factorize on sites leads to mean field approximation. Generalizing an idea originally introduced by [Zhang \(1992\)](#), [Celeux et al. \(2003\)](#) have designed a class of VEM-like algorithm that uses mean field-like approximations for both $f(\cdot | \mathbf{y}, \boldsymbol{\psi}, \mathcal{G})$ and $f(\cdot | \boldsymbol{\theta}, \mathcal{G})$. To put it in simple terms mean field-like approximations refer to distributions for which neighbours of site i are set to constants. Given a configuration $\tilde{\mathbf{x}}$ in \mathcal{X} , the Gibbs distribution $f(\cdot | \boldsymbol{\theta}, \mathcal{G})$ is replaced by

$$\mathbf{P}_{\tilde{\mathbf{x}}}^{\text{MF-like}}(\mathbf{x} | \boldsymbol{\theta}, \mathcal{G}) = \prod_{i \in \mathcal{S}} \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}, \boldsymbol{\theta}, \mathcal{G}).$$

The main difference with the pseudolikelihood (5.2) is that neighbours are not random anymore and setting them to constant values leads to a system of independent variables. From this approximation, the EM path is set up with the corresponding joint distribution approximation

$$\mathbf{P}^{\text{MF-like}}(\mathbf{x}, \mathbf{y} | \boldsymbol{\psi}, \mathcal{G}) = \prod_{i \in \mathcal{S}} g_{\phi}(y_i | x_i) \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}, \boldsymbol{\theta}, \mathcal{G}).$$

Note that this general procedure corresponds to the so-called point-pseudo-likelihood EM algorithm proposed by [Qian and Titterton \(1991\)](#). The flexibility of the approach proposed by [Celeux et al. \(2003\)](#) lies in the choice of the configuration $\tilde{\mathbf{x}}$ that is not necessarily a valid configuration for the model. We refer the reader to [Celeux et al. \(2003\)](#) for further details. When the neighbours $\mathbf{X}_{\mathcal{N}(i)}$ are fixed to their mean value, or more precisely $\tilde{\mathbf{x}}$ is set to the mean field estimate of the complete conditional distribution $f(\mathbf{x} | \mathbf{y}, \boldsymbol{\psi}, \mathcal{G})$, this results in the Mean Field algorithm of [Zhang \(1992\)](#). In practice, [Celeux et al. \(2003\)](#) obtain better performances with their so-called Simulated Field algorithm (see Algorithm 1). In this stochastic version of the EM-like procedure, $\tilde{\mathbf{x}}$ is a realization drawn from the conditional distribution $f(\cdot | \mathbf{y}, \boldsymbol{\psi}^{(t)}, \mathcal{G})$ for the current value of the parameter $\boldsymbol{\psi}^{(t)}$. The latter is preferred to other methods when dealing with maximum-likelihood estimation for hidden Markov random field. This extension of VEM algorithms suffers from a lack of theoretical support due to the propagation of the approximation to the Gibbs distribution $f(\cdot | \boldsymbol{\theta}, \mathcal{G})$. One might advocate in favour of the Monte-Carlo VEM algorithm of [Forbes and Fort \(2007\)](#) for which convergence results are available. However the Simulated Field algorithm provides better results for the estimation of the spatial parameter, as illustrated in [Forbes and Fort \(2007\)](#).

5.3 Approximating model choice criteria

Various approximations have been proposed to overcome intractability of (3.3) but a commonly used one, if only for its simplicity, is the Bayesian Information Criterion (BIC) that is an asymptotic estimate of the evidence based on the Laplace method for integrals ([Schwarz, 1978](#), [Kass and Raftery, 1995](#)). The criterion is a simple penalized function of the maximized log-likelihood

$$\text{BIC}(m) = -2 \log f_m(\mathbf{y} | \hat{\boldsymbol{\psi}}_{\text{MLE}}) + d_m \log(n) \approx -2 \log e(\mathbf{y} | m), \quad (5.7)$$

where d_m is the number of free parameters of model m (usually the dimension of $\boldsymbol{\Psi}_m$) and $n = |\mathcal{S}|$ is the number of sites. The $d_m \log(n)$ term corresponds to a penalty term which

Algorithm 1: Simulated Field algorithm**Input:** an observation \mathbf{y} , a number of iterations T **Output:** an estimate of the likelihood maximum estimator $\hat{\psi}_{\text{MLE}}$ **Initialization:** start from an initial guess $\boldsymbol{\psi}^{(0)} = (\boldsymbol{\theta}^{(0)}, \boldsymbol{\phi}^{(0)})$;**for** $t \leftarrow 1$ **to** T **do****neighbourhood restoration:** draw $\tilde{\mathbf{x}}^{(t)}$ from $\pi(\cdot \mid \mathbf{y}, \boldsymbol{\psi}^{(t-1)}, \mathcal{G})$;**E-step:** compute

$$\widehat{Q}_1(\boldsymbol{\phi}) := \sum_{i \in \mathcal{S}} \sum_{x_i} \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}^{(t)}, y_i, \boldsymbol{\psi}^{(t-1)}, \mathcal{G}) \log g_{\boldsymbol{\phi}}(y_i \mid x_i);$$

$$\widehat{Q}_2(\boldsymbol{\theta}) := \sum_{i \in \mathcal{S}} \sum_{x_i} \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}^{(t)}, y_i, \boldsymbol{\psi}^{(t-1)}, \mathcal{G}) \log \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}^{(t)}, \boldsymbol{\theta}, \mathcal{G});$$

M-step: set $\boldsymbol{\psi}^{(t)} = (\boldsymbol{\theta}^{(t)}, \boldsymbol{\phi}^{(t)})$ where

$$\boldsymbol{\theta}^{(t)} = \arg \max_{\boldsymbol{\theta}} \widehat{Q}_1(\boldsymbol{\theta}) \text{ and } \boldsymbol{\psi}^{(t)} = \arg \max_{\boldsymbol{\psi}} \widehat{Q}_2(\boldsymbol{\psi});$$

end**return** $\boldsymbol{\psi}^{(T)} = (\boldsymbol{\theta}^{(T)}, \boldsymbol{\phi}^{(T)})$

increases with the complexity of the model and the model with the highest posterior probability is the one that minimizes BIC. The criterion is closely related to the Akaike Information Criterion (AIC, [Akaike, 1973](#)) that solely differs in the penalization term. We refer the reader to [Kass and Raftery \(1995\)](#) and the references therein for a more detailed discussion on AIC and for instance to [Burnham and Anderson \(2002\)](#) for comparison between AIC and BIC.

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](#)) or 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 neighbourhood system, [Csiszár and Talata \(2006\)](#) proposed to replace the likelihood by the pseudolikelihood (5.2) and modify the penalty term as the number of all possible configurations for the neighbouring 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 fields.

As already mention, in the context of Markov random fields, difficulties are of two kind. Neither the maximum likelihood estimate $\hat{\boldsymbol{\theta}}_{\text{MLE}}$ nor the likelihood f_m are available. Recall that in the hidden case f_m requires to integrate a Gibbs distribution over the latent space configuration. As regards the simplest 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](#)). Over the past decade, only few works have addressed the model choice issue for hidden Markov random field from that BIC perspective. Below, we describe solutions based

on pseudo-models but other attempts based on simulations techniques have been investigated (Newton and Raftery, 1994).

The central question is the evaluation of the integrated likelihood (3.5). A convenient way to circumvent the issues of computing BIC is to replace the Gibbs distribution by tractable surrogates. Consider a partition of \mathcal{S} into subsets of neighbouring sites, namely

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

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

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

The most straightforward approach is to look for mean-field like approximations which take the following form

$$\mathbf{P}_{\tilde{\mathbf{x}}}(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) = \prod_{i=1}^n \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}, \boldsymbol{\theta}, \mathcal{G}) \approx f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}).$$

Put in other words, the latter is a surrogate in the class of independent probability distributions that factorize with respect to the nodes of \mathcal{G} and for which neighbourhood of each site has been set to a constant field $\tilde{\mathbf{x}}$. One shall remark we recover the mean-field approximations presented in Section 5.2 when $\tilde{\mathbf{x}}$ is set to the mean-field realisation. The integrated likelihood corresponding to $\mathbf{P}_{\tilde{\mathbf{x}}}$ is of the form

$$\mathbf{P}_m^{\text{MF-like}}(\mathbf{y} \mid \boldsymbol{\psi}_m = (\boldsymbol{\theta}_m, \boldsymbol{\phi}_m)) = \prod_{i \in \mathcal{S}} \sum_{x_i} g_{\boldsymbol{\phi}_m}(y_i \mid x_i) \mathbf{P}(x_i; \tilde{\mathbf{x}}_{\mathcal{N}(i)}, \boldsymbol{\theta}_m, \mathcal{G}).$$

This results in the following approximation of BIC

$$\text{BIC}^{\text{MF-like}}(m) = -2 \log \mathbf{P}_m^{\text{MF-like}}(\mathbf{y} \mid \hat{\boldsymbol{\psi}}_{\text{MLE}}) + d_m \log(n). \tag{5.8}$$

This approach includes the Pseudolikelihood Information Criterion (PLIC) of Stanford and Raftery (2002) as well as the mean field-like approximations of BIC proposed by Forbes and Peyrard (2003). The main difference between these criterion lies in the estimation $\boldsymbol{\psi}_{\text{MLE}}$ and the choice of $\tilde{\mathbf{x}}$. The idea of Stanford and Raftery (2002) is to consider as $\tilde{\mathbf{x}}$ a configuration close to the Iterated Conditional Modes (ICM, Besag, 1986) estimate of \mathbf{x} . In its unsupervised version ICM alternates between a restoration step of the latent states and an estimation step of the parameter $\boldsymbol{\psi}$. PLIC is the approximation of BIC based on the output of ICM algorithm

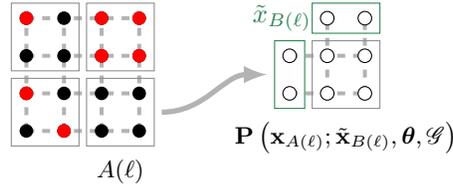


Figure 3: Factorisation of the Gibbs distribution approximation over a set of contiguous block $A(\ell)$ with border $B(\ell)$ set to a constant field $\tilde{\mathbf{x}}$.

$(\mathbf{x}^{\text{ICM}}, \psi^{\text{ICM}})$. The solution proposed by Forbes and Peyrard (2003) is to use for $(\tilde{\mathbf{x}}, \hat{\psi}_{\text{MLE}})$ the output of the VEM-like algorithm based on the mean-field like approximations of (Celeux et al., 2003). More precisely, as regards neighbourhood restoration step, they advocate in favor of the simulated field algorithm (see Algorithm 1).

Stoehr et al. (2016) have extended previous approaches to tractable approximations that factorize over larger sets of nodes, namely blocks of a rectangular lattice, by taking advantage of the general recursion implemented in the R-package GiRaF (see Section 4). They suggest to chose surrogates of the form

$$\mathbf{P}(\mathbf{x}; \tilde{\mathbf{x}}, A(1), \dots, A(C), \boldsymbol{\theta}) = \prod_{\ell=1}^C \mathbf{P}(\mathbf{x}_{A(\ell)}; \mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}, \boldsymbol{\theta}, \mathcal{G}), \quad (5.9)$$

where $B(\ell)$ is the border of $A(\ell)$, see Figure 3, or the empty set. This leads to their so called Block Likelihood Information Criterion (BLIC) which approximates BIC as a summation of tractable normalising constant, namely

$$\text{BLIC}(m) = -2 \sum_{\ell=1}^C \frac{\log Z(\hat{\psi}_{\text{MLE}}, \mathcal{G}, \mathbf{y}_{A(\ell)}, \tilde{\mathbf{x}}_{B(\ell)})}{\log Z(\hat{\boldsymbol{\theta}}_{\text{MLE}}, \mathcal{G}, \tilde{\mathbf{x}}_{B(\ell)})} + d_m \log(|\mathcal{S}|), \quad (5.10)$$

where $Z(\boldsymbol{\theta}, \mathcal{G}, \tilde{\mathbf{x}}_{B(\ell)})$ is the partition function of block $A(\ell)$ with fixed border $\tilde{\mathbf{x}}_{B(\ell)}$ and $Z(\boldsymbol{\theta}, \mathcal{G}, \mathbf{y}_{A(\ell)}, \tilde{\mathbf{x}}_{B(\ell)})$ is the partition function of the conditional random field $\mathbf{X}_{A(\ell)}$ knowing $\mathbf{Y}_{A(\ell)} = \mathbf{y}_{A(\ell)}$ and $\mathbf{X}_{B(\ell)} = \tilde{\mathbf{x}}_{B(\ell)}$. The latter has shown better performances for estimating the number of latent component K as well as for the selection of the underlying dependency structure \mathcal{G} . The criterion nonetheless lack some theoretical support and is limited to regular lattices.

BLIC is to a certain extent related to the RDA approximations of partition functions proposed by Friel et al. (2009). Indeed, let $Z(\boldsymbol{\theta}, \mathcal{G})$ and $Z(\boldsymbol{\psi}, \mathcal{G})$ denote the respective normalizing constants of the latent and the conditional fields. Starting from the Bayes formula, BIC expression turns into

$$\text{BIC}(m) = -2 \frac{\log Z(\boldsymbol{\psi}, \mathcal{G})}{\log Z(\boldsymbol{\theta}, \mathcal{G})} + d_m \log(n).$$

Similarly to Friel et al. (2009), Stoehr et al. (2016) approximate the intractable normalising constants by a product of tractable normalising constant defined on contiguous sub-lattices. Looking

at the issue of estimating the partition function instead of estimating the Gibbs distribution has also been explored by [Forbes and Peyrard \(2003\)](#). They propose to use a first order approximation of the partition function arising from mean field theory. [Forbes and Peyrard \(2003\)](#) argue that the latter is more satisfactory than $\text{BIC}^{\text{MF-like}}(m)$ in the sense it is based on an optimal lower bound for the normalising constants contrary to the mean field-like approximations. However that does not ensure better results as regards model selection.

Regarding the question of inferring the number of latent states, one might advocate in favor of the Integrated Completed Likelihood (ICL, [Biernacki et al., 2000](#)). This opportunity has been explored by [Cucala and Marin \(2013\)](#) but their complex algorithm cannot be extended easily to other model selection problems such as choosing the dependency structure.

6 Sampling methods for Markov random fields

6.1 Sampling from Gibbs distribution

Sampling from a Gibbs distribution can be a daunting task due to the correlation structure on a high dimensional space, and standard Monte Carlo methods are impracticable except for very specific cases. In the Bayesian paradigm, Markov chain Monte Carlo (MCMC) methods have played a dominant role in dealing with such problems, the idea being to generate a Markov chain whose stationary distribution is the distribution of interest. The Ising model is one of these special cases where one can be drawn exactly from the model using coupling from the past ([Propp and Wilson, 1996](#), [Mira et al., 2001](#)). Nevertheless, such perfect sampling scheme is often prohibitively expensive or impossible to carry out for other models. We describe below two popular solutions even though they introduce a bias.

6.1.1 Gibbs sampler

The Gibbs sampler is a highly popular MCMC algorithm in Bayesian analysis starting with the influential development of [Geman and Geman \(1984\)](#). It can be seen as a component-wise Metropolis-Hastings algorithm ([Metropolis et al., 1953](#), [Hastings, 1970](#)) where variables are updated one at a time and for which proposal distributions are the full conditionals themselves, see Algorithm 2. It is particularly well suited to Markov random field since by nature the intractable joint distribution is fully determined by the easy to compute conditional distributions.

Under the irreducibility assumption, the chain converges to the target distribution $f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})$, see for example ([Geman and Geman, 1984](#), *Theorem A*). Note the order in which the components are updated in Algorithm 2 does not make much difference as long as every site is visited. Hence it can be deterministically or randomly modified, especially to avoid possible bottlenecks when visiting the configuration space. A synchronous version is nonetheless unavailable since updating the sites merely at the end of cycle t would lead to incorrect limiting distribution.

We should mention here that Gibbs sampler faces some well known difficulties when it is applied to the Ising or Potts model. The Markov chain mixes slowly, namely long range interactions require many iterations to be taken into account, such that switching the color of a large homogeneous area is of low probability even if the distribution of the colors is exchangeable. This

Algorithm 2: Gibbs sampler**Input:** a parameter θ , a number of iterations T **Output:** a sample \mathbf{x} from the joint distribution $f(\cdot \mid \theta, \mathcal{G})$ **Initialization:** draw an arbitrary configuration $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;**for** $t \leftarrow 1$ **to** T **do** **for** $i \leftarrow 1$ **to** n **do** **draw** $x_i^{(t)}$ from the full conditional $\mathbf{P}(X_i^{(t)} \mid \mathbf{x}_{\mathcal{N}(i)}^{(t-1)}, \theta)$; **end****end****return** the configuration $\mathbf{x}^{(T)}$

peculiarity is even worse when the parameter β is above the critical value of the phase transition, the Gibbs distribution being severely multi-modal (each mode corresponding to a single color configuration). Liu (1996) proposed a modification of the Gibbs sampler that overcome these drawbacks with a faster rate of convergence. Note also that in the context of Gaussian Markov random field some efficient algorithm have been proposed like the fast sampling procedure of Rue (2001).

6.1.2 Auxiliary variables and Swendsen-Wang algorithm

An appealing alternative to bypass slow mixing issues of the Gibbs sampler is the Swendsen-Wang algorithm (Swendsen and Wang, 1987) originally designed to speed up simulation of Potts model close to the phase transition. Swendsen-Wang algorithm iterates two steps : a clustering step and a swapping step (see Algorithm 3), in order to incorporate simultaneous updates of large homogeneous regions (e.g., Besag and Green, 1993). The clustering step relies on auxiliary random variables which aim at decoupling the complex dependence structure between the component of \mathbf{x} and yield a partition of sites into single-valued clusters or connected components. The method set binary (0-1) conditionally independent auxiliary variables U_{ij} which satisfy

$$\mathbf{P}(U_{ij} = 1 \mid \mathbf{x}) = \begin{cases} 1 - \exp(\beta_{ij} \mathbf{1}\{x_i = x_j\}) = p_{ij} & \text{if } i \not\sim j, \\ 0 & \text{otherwise} \end{cases}$$

with $\beta_{ij} \geq 0$ so that p_{ij} takes value between 0 and 1. The latter then represents the probability to keep an edge between neighbouring sites in \mathcal{G} . Auxiliary variables U_{ij} induce on \mathbf{x} a subgraph $\Gamma(\mathcal{G}, \mathbf{x})$ of the dependency graph \mathcal{G} , namely the undirected graph made of edges of \mathcal{G} for which $U_{ij} = 1$, see Figure 4. During the swapping step, each cluster \mathcal{C} of $\Gamma(\mathcal{G}, \mathbf{x})$ is assigned to a new state k with probability $\mathbf{P}(\mathbf{X}_{\mathcal{C}} = k) \propto \exp\{\sum_{i \in \mathcal{C}} \alpha_k\}$. For the special but important case where $\alpha = 0$, new possible states are equally likely. Also for large values of β , the algorithm manages to switch colors of wide areas, achieving a better cover of the configuration space.

For the original proof of convergence, we refer the reader to Swendsen and Wang (1987) and for further discussion see for example Besag and Green (1993). Whilst the ability to change large set of variables in one step seems to be a significant advantage, this can be marred by a slow mixing time, namely exponential in n (Gore and Jerrum, 1999). The mixing time of the algorithm is polynomial in n for Ising or Potts models with respect to the graphs \mathcal{G}_4 and \mathcal{G}_8 but only for

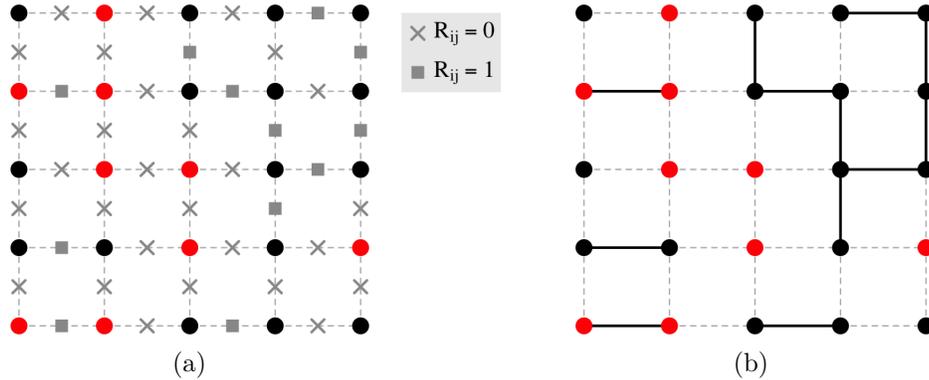


Figure 4: Auxiliary variables and subgraph illustrations for the Swendsen-Wang algorithm. (a) Example of auxiliary variables U_{ij} for a 2-states Potts model configuration on the first order square lattice. (b) Subgraph $\Gamma(\mathcal{G}_4, \mathbf{x})$ of the first order lattice \mathcal{G}_4 induced by the auxiliary variables U_{ij} .

Algorithm 3: Swendsen-Wang algorithm

Input: a parameter θ , a number of iterations T

Output: a sample \mathbf{x} from the joint distribution $f(\cdot \mid \theta, \mathcal{G})$

Initialization: draw an arbitrary configuration $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;

for $t \leftarrow 1$ **to** T **do**

Clustering step: turn off edges of \mathcal{G} with probability $\exp(-\beta_{ij} \mathbf{1}\{x_i^{(t)} = x_j^{(t)}\})$

Swapping step: assign a new state k to each connected component \mathcal{C} of $\Gamma(\mathcal{G}, \mathbf{x}^{(t)})$ with probability $\mathbf{P}(\mathbf{X}_{\mathcal{C}}^{(t)} = k) \propto \exp\{\sum_{i \in \mathcal{C}} \alpha_k\}$;

end

return the configuration $\mathbf{x}^{(T)}$

small enough value of β (Cooper and Frieze, 1999). This was proved independently by Huber (2003) who also derive a diagnostic tool for the convergence of the algorithm to its invariant distribution, namely using a coupling from the past procedure.

The algorithm can be extended to other Markov random field or models (e.g., Edwards and Sokal, 1988, Wolff, 1989, Higdon, 1998, Barbu and Zhu, 2005) but is then not necessarily efficient. In particular, it is not well suited for latent process. The bound to the data corresponds to a non-homogeneous external field that slows down the computation since the clustering step does not make a use of the data. A solution that might be effective is the partial decoupling of Higdon (1993, 1998). More recently, Barbu and Zhu (2005) make a move from the data augmentation interpretation to a Metropolis-Hastings perspective in order to generalize the algorithm to arbitrary probabilities on graphs. Nevertheless, it is not straightforward to bound the Markov chain of such modifications and mixing properties are still an open question despite good results in numerical experiments.

6.2 Monte Carlo maximum likelihood estimator

The use of Monte-Carlo techniques in preference to pseudolikelihood to compute maximum likelihood estimates has been especially highlighted by [Geyer and Thompson \(1992\)](#). Let assume $V_{\boldsymbol{\theta}, \mathcal{G}}$ is continuously differentiable on Θ . The gradient of the log-density (2.2) can be written as

$$\nabla_{\boldsymbol{\theta}} \log f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) = -\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x}) - \nabla_{\boldsymbol{\theta}} \log Z(\boldsymbol{\theta}, \mathcal{G}).$$

Despite closed form is out of reach, forward-simulations from the likelihood taken at each leapfrog step can be used to provide a Monte Carlo estimate of the gradient, using the following identity,

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \log Z(\boldsymbol{\theta}, \mathcal{G}) &= \frac{1}{Z(\boldsymbol{\theta}, \mathcal{G})} \nabla_{\boldsymbol{\theta}} Z(\boldsymbol{\theta}, \mathcal{G}) \\ &= \frac{1}{Z(\boldsymbol{\theta}, \mathcal{G})} \nabla_{\boldsymbol{\theta}} \int_{\mathcal{X}} \exp \{V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x})\} \mu(d\mathbf{x}) \\ &= \int_{\mathcal{X}} \nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x}) \frac{\exp \{V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x})\}}{Z(\boldsymbol{\theta}, \mathcal{G})} \mu(d\mathbf{x}) \\ &= \mathbf{E}_{\boldsymbol{\theta}} \{ \nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{X}) \}. \end{aligned} \tag{6.1}$$

So far, we have only assumed that $V_{\boldsymbol{\theta}, \mathcal{G}}$ is continuously differentiable on Θ . However this identity holds under regularity conditions which allow one to switch the derivative and integral operators (the domain \mathcal{X} of \mathbf{X} is assumed to be independent of $\boldsymbol{\theta}$) and under the assumption that $\nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{X})$ is integrable with respect to $f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})\mu(d\mathbf{x})$. When the function $V_{\boldsymbol{\theta}}$ linearly depends on the vector of parameters $\boldsymbol{\theta}$, that is

$$V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{x}) = -\boldsymbol{\theta}^T \mathbf{S}(\mathbf{x}),$$

where $\mathbf{S}(\mathbf{x}) = (s_1(\mathbf{x}), \dots, s_d(\mathbf{x}))$ is a vector of sufficient statistics, expectation (6.1) is simply the first moment of the statistics with respect to $f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})\mu(d\mathbf{x})$. It is also possible to show that the log-density is concave as its Hessian matrix is the second moment of the statistics with respect to $f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})\mu(d\mathbf{x})$:

$$\nabla_{\boldsymbol{\theta}}^2 \log f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) = -\mathbf{Cov}_{\boldsymbol{\theta}} \{ \mathbf{S}(\mathbf{X}) \}.$$

The maximum likelihood estimator $\hat{\boldsymbol{\theta}}_{\text{MLE}}$ is then the unique zero of the score function $\nabla_{\boldsymbol{\theta}} \log f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})$ and it satisfies

$$\mathbf{S}(\mathbf{x}) - \mathbf{E}_{\hat{\boldsymbol{\theta}}_{\text{MLE}}} \{ \mathbf{S}(\mathbf{X}) \} = 0. \tag{6.2}$$

Hence a solution to solve problem (3.1) is to resort to stochastic approximations on the basis of equation (6.2) (*e.g.*, [Younes, 1988](#), [Descombes et al., 1999](#)). [Younes \(1988\)](#) sets a stochastic gradient algorithm converging under mild conditions. At each iteration t , the algorithm takes the direction of the gradient estimated by the value of the statistic function for one realisation of the random field $\mathbf{x}^{(t)}$, namely

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} + \frac{\delta}{t+1} \left\{ \mathbf{S}(\mathbf{x}) - \mathbf{S}(\mathbf{x}^{(t)}) \right\},$$

where δ is a user-defined threshold. [Younes \(1988\)](#) provides theoretical conditions on δ which guarantee the convergence of the algorithm. However such theoretical value for δ yields in practice a step size too small to ensure the convergence to be achieved in reasonable amount of time. This can be overcome by controlling the probability of non-convergence, see the original paper by [Younes \(1988\)](#) for discussion. Another approach to compute the maximum likelihood estimation is to use direct Monte Carlo calculation of the likelihood such as the MCMC algorithm of [Geyer and Thompson \(1992\)](#). The convergence in probability of the latter toward the maximum likelihood estimator is proven for a wide range of models including Markov random fields. Following that work, [Descombes et al. \(1999\)](#) derive also a stochastic algorithm that, as opposed to [Younes \(1988\)](#), takes into account the distance to the maximum likelihood estimator using importance sampling.

6.3 Computing posterior distributions

Markov Chain Monte Carlo allow to asymptotically sample from analytically intractable posterior distribution π . It provides a very general framework to allow estimation of functionals of the form

$$\int_{\Theta} g(\boldsymbol{\theta})\pi(d\boldsymbol{\theta}),$$

for some function g by generating a Markov chain $(\boldsymbol{\theta}_n)_{n \in \mathbb{N}}$ with transition kernel P which leaves π invariant. The empirical distribution so obtained leads to the following approximation

$$\int_{\Theta} g(\boldsymbol{\theta})\pi(d\boldsymbol{\theta}) \approx \frac{1}{N} \sum_{n=1}^N g(\boldsymbol{\theta}_i).$$

While some Bayesian estimators can be efficiently estimated with such methods via the empirical distribution, the intractability of the likelihood model in (2.2) implies, in particular, that the standard MCMC toolbox is infeasible. Indeed, we are concerned with the situation where the unnormalised posterior distribution, the right-hand-side of (3.2) is intractable. This complication results in what is often termed a doubly-intractable posterior distribution, since the posterior distribution itself is normalised by the evidence (or marginal likelihood) which is typically also intractable. For instance, proposing to move from $(\boldsymbol{\theta})$ to $(\boldsymbol{\theta}')$ in a standard Metropolis-Hastings requires the computation of the unknown normalising constants, $Z(\boldsymbol{\theta}, \mathcal{G})$ and $Z(\boldsymbol{\theta}', \mathcal{G})$,

$$\rho(\boldsymbol{\theta}, \boldsymbol{\theta}') = 1 \wedge \frac{Z(\boldsymbol{\theta}, \mathcal{G})}{Z(\boldsymbol{\theta}', \mathcal{G})} \frac{q(\mathbf{x} | \boldsymbol{\theta}', \mathcal{G}) \nu(\boldsymbol{\theta} | \boldsymbol{\theta}') \pi(\boldsymbol{\theta}')}{q(\mathbf{x} | \boldsymbol{\theta}, \mathcal{G}) \nu(\boldsymbol{\theta}' | \boldsymbol{\theta}) \pi(\boldsymbol{\theta})},$$

where $\nu(\boldsymbol{\theta}' | \boldsymbol{\theta})$ denotes the proposal distribution to move from $\boldsymbol{\theta}$ to $\boldsymbol{\theta}'$. A solution, while being time consuming, is to estimate the ratio of the partition functions using path sampling ([Gelman and Meng, 1998](#)). Starting from equation (6.1), the path sampling identity writes as

$$\log \left\{ \frac{Z(\boldsymbol{\theta}_0, \mathcal{G})}{Z(\boldsymbol{\theta}_1, \mathcal{G})} \right\} = \int_{\boldsymbol{\theta}_0}^{\boldsymbol{\theta}_1} \mathbf{E}_{\boldsymbol{\theta}} \{ \nabla_{\boldsymbol{\theta}} V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{X}) \} d\boldsymbol{\theta}.$$

Hence the ratio of the two normalizing constants can be evaluated with numerical integration. For practical purpose, this approach can barely be recommended within a Metropolis-Hastings scheme since each iteration would require to compute a new ratio.

The exchange algorithm (Murray et al., 2006) is a popular MCMC methods to allow sampling from doubly-intractable distributions. The exchange algorithm samples from an augmented distribution

$$\pi(\boldsymbol{\theta}', \boldsymbol{\theta}, \mathbf{u} \mid \mathbf{x}) \propto f(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G})\pi(\boldsymbol{\theta})\nu(\boldsymbol{\theta}' \mid \boldsymbol{\theta})f(\mathbf{u} \mid \boldsymbol{\theta}', \mathcal{G}).$$

whose marginal distribution in $\boldsymbol{\theta}$ is the posterior distribution of interest. It extends an idea introduced by Møller et al. (2006). The proposal of Møller et al. (2006) consists in including an auxiliary variable \mathbf{U} whose density is the intractable likelihood itself. It follows a method based on single point importance sampling approximations of the partition functions $Z(\boldsymbol{\theta}, \mathcal{G})$ and $Z(\boldsymbol{\theta}', \mathcal{G})$. Murray et al. (2006) develop this work further by directly estimating the ratio $Z(\boldsymbol{\theta}, \mathcal{G})/Z(\boldsymbol{\theta}', \mathcal{G})$ instead of using previous single point estimates. This leads to a clever algorithm to sample from the above augmented distribution, where it turns out that the ratio of intractable normalising constants drops out of the acceptance probability

$$\rho(\boldsymbol{\theta}, \boldsymbol{\theta}', \mathbf{u}) = 1 \wedge \frac{\cancel{Z(\boldsymbol{\theta}, \mathcal{G})} \cancel{Z(\boldsymbol{\theta}', \mathcal{G})} q(\mathbf{u} \mid \boldsymbol{\theta}, \mathcal{G}) q(\mathbf{x} \mid \boldsymbol{\theta}', \mathcal{G}) \nu(\boldsymbol{\theta} \mid \boldsymbol{\theta}') \pi(\boldsymbol{\theta}')}{\cancel{Z(\boldsymbol{\theta}', \mathcal{G})} q(\mathbf{u} \mid \boldsymbol{\theta}', \mathcal{G}) \cancel{Z(\boldsymbol{\theta}, \mathcal{G})} q(\mathbf{x} \mid \boldsymbol{\theta}, \mathcal{G}) \nu(\boldsymbol{\theta}' \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})}$$

Murray et al. (2006) point out that the fraction $q(\mathbf{u} \mid \boldsymbol{\theta}, \mathcal{G})/q(\mathbf{u} \mid \boldsymbol{\theta}', \mathcal{G})$ which appears above, can be considered as an single sample importance estimator of $Z(\boldsymbol{\theta}, \mathcal{G})/Z(\boldsymbol{\theta}', \mathcal{G})$ since it holds that

$$\mathbf{E}_{\boldsymbol{\theta}'} \left\{ \frac{q(\mathbf{U} \mid \boldsymbol{\theta}, \mathcal{G})}{q(\mathbf{U} \mid \boldsymbol{\theta}', \mathcal{G})} \right\} = \frac{Z(\boldsymbol{\theta})}{Z(\boldsymbol{\theta}')}, \quad (6.3)$$

where $\mathbf{E}_{\boldsymbol{\theta}'}$ is the expectation with respect to $\mathbf{U} \sim f(\cdot \mid \boldsymbol{\theta}', \mathcal{G})$. In fact Alquier et al. (2016), consider a generalised exchange algorithm based, at each step of the algorithm, on an improved unbiased estimate of $Z(\boldsymbol{\theta})/Z(\boldsymbol{\theta}')$ including multiple auxiliary draws with respect to the proposed parameter, namely,

$$\frac{\widehat{Z(\boldsymbol{\theta}, \mathcal{G})}}{Z(\boldsymbol{\theta}', \mathcal{G})} = \frac{1}{N} \sum_{n=1}^N \frac{q(\mathbf{u}^{(n)} \mid \boldsymbol{\theta}, \mathcal{G})}{q(\mathbf{u}^{(n)} \mid \boldsymbol{\theta}', \mathcal{G})}, \quad (6.4)$$

where the auxiliary variables $\{\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)}\}$ are drawn from $f(\cdot \mid \boldsymbol{\theta}', \mathcal{G})$. However this so-called noisy exchange algorithm no longer leaves the target distribution invariant, nevertheless it is possible to provide convergence guarantees that the resulting Markov chain is close in some sense to the target distribution. Following the argument by Everitt (2012), any algorithm producing an unbiased estimate of the normalizing constant can thus be used in place of the importance sampling approximation and will lead to a valid procedure. An alternative to previous methods presented but neglected so far in the literature is Russian Roulette sampling (Lyne et al., 2015) which can be used to get an unbiased estimate of $1/Z(\boldsymbol{\theta})$. The idea to apply MCMC

Algorithm 4: Exchange algorithm

Input: an initial guess $(\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}'^{(0)}, \mathbf{u}^{(0)})$, a number of iterations T **Output:** a sample drawn from the augmented distribution $\pi(\boldsymbol{\theta}, \boldsymbol{\theta}', \mathbf{u} \mid \mathbf{x})$ for $t \leftarrow 1$ to T do **draw** $\boldsymbol{\theta}'$ from $\nu(\cdot \mid \boldsymbol{\theta}^{(t-1)})$; **draw** \mathbf{u} from $f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})$; **compute** the Metropolis-Hastings acceptance ratio

$$\rho(\boldsymbol{\theta}^{(t-1)}, \boldsymbol{\theta}', \mathbf{u}) = 1 \wedge \frac{q(\mathbf{u} \mid \boldsymbol{\theta}^{(t-1)}, \mathcal{G})}{q(\mathbf{u} \mid \boldsymbol{\theta}', \mathcal{G})} \frac{q(\mathbf{x} \mid \boldsymbol{\theta}', \mathcal{G})\nu(\boldsymbol{\theta}^{(t-1)} \mid \boldsymbol{\theta}')\pi(\boldsymbol{\theta}')}{q(\mathbf{x} \mid \boldsymbol{\theta}^{(t-1)}, \mathcal{G})\nu(\boldsymbol{\theta}' \mid \boldsymbol{\theta}^{(t-1)})\pi(\boldsymbol{\theta})}$$

Exchange move: set $(\boldsymbol{\theta}^{(t)}, \boldsymbol{\theta}'^{(t)}, \mathbf{u}^{(t)}) = (\boldsymbol{\theta}', \boldsymbol{\theta}^{(t-1)}, \mathbf{u})$ with probability $\rho(\boldsymbol{\theta}^{(t-1)}, \boldsymbol{\theta}', \mathbf{u})$; **Otherwise** set $(\boldsymbol{\theta}^{(t)}, \boldsymbol{\theta}'^{(t)}, \mathbf{u}^{(t)}) = (\boldsymbol{\theta}^{(t-1)}, \boldsymbol{\theta}'^{(t-1)}, \mathbf{u}^{(t-1)})$;

end

return $\left\{ (\boldsymbol{\theta}^{(t)}, \boldsymbol{\theta}'^{(t)}, \mathbf{u}^{(t)}) \right\}_{t=1}^T$

methods to situation where the target distribution can be estimated without bias by using an auxiliary variable construction has appeared in the *generalized importance Metropolis-Hasting* of [Beaumont \(2003\)](#) and has then been extended by [Andrieu and Roberts \(2009\)](#). This brings another justification to the aforementioned methods and possible improvement from sequential Monte Carlo literature (*e.g.*, [Andrieu et al., 2010](#)).

While the motivation is quite similar, the exchange algorithm is more convenient to implement whilst outperforming the single auxiliary variable method (SAVM) of [Møller et al. \(2006\)](#). Indeed, SAVM requires to choose the conditional distribution $q(\mathbf{u} \mid \mathbf{x}, \boldsymbol{\theta})$ which makes it difficult to calibrate. For instance, [Cucala et al. \(2009\)](#) stress out that a suitable choice is paramount and may significantly affect the performances of the algorithm.

A practical difficulty of implementing the exchange algorithm is the requirement to sample \mathbf{u} from Gibbs distribution $f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})$ to guarantee a valid MCMC scheme. In all generality, it is impossible or prohibitively expensive to carry out such perfect sampling. [Everitt \(2012\)](#) has provided convergence results when one uses instead the final draw from a Gibbs sampler with stationary distribution $f(\cdot \mid \boldsymbol{\theta}, \mathcal{G})$ as an approximate realisation. [Everitt \(2012\)](#) has notably pointed out that solely few iterations of the sampler are necessary. This approach has shown good performances in practice (*e.g.*, [Cucala et al., 2009](#), [Caimo and Friel, 2011](#)).

While the application of the exchange algorithm is straightforward for Bayesian parameter inference of a fully observed Markov random field, some work have been devoted to the use of the exchange algorithm for hidden Markov random fields such as the exchange marginal particle MCMC of [Everitt \(2012\)](#) or the estimation procedure in [Cucala and Marin \(2013\)](#). Though these methods produce accurate results they inherit the drawback of the exchange algorithm.

6.4 ABC model selection

Approximate Bayesian computation (ABC) has generated much activity in the literature recently as it offers a way to circumvent the difficulties of models which are intractable but can be simulated from. We refer the reader to [Marin et al. \(2012\)](#) and the references therein for a comprehensive overview on the method. When performing parameter estimation, the method is particularly well suited for problems where the likelihood function does not admit an algebraic form, a situation where MCMC methods are at a loss. We believe that the benefit of ABC for parameter estimation of a Gibbs random field is questionable. Hence we solely focus on model selection in this part.

To approximate \hat{m}_{MAP} , ABC starts by simulating numerous triplets $(m, \boldsymbol{\theta}_m, \mathbf{y})$ from the joint Bayesian model, see Algorithm 5. Afterwards, the algorithm mimics the Bayes classifier (3.4): it approximates the posterior probabilities by the frequency of each model number associated with simulated \mathbf{y} 's in a neighbourhood of \mathbf{y}^{obs} . In a metric space (\mathcal{Y}, ρ) , this neighbourhood is defined by the ball $\mathcal{B}(\epsilon, \mathbf{y}^{\text{obs}})$ of radius ϵ centered at \mathbf{y}^{obs} . If required, we can eventually predict the best model with the most frequent model in the neighbourhood, or, in other words, take the final decision by plugging in (3.4) the approximations of the posterior probabilities. However such a naive implementation is typically infeasible as the data usually lies in a space of high dimension and the algorithm faces the curse of dimensionality. Put in other words, sample dataset in the neighbourhood of \mathbf{y} occurs with a prohibitively low probability. The ABC algorithm performs therefore a (non linear) projection of the observed and simulated datasets onto some Euclidean space of reasonable dimension d via a function s , composed of summary statistics $\mathbf{S}(\cdot) = \{s_1(\cdot), \dots, s_M(\cdot)\}$ that are the concatenation of the summary statistics of each model with cancellation of possible replicates. The neighbourhood of \mathbf{y}^{obs} is thus defined as simulations whose distances to the observation measured in terms of summary statistics, *i.e.*, $\rho\{\mathbf{S}(\mathbf{y}), \mathbf{S}(\mathbf{y}^{\text{obs}})\}$, fall below a threshold ϵ . The accepted particles $(m^{(t)}, \mathbf{y}^{(t)})$ at the end of Algorithm 5 are distributed according to the pseudo-target $\pi(m \mid \rho(\mathbf{S}(\mathbf{y}), \mathbf{S}(\mathbf{y}^{\text{obs}})) \leq \epsilon)$ and the estimate of the posterior model distribution is given by

$$\hat{\pi}_\epsilon(m \mid \mathbf{y}^{\text{obs}}) = \frac{\sum \mathbf{1}\{m^{(t)} = m, \rho(\mathbf{S}(\mathbf{y}^{(t)}), \mathbf{S}(\mathbf{y}^{\text{obs}})) \leq \epsilon\}}{\sum \mathbf{1}\{\rho(\mathbf{S}(\mathbf{y}^{(t)}), \mathbf{S}(\mathbf{y}^{\text{obs}})) \leq \epsilon\}}.$$

ABC hence presents two level of approximations arising from the size of the neighbourhood ϵ and the introduction of summary statistics.

The choice of such summary statistics presents major difficulties that have been especially highlighted for model choice ([Robert et al., 2011](#), [Didelot et al., 2011](#)). When the summary statistics are not sufficient for the model choice problem, [Didelot et al. \(2011\)](#) and [Robert et al. \(2011\)](#) found that the above probability can greatly differ from the genuine $\pi(m \mid \mathbf{y}^{\text{obs}})$. Model selection between fully observed Markov random fields whose energy function is of the form $V_{\boldsymbol{\theta}, \mathcal{G}}(\mathbf{y}) = \boldsymbol{\theta}^T s(\mathbf{y})$ is a surprising example for which ABC is consistent. Indeed [Grelaud et al. \(2009\)](#) have pointed out that vector of summary statistics $\mathbf{S}(\cdot) = \{s_1(\cdot), \dots, s_M(\cdot)\}$ is sufficient for each model but also for the joint parameter across models $(\mathcal{M}, \theta_1, \dots, \theta_M)$. This allows to sample exactly from the posterior model distribution when $\epsilon = 0$. However the fact that the concatenated statistic inherits the sufficiency property from the sufficient statistics of each model is specific to exponential families ([Didelot et al., 2011](#)). When dealing with model choice between hidden Markov random fields, we fall outside of the exponential families due to the bound to the data. Thus we face the major difficulty outlined by [Robert et al. \(2011\)](#): it is almost impossible to build a sufficient statistic of reasonable dimension, *i.e.*, of dimension much lower than the

Algorithm 5: ABC model choice algorithm

Input: an observation \mathbf{y}^{obs} , summary statistics \mathbf{S} , a number of iterations T , an empirical quantile of the distance T_ϵ

Output: a sample from the approximated target of $\pi_\epsilon(\cdot \mid \mathbf{S}(\mathbf{y}^{\text{obs}}))$

for $t \leftarrow 1$ **to** T **do**

draw m from π ;

draw ψ from π_m ;

draw \mathbf{y} from $f_m(\cdot \mid \psi)$;

compute $\mathbf{S}(\mathbf{y})$;

save $\{m^{(t)}, \theta^{(t)}, \mathbf{S}(\mathbf{y}^{(t)})\} \leftarrow \{m, \psi, \mathbf{S}(\mathbf{y})\}$;

end

sort the replicates according to the distance $\rho(\mathbf{S}(\mathbf{y}^{(t)}), \mathbf{S}(\mathbf{y}^{\text{obs}}))$;

keep the T_ϵ first replicates;

return the sample of accepted particles

dimension of \mathcal{X} .

Beyond the seldom situations where sufficient statistics exist and are explicitly known, [Marin et al. \(2014\)](#) provide conditions which ensure the consistency of ABC model choice but the latter are difficult, if not impossible, to check in practice. To answer the absence of sufficient statistics as well as aforementioned theoretical conditions, very few has been accomplished in the context of ABC model choice. One solution would be to rely on the approach of [Prangle et al. \(2014\)](#). The statistics $\mathbf{S}(y)$ reconstructed by [Prangle et al. \(2014\)](#) have good theoretical properties (those are the posterior probabilities of the models in competition). Nevertheless the methods requires a pilot ABC run which is time consuming and can lead to poor approximations. Alternatively, [Stoehr et al. \(2015\)](#) have proposed to use summary statistics based on connected components of induced graphs in order to select between dependency structures of hidden Markov random fields. Beside, the specific result on Markov random fields, they derive an adaptive scheme to select the most informative set of statistics based on a local error rate. The main drawback of the method is that its use is limited to low dimensional vector of statistics. To overcome this issue, the summary statistics proposed by [Stoehr et al. \(2015\)](#) could be used among others within the ABC random forest procedure of [Pudlo et al. \(2015\)](#).

References

- H. Akaike. Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory*, pages 267–281. Akademinai Kiado, 1973.
- M. Alfò, L. Nieddu, and D. Vicari. A finite mixture model for image segmentation. *Statistics and Computing*, 18(2):137–150, 2008.
- P. Alquier, N. Friel, R. Everitt, and A. Boland. Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, 26(1):29–47, 2016. ISSN 1573-1375. doi: 10.1007/s11222-014-9521-x.
- C. Andrieu and G. O. Roberts. The Pseudo-Marginal Approach for Efficient Monte Carlo Computations. *The Annals of Statistics*, 37(2):697–725, 2009.

- C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov Chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010.
- A. Barbu and S.-C. Zhu. Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(8):1239–1253, 2005.
- M. A. Beaumont. Estimation of Population Growth or Decline in Genetically Monitored Populations. *Genetics*, 164(3):1139–1160, 2003.
- J. E. Besag. Nearest-neighbour Systems and the Auto-Logistic Model for Binary Data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(1):75–83, 1972.
- J. E. Besag. Spatial Interaction and the Statistical Analysis of Lattice Systems (with Discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236, 1974.
- J. E. Besag. Statistical Analysis of Non-Lattice Data. *The Statistician*, 24:179–195, 1975.
- J. E. Besag. On the Statistical Analysis of Dirty Pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):259–302, 1986.
- J. E. Besag and P. J. Green. Spatial Statistics and Bayesian Computation. *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(1):25–37, 1993.
- J. E. Besag, J. York, and A. Mollié. Bayesian image restoration, with two applications in spatial statistics. *Annals of the institute of statistical mathematics*, 43(1):1–20, 1991.
- C. Biernacki, G. Celeux, and G. Govaert. Assessing a mixture model for clustering with the integrated completed likelihood. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(7):719–725, 2000.
- K. P. Burnham and D. R. Anderson. *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Science & Business Media, 2002.
- A. Caimo and N. Friel. Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55, 2011.
- G. Celeux, F. Forbes, and N. Peyrard. EM procedures using mean field-like approximations for Markov model-based image segmentation. *Pattern Recognition*, 36(1):131–144, 2003.
- B. Chalmond. An iterative Gibbsian technique for reconstruction of m-ary images. *Pattern Recognition*, 22(6):747–761, 1989.
- P. Clifford. Markov random fields in statistics. *Disorder in physical systems: A volume in honour of John M. Hammersley*, pages 19–32, 1990.
- C. Cooper and A. M. Frieze. Mixing properties of the Swendsen-Wang process on classes of graphs. *Random Structures and Algorithms*, 15(3-4):242–261, 1999.
- I. Csiszár and Z. Talata. Consistent Estimation of the Basic Neighborhood of Markov Random Fields. *The Annals of Statistics*, 34(1):123–145, 2006.
- I. Csiszár, P. C. Shields, et al. The consistency of the BIC Markov order estimator. *The Annals of Statistics*, 28(6):1601–1619, 2000.
- L. Cucala and J.-M. Marin. Bayesian Inference on a Mixture Model With Spatial Dependence. *Journal of Computational and Graphical Statistics*, 22(3):584–597, 2013.

- L. Cucala, J.-M. Marin, C. P. Robert, and D. M. Titterton. A Bayesian Reassessment of Nearest-Neighbor Classification. *Journal of the American Statistical Association*, 104(485): 263–273, 2009.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1): 1–38, 1977.
- X. Descombes, R. D. Morris, J. Zerubia, and M. Berthod. Estimation of Markov random field prior parameters using Markov chain Monte Carlo maximum likelihood. *Image Processing, IEEE Transactions on*, 8(7):954–963, 1999.
- X. Didelot, R. G. Everitt, A. M. Johansen, and D. J. Lawson. Likelihood-free estimation of model evidence. *Bayesian Analysis*, 6(1):49–76, 2011.
- R. G. Edwards and A. D. Sokal. Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm. *Physical Review D*, 38(6):2009, 1988.
- R. G. Everitt. Bayesian Parameter Estimation for Latent Markov Random Fields and Social Networks. *Journal of Computational and Graphical Statistics*, 21(4):940–960, 2012.
- F. Forbes and G. Fort. Combining Monte Carlo and Mean Field-Like Methods for Inference in Hidden Markov Random Fields. *Image Processing, IEEE Transactions on*, 16(3):824–837, 2007.
- F. Forbes and N. Peyrard. 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, 2003.
- O. François, S. Ancelet, and G. Guillot. Bayesian Clustering Using Hidden Markov Random Fields in Spatial Population Genetics. *Genetics*, 174(2):805–816, 2006.
- O. Frank and D. Strauss. Markov graphs. *Journal of the American Statistical Association*, 81(395):832–842, 1986.
- N. Friel and A. N. Pettitt. Likelihood Estimation and Inference for the Autologistic Model. *Journal of Computational and Graphical Statistics*, 13(1):232–246, 2004.
- N. Friel and H. Rue. Recursive computing and simulation-free inference for general factorizable models. *Biometrika*, 94(3):661–672, 2007.
- N. Friel, A. N. Pettitt, R. Reeves, and E. Wit. Bayesian Inference in Hidden Markov Random Fields for Binary Data Defined on Large Lattices. *Journal of Computational and Graphical Statistics*, 18(2):243–261, 2009.
- E. Gassiat. Likelihood ratio inequalities with applications to various mixtures. In *Annales de l’IHP Probabilités et statistiques*, volume 38, pages 897–906, 2002.
- A. Gelman and X.-L. Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, 13(2):163–185, 1998.
- S. Geman and D. Geman. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.
- S. Geman and C. Graffigne. Markov Random Field Image Models and Their Applications to Computer Vision. In *Proceedings of the International Congress of Mathematicians*, volume 1, pages 1496–1517, 1986.

- H. Georgii. *Gibbs Measures and Phase Transitions*. De Gruyter studies in mathematics. De Gruyter, 2011.
- C. J. Geyer and E. A. Thompson. Constrained Monte Carlo Maximum Likelihood for Dependent Data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 54(3):657–699, 1992.
- V. K. Gore and M. R. Jerrum. The Swendsen–Wang process does not always mix rapidly. *Journal of Statistical Physics*, 97(1-2):67–86, 1999.
- P. J. Green and S. Richardson. Hidden Markov Models and Disease Mapping. *Journal of the American Statistical Association*, 97(460):1055–1070, 2002.
- A. Grelaud, C. P. Robert, J.-M. Marin, F. Rodolphe, and J.-F. Taly. ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 4(2):317–336, 2009.
- G. R. Grimmett. A theorem about random fields. *Bulletin of the London Mathematical Society*, 5(1):81–84, 1973.
- W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- D. M. A. Haughton. On the Choice of a Model to Fit Data from an Exponential Family. *The Annals of Statistics*, 16(1):342–355, 1988.
- D. M. Higdon. Discussion on the Meeting on the Gibbs Sampler and Other Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society. Series B*, 55(1):78, 1993.
- D. M. Higdon. Auxiliary variable methods for Markov chain Monte Carlo with applications. *Journal of the American Statistical Association*, 93(442):585–595, 1998.
- M. Huber. A bounding chain for Swendsen-Wang. *Random Structures & Algorithms*, 22(1):43–59, 2003.
- E. Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift fur Physik*, 31:253–258, 1925.
- T. S. Jaakkola and M. I. Jordan. Bayesian parameter estimation via variational methods. *Statistics and Computing*, 10(1):25–37, 2000. ISSN 1573-1375. doi: 10.1023/A:1008932416310. URL <http://dx.doi.org/10.1023/A:1008932416310>.
- C. Ji and L. Seymour. A consistent model selection procedure for Markov random fields based on penalized pseudolikelihood. *The annals of applied probability*, pages 423–443, 1996.
- M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An Introduction to Variational Methods for Graphical Models. *Machine learning*, 37(2):183–233, 1999.
- R. E. Kass and A. E. Raftery. Bayes factors. *Journal of the american statistical association*, 90(430):773–795, 1995.
- C. Keribin. Consistent Estimation of the Order of Mixture Models. *Sankhy: The Indian Journal of Statistics, Series A (1961-2002)*, 62(1):49–66, 2000.
- I. Lanford, O.E. and D. Ruelle. Observables at infinity and states with short range correlations in statistical mechanics. *Communications in Mathematical Physics*, 13(3):194–215, 1969.
- B. G. Lindsay. Composite likelihood methods. *Contemporary Mathematics*, 80(1):221–39, 1988.
- J. S. Liu. Peskun’s theorem and a modified discrete-state gibbs sampler. *Biometrika*, 83(3):681–682, 1996.

- A.-M. Lyne, M. Girolami, Y. Atchadé, H. Strathmann, and D. Simpson. On Russian Roulette Estimates for Bayesian Inference with Doubly-Intractable Likelihoods. *Statistical Science*, 30(4):443–467, 2015.
- J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate Bayesian Computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- J.-M. Marin, N. S. Pillai, C. P. Robert, and J. Rousseau. Relevant statistics for Bayesian model choice. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(5):833–859, 2014.
- V. Matveev and R. Shrock. Complex-temperature singularities in Potts models on the square lattice. *Physical Review E*, 54(6):6174, 1996.
- C. A. McGrory, D. M. Titterton, R. Reeves, and A. N. Pettitt. Variational Bayes for estimating the parameters of a hidden Potts model. *Statistics and Computing*, 19(3):329–340, 2009.
- N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The journal of Chemical Physics*, 21(6):1087–1092, 1953.
- A. Mira, J. Møller, and G. O. Roberts. Perfect slice samplers. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63(3):593–606, 2001.
- J. Møller, A. N. Pettitt, R. Reeves, and K. K. Berthelsen. An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika*, 93(2):451–458, 2006.
- M. T. Moores, C. E. Hargrave, F. Harden, and K. Mengersen. Segmentation of cone-beam CT using a hidden Markov random field with informative priors. *Journal of Physics : Conference Series*, 489, 2014.
- I. Murray, Z. Ghahramani, and D. J. C. MacKay. MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, pages 359–366. AUAI Press, 2006.
- R. D. Neidinger. Introduction to Automatic Differentiation and MATLAB Object-Oriented Programming. *SIAM Review*, 52(3):545–563, 2010.
- M. A. Newton and A. E. Raftery. Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 3–48, 1994.
- L. Onsager. Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition. *Physical Review*, 65:117–149, 1944.
- A. N. Pettitt, N. Friel, and R. W. Reeves. Efficient calculation of the normalizing constant of the autologistic and related models on the cylinder and lattice. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 65(1):235–246, 2003.
- R. B. Potts. Some generalized order-disorder transformations. In *Mathematical proceedings of the cambridge philosophical society*, volume 48, pages 106–109. Cambridge Univ Press, 1952.
- D. Prangle, P. Fearnhead, M. P. Cox, P. J. Biggs, and N. P. French. Semi-automatic selection of summary statistics for ABC model choice. *Statistical Applications in Genetics and Molecular Biology*, 13(1):67–82, 2014.

- J. G. Propp and D. B. Wilson. Exact Sampling with Coupled Markov chains and Applications to Statistical Mechanics. *Random structures and Algorithms*, 9(1-2):223–252, 1996.
- P. Pudlo, J.-M. Marin, A. Estoup, J.-M. Cornuet, M. Gauthier, and C. P. Robert. Reliable abc model choice via random forests. *Bioinformatics*, 2015.
- W. Qian and D. Titterton. 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, 1991.
- R. Reeves and A. N. Pettitt. Efficient recursions for general factorisable models. *Biometrika*, 91(3):751–757, 2004.
- C. P. Robert, J.-M. Cornuet, J.-M. Marin, and N. S. Pillai. Lack of confidence in approximate Bayesian computation model choice. *Proceedings of the National Academy of Sciences*, 108(37):15112–15117, 2011.
- G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph (p^*) models for social networks. *Social networks*, 29(2):173–191, 2007.
- H. Rue. Fast Sampling of Gaussian Markov Random Fields. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63(2):325–338, 2001.
- T. Rydén and D. Titterton. Computational Bayesian analysis of hidden Markov models. *Journal of Computational and Graphical Statistics*, 7(2):194–211, 1998.
- G. Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.
- L. Seymour and C. Ji. Approximate Bayes model selection procedures for Gibbs-Markov random fields. *Journal of Statistical Planning and Inference*, 51(1):75–97, 1996.
- D. C. Stanford and A. E. Raftery. Approximate Bayes factors for image segmentation: The pseudolikelihood information criterion (PLIC). *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(11):1517–1520, 2002.
- J. Stoehr, P. Pudlo, and L. Cucala. Adaptive ABC model choice and geometric summary statistics for hidden Gibbs random fields. *Statistics and Computing*, 25(1):129–141, 2015.
- J. Stoehr, J.-M. Marin, and P. Pudlo. Hidden Gibbs random fields model selection using Block Likelihood Information Criterion. *Stat*, 5(1):158–172, 2016.
- R. H. Swendsen and J.-S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Physical Review Letters*, 58(2):86–88, 1987.
- C. Varin, N. Reid, and D. Firth. An overview of composite likelihood methods. *Statistica Sinica*, 21(1):5–42, 2011.
- S. Wasserman and P. Pattison. Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p^* . *Psychometrika*, 61(3):401–425, 1996.
- G. C. G. Wei and M. A. Tanner. A Monte Carlo Implementation of the EM Algorithm and the Poor Man’s Data Augmentation Algorithms. *Journal of the American Statistical Association*, 85(411):699–704, 1990.
- U. Wolff. Collective Monte Carlo updating for spin systems. *Physical Review Letters*, 62(4):361, 1989.
- C. J. Wu. On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11(1):95–103, 1983.

- F.-Y. Wu. The Potts model. *Reviews of modern physics*, 54(1):235, 1982.
- L. Younes. Estimation and annealing for Gibbsian fields. *Annales de l'Institut Henri Poincaré*, 24:269–294, 1988.
- J. Zhang. The mean field theory in EM procedures for Markov random fields. *Signal Processing, IEEE Transactions on*, 40(10):2570–2583, 1992.