



The node-wise Pseudo-marginal method: model selection with spatial dependence on latent graphs

Denishrouf Thesingarajah¹ · Adam M. Johansen¹

Received: 17 September 2021 / Accepted: 30 April 2022
© The Author(s) 2022

Abstract

Motivated by problems from neuroimaging in which existing approaches make use of “mass univariate” analysis which neglects spatial structure entirely, but the full joint modelling of all quantities of interest is computationally infeasible, a novel method for incorporating spatial dependence within a (potentially large) family of model-selection problems is presented. Spatial dependence is encoded via a Markov random field model for which a variant of the pseudo-marginal Markov chain Monte Carlo algorithm is developed and extended by a further augmentation of the underlying state space. This approach allows the exploitation of existing unbiased marginal likelihood estimators used in settings in which spatial independence is normally assumed thereby facilitating the incorporation of spatial dependence using non-spatial estimates with minimal additional development effort. The proposed algorithm can be realistically used for analysis of moderately sized data sets such as 2D slices of whole 3D dynamic PET brain images or other regions of interest. Principled approximations of the proposed method, together with simple extensions based on the augmented spaces, are investigated and shown to provide similar results to the full pseudo-marginal method. Such approximations and extensions allow the improved performance obtained by incorporating spatial dependence to be obtained at negligible additional cost. An application to measured PET image data shows notable improvements in revealing underlying spatial structure when compared to current methods that assume spatial independence.

Keywords Bayesian inference · Monte Carlo · Positron emission tomography · Spatial modelling

1 Introduction

Incorporating spatial dependence in the process of analysing large image, and other spatial image-like, data sets can be a difficult problem largely due to the computational requirements. An important example of such data is provided by PET (Positron Emission Tomography) imaging of the brain, where a whole image typically requires analysis of up to 10^6 time series (Hammers et al. 2007). Current state of the art analysis, e.g. Fan et al. (2021); Zhou et al. (2016), of such data generally either assumes spatial independence of pixels or performs large-scale aggregation over space (Zhou et al. 2002) to overcome computational restrictions. In this paper,

a novel computational method to incorporate spatial dependence in the process of model selection is presented.

The technique presented here extends the well-known pseudo-marginal Markov Chain Monte Carlo (MCMC) algorithm originally presented by Beaumont (2003), and characterised by Andrieu and Roberts (2009). More specifically, a Potts model (Potts 1952) is first utilised as a prior distribution to encode spatial dependence over a *local* model associated with each node of a graph. Imposing a further assumption of conditional spatial independence for local model parameters given the local model then gives rise to significant simplifications in computations.

This Markov random field model is used on the discrete space of model orders to describe spatial dependence between neighbouring nodes (pixels). Finally, a standard component-wise MCMC updating scheme is used in conjunction with the above assumptions. This then allows for the use of node-level marginal likelihood estimators to be used within a pseudo-marginal MCMC method. The end result is a tractable, accessible computational method that uses model

✉ Denishrouf Thesingarajah
d.thesingarajah@warwick.ac.uk

Adam M. Johansen
a.m.johansen@warwick.ac.uk

¹ Department of Statistics, University of Warwick, Coventry CV4 7AL, UK

selection at the individual (node/pixel) level to perform spatially dependent model selection on the whole (graph/image) data set. This method provides a flexible but efficient algorithm that can be readily implemented—in fact, the proposed method allows for the use of existing unbiased non-spatial estimators to be used within a framework which incorporates spatial dependence.

In addition, the proposed method can be further extended through augmentation of the state space of the MCMC chain; in this specialised setting careful specification of the proposal distributions can lead to many further computational approaches and techniques.

The remainder of this paper is structured as follows. Section 2 presents relevant background material including the motivating problem and reviews the use of Potts models in related contexts. Section 3 reviews existing methods for sampling from distributions of this type. Section 4 presents a hierarchical model and develops methods for using it in the context of interest alongside a formal justification of the approach and a number of extensions. Finally, the performance of the developed method, and some variants thereof, is studied empirically in applied settings in Sect. 5.

2 Background

2.1 Preliminaries

A graph $G = (V, E)$ is the pair of sets of nodes, or vertices, denoted V and edges denoted E . In particular, each element of the edge-set E is some pair of elements of the nodes-set $u, v \in V$, denoted $\langle u, v \rangle \in E$. Two nodes u, v are said to be connected by the edge $\langle u, v \rangle$ if $\langle u, v \rangle \in E$. In this case we may say that u and v are neighbours, or adjacent, and denote this by the relation $u \sim v$. Here, we will look only at undirected graphs, so the relation \sim will be symmetric and $\langle u, v \rangle$ is an *unordered* pair. Denote by $\partial(v) = \{u \in V : v \sim u\}$ the set of neighbours of v , by convention v is not a neighbour of itself.

Given a graph, G , a collection of random variables, $\mathbf{X} = (X_v : v \in V)$, indexed by the nodes-set, V , is called a random field on G . Let P be the law of \mathbf{X} ; and define, for any $A \subset V$, $\mathbf{X}_A = \{X_v : v \in A\}$ and $\mathbf{X}_{-A} = \{X_v : v \in V \setminus A\}$. With a slight abuse of this notation, we will write \mathbf{X}_{-v} to denote $\mathbf{X}_{V \setminus \{v\}}$ for $v \in V$. \mathbf{X} is called a Markov random field (Besag 1974) on a discrete graph if and only if we have that:

$$P(X_v | \mathbf{X}_{-v}) = P(X_v | \mathbf{X}_{\partial(v)}),$$

where each component X_v takes value in some set \mathcal{X} . Thus, \mathbf{X} takes values in \mathcal{X}^V , the collection of maps from V to \mathcal{X} . In this paper, \mathcal{X} will always be finite.

Denote a parametric model for data $Y \in \mathcal{Y}$, by $S = (\mathcal{Y}, \{P_\theta : \theta \in \Theta\})$; where P_θ is a probability distribution

over \mathcal{Y} , with θ taking values in the parameter space $\Theta \subset \mathbb{R}^d$ with $d < \infty$. We will use $f(\cdot | \theta)$ to denote the density (w.r.t some dominating measure, usually the Lebesgue or counting measures) of the distribution P_θ . Allow $Y \sim P_\theta$ and, where the density exists, $Y \sim f(\cdot | \theta)$, to mean that Y is distributed according to the distribution P_θ .

A finite set of statistical models, called the model space, is denoted $\mathcal{S} = \{S_M : M \in \mathcal{M}\}$. Here, each model is indexed by the model order M , which is an element of a finite index set $\mathcal{M} \subset \mathbb{N}$. For a model S_M , of order M , let $f_M(\cdot | \theta)$ for parameter $\theta \in \Theta_M$, be the associated density. We write Θ_M to emphasis that the parametric space is often dependent on the model order. A straightforward example of such models are the compartmental models, see Gunn et al. (2001) or Sect. 2.2 below. Bayesian model selection, averaging and comparison involves inference of the model order $M \in \mathcal{M}$, where M is a random variable to which some prior distribution is attached (Robert 2007, Chapter 7). An important component of the Bayesian model selection process is the *posterior model probability*, denoted $\pi(M | Y)$. Specifically, a central objective here is to characterise this distribution in the setting for image and image-like data, while accounting for spatial dependence. The definition of the posterior distribution (particularly for the setting we are interested in) will be detailed in the sequel.

2.2 Position emission tomography

An application that motivates the methodological development that follows is Positron Emission Tomography (PET; Phelps et al. (2006)).

Dynamic PET is an increasingly important neuroimaging technique used to study different functions of the brain *in vivo*. PET images have been used to investigate many biochemical processes in the brain, in both control and atypical subjects. This invasive imaging modality uses the detection of high-energy photons, released as a consequence of an injected tracer's positron emissions, as a signal to form a dynamical three-dimensional image.

Given that there will be greater radioactivity in regions where there are higher amounts of the tracer, the signal intensity in each voxel will be proportional to the concentration of the tracer. Thus, by taking a sequence of measurements over time, the constructed image represents the time course of the tissue concentration of the tracer *in vivo*.

As a result of the nature and number of observations recorded in the reconstructed PET image, the data is usually modelled using a system of linear differential equations such as compartmental models (Gunn et al. 2001). We restrict our attention here to the plasma input compartment models. Given data $y = (y_1, \dots, y_k)^\top$, at a voxel, an M -compartment model can be formally written as:

$$y_j = C_T(t_j; \phi_{1:M}, \vartheta_{1:M}) + \sqrt{\frac{C_T(t_j; \phi_{1:M}, \vartheta_{1:M})}{t_j - t_{j-1}}} \epsilon_j, \quad (1)$$

$$C_T(t_j; \phi_{1:M}, \vartheta_{1:M}) = \sum_{i=1}^M \phi_i \int_0^{t_j} C_P(s) e^{-\vartheta_i(t_j-s)} ds, \quad (2)$$

for $j = 1, \dots, k$; where C_T is the tissue concentration, t_j is the measurement time of observation y_j , ϵ_j is the additive measurement error and the plasma input function C_P is treated as known. The parameters $\phi_{1:M} = (\phi_1, \dots, \phi_M)$ and $\vartheta_{1:M} = (\vartheta_1, \dots, \vartheta_M)$ determine the dynamics of the model. Importantly, these parameters are indistinguishable, (Gunn et al. 2001, Theorem 2.2), and are often called the micro-parameters. In contrast, the macro-parameter

$$V_D = \sum_{i=1}^M \frac{\phi_i}{\vartheta_i},$$

termed the volume of distribution is uniquely identifiable. This macro-parameter is very often the quantity of principal interest in statistical analysis of PET data.

The additive error variable ϵ_j is typically modelled using a normal distribution:

$$\epsilon_j \sim \mathcal{N}(0, \sigma^2),$$

where $\mathcal{N}(0, \sigma^2)$ is the normal distribution with mean zero and variance σ^2 . This class of statistical models have proven to be very effective in Bayesian analysis of PET data, as studied and presented by Peng et al. (2008) and Jiang et al. (2009).

One ubiquitous method for PET analysis is the non-negative non-linear least squares (NNLS; Cunningham and Jones (1993)). This approach is often used due to the non-negative nature of the rate constants. The NNLS method involves minimal modelling assumptions, imposing no prior assumptions on the number of components. Instead, the time-series at each voxel is interpreted as a noisy measurement of exponential decay. The method involves the formulation of a constrained linear optimisation problem, where: M will represent the maximum number of terms to be included, for example (Cunningham and Jones 1993) use $M = 100$; Each component of the ϑ parameter is fixed and predetermined to lie within some range that is physiologically meaningful. Next, the optimal value of ϕ_j for each ϑ_j is estimated using a numerical method, subject to $\phi_i \geq 0, j = 1, \dots, m$.

Recently, computer-intensive Monte Carlo methods have also been successfully applied to these models towards meaningful statistical inference of PET data. Specifically, Zhou et al. (2013) investigated an application specific MCMC approach to Bayesian model comparison for this class of models; and studied both vague non-informative and biolog-

ically informed priors. In this work, samples from a MCMC chain were used in the generalized harmonic mean estimator to approximate the model evidence in order to perform model comparison, selection and averaging. This work also demonstrated the possibility that Monte Carlo approaches could be used to investigate and meaningfully compare more complex models, containing up to $M = 3$ compartments, rather than the typical up to $M = 2$ compartmental models. Zhou et al. (2013) further showed that within a Bayesian framework, a t -distributed error structure may be far more plausible. This can be written as

$$\epsilon_j \sim \mathcal{T}(0, \tau, \nu),$$

where $\mathcal{T}(0, \tau, \nu)$ denotes the Student's t -distribution with mean zero, scale τ , and ν degrees of freedom. Expressions for prior and posterior densities, of both cases, can be found in Appendix A.

Importantly, with regard to model selection within this context, Zhou et al. (2013) also showed that a Bayesian approach *exhibited some spatial structure* despite assuming spatial independence. This was in contrast to using the Akaike Information Criteria (AIC; Akaike (1973)) for model selection with a non-linear least square (NLS) method (as well as the NNLS method) used to approximate the MLE which showed no obvious spatial structure, see also Sect. 5.3.2.

More recently, Zhou et al. (2016) used an adaptive variant of the sequential Monte Carlo (SMC) sampler algorithm of Del Moral et al. (2006) to estimate the model evidence (marginal likelihood). The class of algorithms presented in this work, use adaptive MCMC kernels together with adaptive annealing schemes to minimise tuning requirements to further the accessibility of Monte Carlo approaches to this problem. In addition, the model evidence can be directly approximated using the set of weights from the sampler. The computational method proposed in this present paper, builds upon this approach and thus inherits many of these advantages. Further, Zhou et al. (2016) presented empirical results suggesting that this SMC method produces similar performance to the MCMC method proposed by Zhou et al. (2013); As such, in this paper, we focus on comparing performance to the SMC sampler method—which we refer to as the (spatially) independent SMC method.

Similarly, Castellaro et al. (2017) propose a method based on a Variational Bayes (VB) approach for parameter estimation in PET data analysis; and also present an empirical comparison to the MCMC method of Zhou et al. (2013). Specifically, in this empirical comparison the $M = 2$ -compartment model was used to represent measured PET data. The VB method was adapted to kinetic modelling of PET data, applied to the $M = 2$ -compartment model and used to estimate the V_D . Next, the MCMC method, restricted

to inference from the $M = 2$ -compartment model, was also used to analyse this data. Castellaro et al. concluded that there was a strong agreement between the outputs of these two methods. In particular the V_D maps of both the VB and MCMC methods showed a very high correlation (Pearson $R^2 = 0.99$ (see, Figure S2, Castellaro et al. (2017))).

There continues to be further development in Bayesian approaches in analysis of PET data, albeit in almost all cases they are non-spatial approaches. The most recent example is Fan et al. (2021), who propose a simple and intuitive algorithm, based on Approximate Bayesian Computation (ABC), for analysis of PET data. As with the above Bayesian approaches, this method, termed PET-ABC, assumes voxels are spatially independent. PET-ABC works by simulating from the parameter space using a simple rejection scheme: Firstly, proposals are sampled from the prior distribution, typically this is the uniform density with a physiologically meaningful range. Each proposal is used as a trial value to estimate the signal C_T . Next, the error between a summary statistics of the estimated signal and the de-noised data y is computed. In their study, Fan et al. (2021) suggest using spline smoothed estimates of C_T as the summary statistic for both the estimated and observed signal. The proposed value of the parameter is accepted if the above error is below a pre-determined threshold. The generated parameter sample can then be used for point estimation and to quantify uncertainty. Bayesian model selection follows the above method, with the additional step of proposing a model at each iteration. Fan et al. (2021) state that using PET-ABC produced estimates with lower variance when compared to the weighted NLS method.

2.3 Potts model

The Potts model (Potts 1952), a generalisation of the Ising model (Ising 1925), was used originally to model interacting spins on a lattice. However, these models have been shown to be also very effective in analysis of image data; for a detailed discussion and review of such applications see Winkler (1995), Geman (1990) and references therein. For example, Geman and Geman (1984) is an early study demonstrating the effectiveness of the Ising model for restoration of images under a Bayesian framework. More recently, the Potts model has been used very successfully within the more broader but still growing sub-field of Bayesian image analysis Hurn et al. (2003).

In the present context, the Potts model is an ideal prior for model orders, due to its discrete state space, minimal parametrisation and ability to encode the general principle that nearby vertices are *a priori* likely to be best described by the same model. Following the Bayesian approach, in this study, the model order of the data will be Markov random field with a Potts prior distribution.

For a review of Potts distributions, particularly in the image analysis context, see Hurn et al. (2003).

Given a graph $G = (V, E)$, finite state space \mathcal{X} and coupling constant $J > 0$, the Potts model specifies a family of parametric probability distributions on \mathcal{X}^V ; characterised by the joint mass function:

$$p(\mathbf{x}|J, G) = \frac{1}{\zeta(J)} \exp\left(J \sum_{v \sim u} \delta_{x_v, x_u}\right).$$

Here, recall $v \sim u$ denotes neighbouring pairs and δ_{x_v, x_u} is the Kronecker delta, i.e. δ_{x_v, x_u} is one if $x_v = x_u$, and zero otherwise. The intractable normalising constant (or partition function), written

$$\zeta(J) = \sum_{\mathbf{x}' \in \mathcal{X}^V} \exp\left(J \sum_{v \sim u} \delta_{x'_v, x'_u}\right),$$

is a function of J .

The parameter J dictates how likely neighbouring random variable X_v and X_u are to have the same value. Given that the normalisation depends upon this parameter, J is typically very difficult to infer; and will be treated as known in this setting as discussed at the end of this section.

Given some graph $G = (V, E)$, and the associated random variable $\mathbf{X} = (X_v \in \mathcal{X} : v \in V)$, we write

$$\mathbf{X} \sim \text{Potts}(J, G, \mathcal{X})$$

to mean that the random variable \mathbf{X} , taking values in state space \mathcal{X}^V , is distributed according to the Potts model with given coupling constant J . Where no ambiguity arises, the parameters will be omitted from notation in the interests of clarity.

The Potts model exhibits phase transition behaviour. This has significant implications when using single-site update scheme MCMC methods, which could result in slow mixing. That is, if the chain is in a configuration such as the case described above, changes to single variables will mean proposing to move to states of lower probability. Furthermore, phase transition behaviour happens close to or higher than critical values of the coupling constant, which we denote as J_{critical} . In the Monte Carlo setting, there is a sharp transition from fast mixing below J_{critical} and slow mixing above it. Onsager (1944) showed that for the Ising model on the two dimensional first order square lattice the exact value was

$$J_{\text{critical}} = \log\left(1 + \sqrt{2}\right) \approx 0.881.$$

More recently, see for instance Matveev and Shrock (1996) or Wu (1982), this has been extended: Let $D = |\mathcal{X}|$, then for a D -state Potts model,

$$J_{\text{critical}} = \log(1 + \sqrt{D}).$$

In applications, since the intractable normalising constant $\zeta(J)$ is dependent on J , it is difficult to infer in the model fitting process (Everitt 2012). Within the Monte Carlo context, these types of distributions are referred to as *doubly-intractable*. Møller et al. (2006) introduced an ingenious method to address such problems for the Ising model. For a more recent work, and for a concise but detailed review of developments since Møller et al. (2006), see Moores et al. (2020) and references therein.

Additionally, J is a parameter of a prior distribution, thus under the Bayesian analysis framework it would not be too unreasonable for it to be pre-specified. As such, in what follows, we will treat J as known. For example, in the numerical studies below we use a value which gives good performance in all our pilot studies, with a preference for choosing a smaller value of J to avoid imposing too much spatial structure via this prior distribution.

3 Existing Monte Carlo methods

Having established the problem of interest and reviewed the relevant distribution for encoding spatial information, we now turn our attention to a brief review of existing computational methods for investigating the characteristics of such complex distributions. We begin by looking at the relatively simple Gibbs sampling method for the Potts model. However, ultimately we would like to sample from, and subsequently characterise, a posterior distribution that consists of the Potts model as a prior, together with an intractable likelihood (detailed in Sect. 4). To this end, a common and standard approach to sampling from densities expressed in terms of intractable integrals, known as the pseudo-marginal method, is also discussed. In this paper, we use the SMC normalising constant estimator, to approximate said intractable integrals. Thus, the section ends with a brief discussion on the SMC estimator and its properties. However, generic pseudo-marginal methods would be computationally infeasible in the context of high-dimensional discrete latent spaces of the sort considered here; this then leads to the proposal of a novel computational method in the subsequent section.

3.1 Gibbs sampler for the Potts model

There exist efficient samplers for the Potts model, most notably Glauber Dynamics Glauber (1963) corresponding to single-site Gibbs updates, and the elegant Swendsen-Wang algorithm (Swendsen and Wang 1987). Typically, the Swendsen-Wang algorithm is used in simple Potts models, where there is an absence of a “strong external field” i.e. the (graph) likelihood in this context. In this regime, the Swendsen-Wang algorithm has appealing convergence prop-

erties for all values of the coupling strength, J , see Hurn et al. (2003). However, this excellent performance deteriorates markedly in the presence of an external field (Higdon 1998) and hence is less appealing in the presence of an informative likelihood. In settings in which the likelihood is relatively uninformative, there might be benefits in combining this type of cluster update with the local moves explored here.

We focus on the Gibbs sampling approach (Glauber 1963; Geman and Geman 1984), variants of which will extend naturally to the context of interest in this paper. The full conditionals required for a Gibbs sampler are readily computed for the Potts model:

$$p(x_v | \mathbf{x}_{-v}) = \frac{\exp\left(J \sum_{u \in \partial(v)} \delta_{x_u, x_v}\right)}{\sum_{x' \in \mathcal{X}} \exp\left(J \sum_{u \in \partial(v)} \delta_{x_u, x'}\right)}.$$

As is typical in Gibbs samplers, the use of full conditionals results in a simple node-wise update schedule.

However, computing full conditionals for posterior densities such as that of interest here is not possible because the likelihood cannot be evaluated even up to a normalising constant. So we turn to another Monte Carlo method to tackle this problem.

3.2 Pseudo-marginal Monte Carlo Markov chain

Consider, first, MCMC in the simpler case of a tractable target density; one which can be evaluated point-wise.

Denote the target probability density $\mu(x)$, where $x \in \mathcal{X}$, and $\mathcal{X} \subseteq \mathbb{R}^d$ is the state space. MCMC algorithms construct a μ -invariant Markov chain denoted $(X^{(i)})_{i=1}^n$ and like all time-homogeneous Markov chains are completely characterised by their initial distribution and their transition kernel.

The Metropolis-Hastings (MH) (Metropolis et al. 1953; Hastings 1970) algorithm is a particularly widely-used MCMC algorithm that may be used to target any analytically tractable density μ . The algorithm can be described very briefly as follows:

Given that the MH Markov chain is at some state $x \in \mathcal{X}$, a new state $x^* \in \mathcal{X}$ is proposed from a proposal density, denoted here $q(x, \cdot)$. The proposed state x^* is accepted with probability $\min\{1, R(x, x^*)\}$, where the acceptance ratio is defined

$$R(x, x^*) := \frac{\mu(x^*)q(x^*, x)}{\mu(x)q(x, x^*)};$$

or remains in the current state otherwise. The MH algorithm specifies the kernel to be of the form:

$$K_{\text{MH}}(x, dx^*) := \min\{1, R(x, x^*)q(x, x^*)\} + \alpha_x \delta_x(dx^*),$$

where,

$$\alpha_x := 1 - \int \min\{1, R(x, x^*)\}q(x, x^*)dx^*,$$

is the probability of remaining in the same state and δ_x denotes a probability measure concentrated at x .

The requirement of μ being known point-wise arises when computing R . Thus, a very natural approach for intractable target densities is to approximate μ instead. Beaumont (2003) and Andrieu and Roberts (2009) show that as long as the approximation of μ is unbiased, the Markov chain of this approximate version will have the same invariant distribution as when the “exact” target distribution had been used. This approximation results in a class of techniques called “pseudo-marginal” algorithms.

Many intractable target densities, including those of interest here, can be expressed in terms of marginals of tractable densities: integrals that cannot be evaluated analytically. An obvious example is when we wish to use marginal likelihoods as will be the case in our context; see (10) below.

More formally, let the target density be of the form

$$\mu(x) = \int_{\Theta} \mu(x, \theta)d\theta,$$

where the integrals cannot be solved analytically; and $\theta \in \Theta$ may be considered a latent (or nuisance) variable that is not necessarily of current interest.

Let $\hat{\mu}(x)$ be an estimator of $\mu(x)$ and suppose that $\hat{\mu}(x)$ is unbiased, for all $x \in \mathcal{X}$. That is, if we let $g(\cdot|x)$ denote the density of $\hat{\mu}(x)$, then we have that $\mathbb{E}[\hat{\mu}(x)] = \mu(x)$ for every $x \in \mathcal{X}$, where the expectation is taken with respect to $g(\cdot|x)$. Note that g need not be known.

The estimator $\hat{\mu}(x)$ is random, and its variance will play an important part in producing accurate results—we postpone further discussion until Sect. 4.2. The normalising constant estimator of the sequential Monte Carlo sampler, as detailed in Sect. 3.3, is a common and popular marginal likelihood estimator.

Grouped independence Metropolis Hastings (GIMH), a pseudo-marginal algorithm first presented by Beaumont (2003) and subsequently interpreted by Andrieu and Roberts (2009), can now be presented. Using the simpler reformulation (without auxiliary variables) given by Andrieu and Vihola (2015), the pseudo-code description of GIMH is presented in Algorithm 1.

Here, we have used subscripts (i) and $*$ for $\hat{\mu}$ to emphasise precisely where the marginal likelihood has been (re-)estimated. Henceforth these subscripts will be suppressed for notational clarity.

Note in particular, Algorithm 1 is analogous to the marginal MH algorithm, with the difference being that the acceptance ratio, R , is now approximated by

Algorithm 1 The GIMH Pseudo-Marginal Algorithm

1. Given $x^{(i-1)}$ and unbiased estimate $\hat{\mu}_{(i-1)}(x^{(i-1)})$.
2. Sample:
 - (a) $x^* \sim q(x^{(i-1)}, \cdot)$
 - (b) $\hat{\mu}_*(x^*)|x^* \sim g(\cdot|x^*)$
3. Compute

$$\hat{r}(x^{(i-1)}, x^*) := \frac{\hat{\mu}_*(x^*)q(x^*, x^{(i-1)})}{\hat{\mu}_{(i-1)}(x^{(i-1)})q(x^{(i-1)}, x^*)}$$

4. With probability $\min\{1, \hat{r}\}$ let:

$$x^{(i)} = x^* \text{ and } \hat{\mu}_{(i)}(x^{(i)}) = \hat{\mu}_*(x^*);$$

otherwise

$$x^{(i)} = x^{(i-1)} \text{ and } \hat{\mu}_{(i)}(x^{(i)}) = \hat{\mu}_{(i-1)}(x^{(i-1)}).$$

$$\widehat{R}(x, x^*) := \frac{\hat{\mu}(x^*)q(x^*, x)}{\hat{\mu}(x)q(x, x^*)}.$$

As Andrieu and Roberts (2009) point out, the pseudo-marginal Markov chain should be intuitively thought of as a Markov chain on the extended space of the random variable $(X, \hat{\mu}(X))$, rather than just X itself. They further show that a Markov kernel with this acceptance ratio produces a Markov chain with (and converging to) invariant distribution with marginal μ . A similar justification applies to the algorithm proposed below which is a natural extension of GIMH specialised to the setting of interest. We detail this in Sect. 4.2.2, Proposition 1.

3.3 Estimating the marginal likelihood using an SMC Sampler

Unbiased likelihood estimators are commonly constructed using particle filters, or SMC samplers, as in Andrieu et al. (2010).

SMC samplers (Del Moral et al. 2006) are a class of algorithms that generate collections of weighted samples approximating each of a sequence of target distributions $\{\mu_t\}_{t \geq 1}$ over essentially any random variable defined on some spaces $(\mathcal{X}_t)_{t \geq 1}$. Standard SMC methods achieve this through a combination of sequential IS and variance reduction resampling methods on spaces with increasing dimensions (see, for example, Doucet and Johansen (2011)). However, SMC samplers allow for sequence of distributions defined on some common space \mathcal{X} .

SMC samplers accomplish this by creating a sequence of auxiliary distributions on spaces of increasing dimensions. SMC samplers can also be readily applied to simpler spaces through the use of annealing schemes. See Zhou et al. (2016) and Sect. 5 for details on specification of the annealing scheme within the present Bayesian analysis context.

The sequence of auxiliary distributions are specified using Markov kernels called “backward kernels”, which critically impact the variance of the estimators. See Del Moral et al. (2006) for more details on these so-called backward kernels.

Conventional sequential importance re-sampling algorithms can then be applied to the auxiliary distribution sequence. More precisely, assume that at iteration $t - 1$ the set of N weighted particles approximate $\tilde{\mu}_{t-1}$, denote this $\{W_{t-1}^{(i)}, X_{0:t-1}^{(i)}\}_{i=1}^N$. Here, $W_t^{(i)}$ denotes the *normalised* weight of particle $X_{0:t}^{(i)}$ for all $i = 1, \dots, N$. For iteration t , the path of each particle $X_{0:t-1}^{(i)}$ is extended with a Markov kernel say, $K_t(x_{t-1}, x_t)$, yielding the set of particles $\{X_{0:t}^{(i)}\}_{i=1}^N$ to which IS is applied. Subsequently, the weights are updated by a factor called incremental weights.

Importantly, particularly when interested in approximating marginals such as the model evidence (marginal likelihood); if $\mu_t(x_t) = \gamma_t(x_t)/Z_t$ is only known up to a normalising constant, the unnormalised incremental weights

$$w_t(x_{t-1}, x_t) = \frac{\gamma_t(x_t)}{\gamma_{t-1}(x_{t-1})}$$

can be used. Note: here, and as practised in the studies below, we assume that K_t is constructed to be μ_t -invariant, $\mu_t \ll \mu_{t-1}$ and the associated time-reversal kernel is used for the backward kernel.

In fact, this collection of algorithms provide unbiased (see (Del Moral 2004, Proposition 7.4.1)) estimates of Z_t/Z_{t-1} via

$$\frac{\widehat{Z}_t}{Z_1} = \prod_{p=2}^t \frac{\widehat{Z}_p}{Z_{p-1}} = \prod_{p=2}^t \sum_{i=1}^N W_{p-1}^{(i)} w(X_{p-1:p}^{(i)}). \tag{3}$$

These normalising constant estimators satisfy a Central Limit Theorem (Del Moral 2004, Proposition 9.4.1) and have asymptotic variance that are inversely proportional to the number of particles N (see also Chopin and Papaspiliopoulos (2020) for a recent review). As such, N is an important tuning parameter when this sampler is used within a pseudo-marginal algorithm (Doucet et al. 2015).

4 Methodology

We now turn our attention to the development and presentation of the models to encode spatial dependence together with computational methods that can be used to perform inference under, and characterise, these models.

There is a considerable literature on spatial modelling in Bayesian data analysis (see, for example, Gelfand et al. (2010)). However, some specialised settings, such as that encountered when dealing with PET images, are yet to

be explored to the same extent. In these cases, there may be significantly large amounts of data to be analysed and relatively sophisticated models may be required for meaningful investigation. Such approaches may need to incorporate application-specific structure and information. For example, Bezener et al. (2018) presents a method of variable selection which models spatial dependence using hierarchical spatial priors and parcellation for the analysis of MRI images.

Studies that attempt to model spatial relationships in PET images are scarce, and tend not to use a Bayesian approach. One such example is Zhou et al. (2002) who proposes a method to allow for spatial dependence. That is, spatial constraint using nonlinear ridge regression is used to improve upon parametric images produced by conventional weighted NLS methods. They showed that doing so reduced the percentage mean square error by 60–80% in simulation studies. These promising results motivate the incorporation of spatial dependence when analysing PET images, particularly in Bayesian frameworks where such studies have been largely absent.

On the other hand, there exists many computationally efficient methods for complex settings, such as PET data analysis. These include, in the PET context, among many others: Gunn et al. (2002); Peng et al. (2008); Jiang et al. (2009); Zhou et al. (2013, 2016). Computational tractability in these methods is typically attained through the assumption of spatial independence.

A natural strategy is, therefore, to adapt existing non-spatial methods to a model that does incorporate spatial dependence. This strategy is particularly natural when modelling spatial dependence via a Potts model as it can be efficiently targeted by MCMC methods with single-site update schemes. In this section we construct a class of models that allows for such computational methods.

Having presented the model that allows for effective incorporation of spatial dependence, we then present a natural extension of the generic pseudo-marginal approach, discussed in Sect. 3 above, to this specialised setting in Sect. 4.2 along with the theoretical justification of the method, which follows from that of the standard pseudo-marginal approach and some additional considerations. Further extensions of the method are described in Sect. 4.3, and a straightforward approximation of the proposed algorithm is introduced in Sect. 4.4.

4.1 Spatial Bayesian model selection using a Potts prior

Given spatial data (or spatio-temporal data, such as PET) $Y = (Y_v : Y_v \in \mathbb{R}^k, v \in V)$, associate a graph $G_Y = (V_Y, E_Y)$ that encodes the geometry of spatial dependence. That is, data points Y_u is conditionally independent of Y_v given $Y_{-\{u,v\}}$ unless $(u, v) \in E_Y$. For example, for image

data a simple (finite) lattice would suffice. We will call Y_v the node data point at the node $v \in V$.

Building good statistical models for Y can be difficult; model selection would involve selecting among many complex and computationally intensive models from the model space. Analysis in this setting can be simplified by using existing simpler models such as those that assume spatial independence and model data at the node-level.

Such an approach has two significant strengths. Firstly, this approach can readily exploit many existing techniques which make spatial-independence assumptions, allowing them to incorporate spatial dependence. Secondly, using this approach allows for considerable computational simplifications as will be seen later.

4.1.1 A generic model for spatial model selection

Formally, consider parametric models for the node data point Y_v , for all $v \in V_Y$, rather than the whole data set Y . Denote the parameter of these models by θ_v , with parametric space Θ_v for each node $v \in V$.

It is immediate that a Potts model can be used to encode spatial relationships between models at different locations in space; define $X = (X_v \in \mathcal{X} : v \in V_Y)$ a Markov random field with respect to the graph G_Y . Here, the coupling constant J is treated as known and constant.

The parametric distribution for Y_v , now has parameters (θ_v, X_v) . Here, X_v is a component of the Potts random variable X , and so has spatial dependence dictated by E_Y . The model order at node v is given by X_v . Given a family of parameterised prior densities for each possible model denoted generically p , over the parameter space; we may compactly summarise this hierarchical model as:

$$X|J, G_Y \sim \text{Potts}(J, G_Y), \tag{4}$$

$$\theta_v|X_v, \xi_v \sim p_{X_v}(\cdot|\xi_v) \text{ for all } v \in V_Y, \tag{5}$$

$$Y_v|\theta_v, X_v \sim f_{X_v}(\cdot|\theta_v) \text{ for all } v \in V_Y. \tag{6}$$

Here, recall that f_M denotes the density associated with statistical model S_M with model order M . A similar convention is used for the (parameter) prior to denote that a different prior distribution must be used for each model order in the Bayesian context. Specifically, here we have that the model order is $M = X_v$ for each node $v \in V$. This generic model is represented as a plate diagram in Fig. 1.

Here, ξ_v is a parameter of the distribution p . In principle, a hyperprior could be attached to this parameter with no particular difficulty; we work with this simple setting in the interest of parsimony. To this end, we fix ξ_v to be a known hyperparameter common to all nodes, v , and so we write ξ instead of ξ_v in what follows.




Fig. 1 Plate diagram of the proposed generic model

Additionally, given that we now treat the model order as a random variable, henceforth we use the convention that $f(\cdot|\theta_v, M) = f_M(\cdot|\theta_v)$. Note that, at the nodal (pixel) level, Bayesian model selection would often involve the computation of the marginal likelihood, $f(y_v|M) = \int_{\Theta_v} f(y_v|\theta_v, M)p(\theta_v|M\xi)d\theta_v$. We are interested in the marginal likelihood of the whole image Y , which may involve multiple intractable high-dimensional integrals. We impose an assumption of conditional independence, discussed next, which allows us to make significant simplifications. For a simple example, see the toy model described in Sect. 5.1; This toy model will be used for simulation experiments to evaluate the proposed method.

In realistic settings, it may be that the parameter $\theta = (\theta_v)_{v \in V}$ also exhibits spatial dependence. The above hierarchical model, makes the underlying assumption that it does not and this setting is the focus of the present paper; where we have found the incorporation of spatial structure at the level of model sufficient to improve upon the state of the art for problems of interest. Further generalisation would be possible and provides an interesting direction for further work. We encode this in Assumption 1 which will be a standing assumption throughout this paper and is discussed in Sect. 4.1.3.

Assumption 1 (Conditional Independence)

$$\theta_v|X_v, \xi \perp X_{-v}, \theta_{-v} \text{ for all } v \in V.$$

The primary aim of the proposed methodology is to infer X . Although θ is inferred as a by-product of the proposed method, it can be thought of as a latent variable in this framework and may be a nuisance parameter in some settings which mitigates the impact of modelling the parameters in this way.

Finally, note that, as a direct consequence of the above model and assumption, the marginal likelihood of \mathbf{Y} can be written as the product of the marginal likelihoods of Y_v over all $v \in V$ —see Sect. 4.1.4 below.

4.1.2 Spatial bayesian model selection

To exploit the presented model, denote by $\mathcal{S}_v = \{S_{v, M_v} : M_v \in \mathcal{M}\}$ the model space used at each nodal data point Y_v . I.e. every node in the graph is associated with a model M_v from a set common to all nodes, \mathcal{M} . Thus a statistical model is associated with each node in the graph and hence each data point,

$$S_{v, M_v} = \{f(\cdot | \theta_v, M_v) : \theta_v \in \Theta_{v, M_v}\}.$$

The notation Θ_{v, M_v} is used here to emphasis that since M_v is the model order at node v , the parameter space will be dependent on it. For example, for compartmental models, the model order dictates the dimension of the parameter space. These collections of models can be used to generate a model space for the whole data set \mathbf{Y} . Specifically, a set of candidate models for \mathbf{Y} can be formulated as

$$\mathcal{S}_Y = \{(S_{v, M_v}) : M_v \in \mathcal{M}\}_{v \in V}.$$

Writing $\mathbf{M} = (M_v)_{v \in V} \in \mathcal{M}^V$, Bayesian model selection of spatial data \mathbf{Y} can be thought of as inference of the model order parameter \mathbf{M} in the space \mathcal{M}^V . Each realisation of \mathbf{M} is called a configuration, note that there are $|\mathcal{M}|^{|V|}$ candidate models for \mathbf{Y} .

As mentioned before, the Potts model is a natural choice for the prior distribution over model order. For model selection, we adapt the hierarchical model above such that:

1. $\mathbf{M} = (M_v)_{v \in V}$ is a Markov random field with a Potts distribution, with spatial dependence represented by E_Y as before;
2. Each θ_v , for all $v \in V$, is dependent on M_v (e.g. model order determines parameter dimension in some cases, accordingly we denote the prior $p(\cdot | \xi, M_v)$) and known constant hyper-parameter ξ , importantly it is spatially independent (given the model order M_v);
3. Y_v , the observed data, has likelihood $f(\cdot | \theta_v, M_v)$, where the model M_v dictates which model is selected, for all $v \in V$.

Formally, the above can be summarised:

$$\mathbf{M} \sim \text{Potts}(J, G_Y) \tag{7}$$

$$\theta_v | M_v, \xi \sim p(\cdot; M_v, \xi) \text{ for all } v \in V, \tag{8}$$

$$Y_v | \theta_v, M_v \sim f(\cdot; \theta_v, M_v) \text{ for all } v \in V. \tag{9}$$

4.1.3 Discussion of Assumption 1

Assumption 1 may not hold in general: if there is spatial dependence between the generative model which describes each vertex then there may also be dependence between the parameters of those models. In some settings it can be viewed as an approximation which may be tolerable in order to allow inference under a model which is, at least, better than the assumption of full spatial independence. In other settings, this assumption may not be unrealistic. It is also noteworthy that, if the model order dictates the dimensions of the parameter space at that node, and the space contains dependent parameters, it may be difficult to incorporate spatial dependence at a the parameter level in a meaningful manner.

For example, if considering compartmental models: if two adjacent pixels were to contain different number of compartments, it may be difficult to say anything meaningful about the spatial dependence of the micro-parameters such as the transfer rates. Such considerations make it difficult to impose some more general Markov random field over the parameters.

Furthermore, incorporating spatial dependence at a model order level will typically take some precedence over spatial dependence at the parameter level. For instance, in the PET setting, generally we would like to know if two adjacent pixels have the same compartments (i.e. whether or not the localised area contains receptors) before we think about the transfer rates.

To summarise the problem of interest in general terms: the image \mathbf{Y} is modelled node-wise using the density $f(\cdot; \theta_v, M_v)$, $v \in V_Y$, with priors over the parameters θ and \mathbf{M} . In particular, we propose the use of the Potts model as a prior over discrete parameters \mathbf{M} to allow for tractable incorporation of spatial dependence in images. As mentioned above, θ will be treated as a latent variable—this is discussed in detail next.

4.1.4 Inference from the proposed model

Having specified this class of models, we now turn to the task of inference. Let $p(\mathbf{M})$ denote the Potts prior probability mass function over $\mathbf{M} \in \mathcal{M}^V$. Then, the model posterior density, for \mathbf{Y} , is denoted

$$\pi(\mathbf{M} | \mathbf{y}) \propto f(\mathbf{y} | \mathbf{M}) p(\mathbf{M}).$$

Here, $f(\mathbf{y} | \mathbf{M})$ will be called the *graph* marginal likelihood of \mathbf{y} ; it encodes a generative description of the data under the proposed model.

The proposed method allows us to incorporate spatial dependence at the level of the discrete parameter \mathbf{M} only. Subsequently, for computational tractability we impose Assumption 1: given the model order M_v the random variable Y_v at each node $v \in V$ is independent. In other words,

as per (5) and (6), we have that

$$f(\mathbf{y}|\mathbf{M}) = \prod_{v \in V} f(y_v|M_v).$$

Here $f(y_v|M_v)$ denotes the likelihood at each node v , and will be henceforth called the *node-wise* likelihood. This assumption is important as it makes the graph likelihood tractable, and provides computational simplifications discussed later. Typically, as described above, a parametric model is used, thus $f(y_v|M_v)$ will in fact be a marginal likelihood or evidence (i.e. an integral, which is likely to be analytically intractable).

We finally have the probability mass function of primary computational interest

$$\begin{aligned} \pi(\mathbf{M}|\mathbf{y}) &\propto \left\{ \prod_{v \in V} f(y_v|M_v) \right\} p(\mathbf{M}) \\ &\propto \left\{ \prod_{v \in V} \int_{\Theta_v} f(y_v|\theta_v, M_v) p(\theta_v|M_v, \xi) d\theta_v \right\} \\ &\quad \times \exp\left(J \sum_{v \sim u} \delta_{M_v, M_u} \right). \end{aligned} \tag{10}$$

Recall the mild assumption that hyper-parameter ξ is known and henceforth suppressed in notation.

Even in this general formulation, this mass function is not particularly attractive. There are some positives: It is straightforward to evaluate up to the intractable normalising constant; and Assumption 1 does seemingly lead to a simpler computation even at this stage, as only the node-wise marginal likelihoods $f(y_v|M_v)$ need to be computed rather than the high-dimensional integration of the full marginal likelihood $f(\mathbf{y}|\mathbf{M})$.

However, even in this simpler setting, the graph marginal likelihood poses two computational difficulties: (i) it is a product of $|V|$ integrals and (ii) these integrals in general are analytically intractable in most cases of interest.

4.2 The node-wise Pseudo-marginal algorithm

We now turn to extending the techniques discussed in Sect. 3, to build the novel computational method for use in the class of models constructed above.

The aim is to characterise the posterior density, $\pi(\mathbf{M}|\mathbf{y})$, of the model given in (7)–(9). Suppose, for now, that a generic MH algorithm could be used in this setting. That is, to generate a MH Markov chain that targets the posterior density. This then gives the acceptance ratio:

$$R(\mathbf{M}, \mathbf{M}^*) = \frac{\pi(\mathbf{M}^*|\mathbf{y})q(\mathbf{M}^*, \mathbf{M})}{\pi(\mathbf{M}|\mathbf{y})q(\mathbf{M}, \mathbf{M}^*)}.$$

Even with a tractable marginal likelihood, computing this acceptance ratio would be costly. Under Assumption 1 (given \mathbf{M} the parameters $\boldsymbol{\theta} = (\theta_v)_{v \in V}$, of the model over the data \mathbf{y} , are independent) the computational complexity of the problem is reduced considerably; now, only the node-wise marginal likelihoods need to be computed. However, every time a new configuration is proposed, up to $|V|$ new integrals would then need to be computed and the acceptance probability is likely to be very small. Given that the number of possible Potts configurations is $|\mathcal{M}|^{|V|}$, growing geometrically with the number of nodes (i.e. pixels, for image data), the computational costs are not appealing.

Under Assumption 1 it is natural to employ a Metropolis-within-Gibbs approach to mitigate these difficulties and to obtain significant computational simplifications. Suppose that the MCMC sampler is currently at some configuration $\mathbf{M} = (M_v)_{v \in V}$. Under a single variable updating schedule: *only* the model order M_v , at some node $v \in V$, will be proposed to be changed in a given step. We will refer to this as node-wise updating schedule, since we update the model order at a single node v only. Denote the proposed state to be M_v^* ; and $\mathbf{M}_{[v]}^*$ to be the vector equivalent to \mathbf{M} but with M_v^* as the v -th coordinate. For simplicity, assume that the proposal distribution is symmetric, and the acceptance ratio is

$$\begin{aligned} \frac{\pi(\mathbf{M}_{[v]}^*|\mathbf{y})}{\pi(\mathbf{M}|\mathbf{y})} &= \frac{f(\mathbf{y}|\mathbf{M}_{[v]}^*)p(\mathbf{M}_{[v]}^*)}{f(\mathbf{y}|\mathbf{M})p(\mathbf{M})} \\ &= \frac{f(y_v|M_v^*)p(\mathbf{M}_{[v]}^*)}{f(y_v|M_v)p(\mathbf{M})} \\ &= \frac{\int f(y_v|\theta_v, M_v^*)p(\theta_v|M_v^*)d\theta_v p(\mathbf{M}_{[v]}^*)}{\int f(y_v|\theta_v, M_v)p(\theta_v|M_v)d\theta_v p(\mathbf{M})}. \end{aligned}$$

Thus, the acceptance probability of each Metropolis-within-Gibbs step requires the computation of only a single integral. A complete sweep over the graph thus requires $|V|$ such integrals but one could expect a reasonable proportion of these proposed moves to be accepted, in contrast to a global proposal which sought to update every node simultaneously.

Typically, even the node-wise marginal likelihood $f(y_v|M_v)$ will be difficult or impossible to evaluate analytically. A pseudo-marginal MH algorithm is a natural choice in such cases and is explored in the next section.

4.2.1 Graph model selection using node-wise likelihood marginal estimates

The node-wise update schedule, gives rise to nested iterations. For clarity, term the outer iteration, indexed with i , of a full pass through the whole graph as the graphical iteration. The inner iteration over V will be termed the node-wise iteration.

Suppose that the MH Markov chain denoted, $(\mathbf{M}^{(i)})_{i=1}^n$, is at graphical iteration i and that the state at node v is being proposed to be changed. Next, since the node-wise marginal likelihood is a normalising constant, for simplicity denote

$$Z_v(M) = f(y_v|M_v = M) \text{ for all } M \in \mathcal{M}.$$

That is, $Z_v(M)$ is the normalising constant at node v when the model order $M_v = M$.

At graphical iteration i , at node v , for given model order proposal $M_v^{(i)}$, let $(\widehat{Z}_v)^{(i)}$ denote the approximation of the marginal $f(y_v|M_v = M_v^{(i)})$. Associate with this generic estimator, the computation cost parameter N (for the SMC sampler this was the number of particles). As mentioned before this is a tuning parameter for the pseudo-marginal, thus also the algorithm presented below.

It is important to note the slight abuse of notation here— $(\widehat{Z}_v)^{(i)}$ is dependent on the model order proposal $M_v^{(i)}$. As previously stated, we require that $(\widehat{Z}_v)^{(i)}$ is unbiased. Additionally, recall we denote $g(\cdot|M_v = M_v^{(i)})$ the possibly unknown density of $(\widehat{Z}_v)^{(i)}$.

Following the node-wise updating schedule, whenever a new state $M_v^* \sim q(M_v^{(i-1)}, \cdot)$ is proposed, a new node-wise marginal estimation, denoted $(\widehat{Z}_v)^*$, is recomputed. The node-wise acceptance ratio, $\widehat{r}(\mathbf{M}^{(i-1)}, \mathbf{M}_{[v]}^{(i-1)*})$, is then

$$\frac{p(\mathbf{M}_{[v]}^*)(\widehat{Z}_v)^* q(M_v^*, M_v^{(i-1)})}{p(\mathbf{M}^{(i-1)})(\widehat{Z}_v)^{(i-1)} q(M_v^{(i-1)}, M_v^*)}.$$

We emphasise that the computation of \widehat{r} , precisely speaking, involves not just the standard MH proposal of a new configuration in \mathcal{M}^V , but also a proposal of the random variable \widehat{Z}_v : it is a pseudo-marginal algorithm in the sense of Andrieu and Roberts (2009), with the added subtlety that the unbiased estimates of the marginal likelihood associated with every other node in the graph are also retained as a part of the extended state. The justification for using this acceptance ratio will be detailed in Sect. 4.2.2

The above choices and assumptions lead to the algorithm which we term the “Node-Wise Pseudo-Marginal” (NWPM) algorithm; presented in pseudo-code in Algorithm 2.

Within the context of this paper, the NWPM algorithm proposed here is a broadly applicable extension of the GIMH algorithm described above in Algorithm 1.

The output of this algorithm, the Monte Carlo sample $(\mathbf{M}^{(i)})_{i=1}^n$, can be used to perform model selection in a principled manner. For example, this could be done using the node-wise marginal modal model order: i.e., at each node $v \in V$ the model order which occurs the most in the (marginal) chain $(M_v^{(i)})_{i=1}^n$ is selected. Additionally, samples of the latent variables θ are typically available as byproducts of the marginal likelihood estimates and can be used to per-

Algorithm 2 The Node-Wise Pseudo-Marginal Algorithm

Given: (i) target density $\mu(\mathbf{M}) := \pi(\mathbf{M}|y)$; (ii) unbiased node-wise marginal likelihood estimators $\widehat{Z}_v(M) \approx \int f(y_v|M, \theta_v) p(\theta_v|M, \xi) d\theta_v$ for all $M \in \mathcal{M}$, $v \in V$; and (iii) coupling constant J , tuning parameter N .

1. At $i = 1$:
 - (a) Initialise $\mathbf{M}^{(1)}$.
 - (b) Sample $(\widehat{Z}_v)^{(1)} \sim g(\cdot|M_v^{(1)})$ for each $v \in V$.
2. For $i = 2, \dots$:
 - (a) For $v \in V$:
 - (i) Sample: $M_v^* \sim q(M_v^{(i-1)}, \cdot)$.
 - (ii) Sample: $(\widehat{Z}_v)^* | m_v^* \sim g(\cdot|M_v^*)$.
 - (iii) Compute:

$$\widehat{r} = \frac{p(\mathbf{M}_{[v]}^*)(\widehat{Z}_v)^* q(M_v^*, M_v^{(i-1)})}{p(\mathbf{M}^{(i-1)})(\widehat{Z}_v)^{(i-1)} q(M_v^{(i-1)}, M_v^*)}$$

- (iv) With probability $\min\{1, \widehat{r}\}$ take:

$$M_v^{(i)} = M_v^* \text{ and } (\widehat{Z}_v)^{(i)} = (\widehat{Z}_v)^*$$

otherwise,

$$M_v^{(i)} = M_v^{(i-1)} \text{ and } (\widehat{Z}_v)^{(i)} = (\widehat{Z}_v)^{(i-1)}$$

form parameter inference as demonstrated in the numerical studies below.

An advantage of this pseudo-marginal based approach in this setting is its flexibility. It can be used with essentially any unbiased estimator of marginal likelihood and hence allows existing technology from any application domain to be employed. The incorporation of spatial dependence via the modelling framework and algorithm developed here can come at the very little computational cost by approximating or making pragmatic adaptations to Algorithm 2. We employ such techniques in the simulation studies presented in Sect. 5 and discuss in more detail in Sect. 4.4 below.

Finally, note that for convenience we assume that the tuning parameter N is fixed to the same value for all $v \in V$. This parameter typically determines the variance of the estimator $\widehat{Z}_v(M)$ and additional flexibility could be obtained by allowing it to vary between vertices. Unsurprisingly, a larger N should give a better performance at the expense of greater computational cost. Indeed Andrieu and Roberts (2009), and more recently Andrieu and Vihola (2016), showed that when the dispersion of the marginal likelihood estimates are larger the mixing rate of any pseudo-marginal MH algorithm decreases. Of course, there is a higher computational cost when using larger N to reduce the estimate variance. The effects of this tuning parameter, and the trade off, on the mixing of the Markov chain in a standard pseudo-marginal algorithm have been studied by Sherlock et al. (2015); Doucet et al. (2015); Sherlock (2016). These works give theoretical justification and numerical studies for the recommendation

that N should be chosen such that the variance of the log-likelihood estimator is close to one.

4.2.2 Marginal invariant distribution of NWPM

Here we establish the formal validity of a class of algorithms which includes the NWPM, allowing for more general moves than those described in Algorithm 2. We show that this class of algorithms leads to a Markov chain with an invariant distribution coinciding with π , as given in (10). For further discussion on specifying proposals within this setting see Sect. 4.3.

Recalling Assumption 1, we may write the target distribution of the NWPM algorithm as

$$\mu(\mathbf{M}) \propto p(\mathbf{M}) \prod_{v \in V} \int_{\Theta_v} f(y_v | \theta_v, M_v) p(\theta_v | M_v, \xi) d\theta_v,$$

for $\mathbf{M} = (M_v)_{v \in V} \in \mathcal{M}^V$.

Given \mathbf{M} , denote the random vector of the unbiased node-wise marginal likelihoods estimators as $\widehat{\mathbf{Z}} | \mathbf{M} \sim g(\cdot | \mathbf{M})$, i.e.,

$$\widehat{\mathbf{Z}} = (\widehat{Z}_v(M_v) : v \in V).$$

Henceforth using the shorthand $\widehat{Z}_v^M := \widehat{Z}_v(M)$ for all $M \in \mathcal{M}$; we note that, due to independence,

$$g(\widehat{\mathbf{Z}} | \mathbf{M}) \propto \prod_{v \in V} g(\widehat{Z}_v^{M_v} | M_v).$$

Finally, define the extended target density

$$\widehat{\mu}(\mathbf{M}, \widehat{\mathbf{Z}}) = \frac{p(\mathbf{M}) \prod_{v \in V} \widehat{Z}_v^{M_v} g(\widehat{Z}_v^{M_v} | M_v)}{\sum_{\mathbf{M}' \in \mathcal{M}^V} p(\mathbf{M}') \prod_{v \in V} f(y_v | M'_v)} \quad (11)$$

$$= \frac{p(\mathbf{M}) \prod_{v \in V} \widehat{Z}_v^{M_v} g(\widehat{Z}_v^{M_v} | M_v)}{f(\mathbf{y})}. \quad (12)$$

Here, we let $f(\mathbf{y})$ denote the marginal probability of the observed data, integrating out unknown parameters and unknown model orders; Given observation $\mathbf{Y} = \mathbf{y}$, we treat it as a normalising constant. The NWPM algorithm has similar theoretical properties as GIMH. As such, the remainder of the argument of formal justification then follows the same as the GIMH case, see Andrieu and Roberts (2009).

Proposition 1 *Let \mathbf{M} denote a random variable in the graph model order space \mathcal{M}^V . Let $\widehat{\mathbf{Z}} | \mathbf{M}$, g , μ and $\widehat{\mu}$ be defined as above.*

For any $U \subset V$, let q_U denote a Markov kernel on $(\mathcal{M} \times \mathbb{R}_+)^V$ of the form

$$q_U((\mathbf{M}, \widehat{\mathbf{Z}}), (\mathbf{M}^*, \widehat{\mathbf{Z}}^*)) = \delta_{(\mathbf{M}_{-U}, \widehat{\mathbf{Z}}_{-U})}(\mathbf{M}_{-U}^*, \widehat{\mathbf{Z}}_{-U}^*)$$

$$\times q_U^{\mathcal{M}}(\mathbf{M}_U, \mathbf{M}_U^*) \prod_{v \in U} g(\widehat{Z}_v^* | M_v^*),$$

where $q_U^{\mathcal{M}}$ denotes a Markov kernel on \mathcal{M}^U and we slightly abuse density notation using the Dirac delta functions to indicate that variables associated with nodes outside U are unchanged (absolute continuity of the numerator and denominator of the Metropolis-Hastings ratio is ensured by the symmetry of this singular part of the kernel).

The standard Metropolis-Hasting acceptance probability with target distribution $\widehat{\mu}$ can be expressed as:

$$1 \wedge \frac{p(\mathbf{M}^*) (\prod_{v \in U} \widehat{Z}_v^*) q_U^{\mathcal{M}}(\mathbf{M}_U^*, \mathbf{M}_U)}{p(\mathbf{M}) (\prod_{v \in U} \widehat{Z}_v) q_U^{\mathcal{M}}(\mathbf{M}_U, \mathbf{M}_U^*)} \quad (13)$$

and the marginal distribution of \mathbf{M} under $\widehat{\mu}$ is μ .

Proof First, consider the acceptance ratio of a MH (Metropolis-Hastings) algorithm with target distribution $\widehat{\mu}$ and proposal kernel q_U :

$$\widehat{r} = \frac{\widehat{\mu}(\mathbf{M}^*, \widehat{\mathbf{Z}}^*) q_U((\mathbf{M}^*, \widehat{\mathbf{Z}}^*), (\mathbf{M}, \widehat{\mathbf{Z}}))}{\widehat{\mu}(\mathbf{M}, \widehat{\mathbf{Z}}) q_U((\mathbf{M}, \widehat{\mathbf{Z}}), (\mathbf{M}^*, \widehat{\mathbf{Z}}^*))}$$

upon inserting the definition of q_U one observes that the numerator is absolutely continuous with respect to the denominator and the singular elements simply impose that $\mathbf{M}_{-U} = \mathbf{M}_{-U}^*$ and $\widehat{\mathbf{Z}}_{-U} = \widehat{\mathbf{Z}}_{-U}^*$, and we have:

$$\begin{aligned} \widehat{r} &= \frac{\widehat{\mu}(\mathbf{M}^*, \widehat{\mathbf{Z}}^*) q_U^{\mathcal{M}}(\mathbf{M}_U^*, \mathbf{M}_U) \prod_{v \in U} g(\widehat{Z}_v | M_v)}{\widehat{\mu}(\mathbf{M}, \widehat{\mathbf{Z}}) q_U^{\mathcal{M}}(\mathbf{M}_U, \mathbf{M}_U^*) \prod_{v \in U} g(\widehat{Z}_v^* | M_v^*)} \\ &= \frac{p(\mathbf{M}^*) \prod_{v \in V} \widehat{Z}_v^* g(\widehat{Z}_v^* | M_v^*)}{p(\mathbf{M}) \prod_{v \in V} \widehat{Z}_v g(\widehat{Z}_v | M_v)} \\ &\quad \times \frac{q_U^{\mathcal{M}}(\mathbf{M}_U^*, \mathbf{M}_U)}{q_U^{\mathcal{M}}(\mathbf{M}_U, \mathbf{M}_U^*)} \prod_{v \in U} \frac{g(\widehat{Z}_v | M_v)}{g(\widehat{Z}_v^* | M_v^*)} \\ &= \frac{p(\mathbf{M}^*) \prod_{v \in U} \widehat{Z}_v^*}{p(\mathbf{M}) \prod_{v \in U} \widehat{Z}_v} \\ &\quad \times \underbrace{\frac{q_U^{\mathcal{M}}(\mathbf{M}_U^*, \mathbf{M}_U)}{q_U^{\mathcal{M}}(\mathbf{M}_U, \mathbf{M}_U^*)} \prod_{v \notin U} \frac{\widehat{Z}_v^* g(\widehat{Z}_v^* | M_v^*)}{\widehat{Z}_v g(\widehat{Z}_v | M_v)}}_{=1}, \end{aligned}$$

where the final factor is equal to one almost surely under the proposal distribution, and the result follows.

The marginal distribution of \mathbf{M} follows by the essentially the same argument as in the standard pseudo-marginal context.

$$\int_{\mathbb{R}_+^V} \widehat{\mu}(\mathbf{M}, \widehat{\mathbf{Z}}) d\widehat{\mathbf{Z}}$$

$$\begin{aligned}
 &= \frac{p(\mathbf{M})}{f(\mathbf{y})} \prod_{v \in V} \int_{\mathbb{R}_+} \widehat{Z}_v^{M_v} g(\widehat{Z}_v^{M_v} | M_v) d\widehat{Z}_v^{M_v} \\
 &= \frac{p(\mathbf{M})}{f(\mathbf{y})} \prod_{v \in V} \mathbb{E}_g[\widehat{Z}_v(M_v)] \\
 &= \frac{p(\mathbf{M}) \prod_{v \in V} f(y_v | M_v)}{f(\mathbf{y})} \\
 &= \mu(\mathbf{M}).
 \end{aligned}$$

□

Clearly, the node-wise proposals described previously fit this framework with $U = \{u\}$ being the single node being updated, although this framework would allow a somewhat broader class of proposals and blocked Metropolis-within-Gibbs type strategies to be explored. Standard arguments mean that a mixture or cycle of such kernels will also preserve $\widehat{\mu}$ as the invariant distribution.

One might anticipate that the theoretical properties of the NWPM algorithm might be further characterised by extending techniques from the standard pseudo-marginal setting such as Andrieu and Vihola (2015, 2016); Andrieu et al. (2021) to this one, indeed it would be surprising were it to behave substantially differently, but doing this rigorously would involve some delicate technical work and is beyond the scope of this paper.

Next, we discuss some immediate extensions and innovations of this proposed methodology; resulting in further algorithms, approximations and future approaches.

4.3 Multiple augmentation Pseudo-marginal algorithms

In this section we will explore the potential of further augmentation of the state space. The specialised setting of the problem, allows approaches which to the authors’ knowledge, have not been previously studied. The strategy developed here allows for a number of further extensions. Some simple approaches are discussed below—methods based on these approaches will also be evaluated in the numerical studies.

The augmentation of the state space used to justify node-wise pseudo-marginal algorithms can be further extended by adding an estimate of the marginal likelihood associated with every possible model at every node to the state space. Although doing so may seem counter-intuitive and leads to a rather large state space, it allows a number of algorithmic innovations. In particular, by considering the following extended state space it will be possible to use a variety of standard MH moves in order to explore this space.

Specifically, the further extended space allows us to decouple the updating of the marginal likelihood estimates from those of the model configurations themselves; the standard

pseudo-marginal update would require that these two updates occur simultaneously. We exploit this decoupling when we go on to consider the possibility of updating the marginal likelihood estimates with lower frequency than the model configurations.

Let $\bar{\mathbf{Z}} := (\widehat{Z}_v(M) : v \in V, M \in \mathcal{M})$ be a vector of marginal likelihood estimator for each model order $M \in \mathcal{M}$, at every node $v \in V$. Consider, next, a further extended form of the target density (12):

$$\bar{\mu}(\mathbf{M}, \bar{\mathbf{Z}}) = \frac{p(\mathbf{M})}{f(\mathbf{y})} \prod_{v \in V} \widehat{Z}_v^{M_v} \left(\prod_{m' \in \mathcal{M}} g(\widehat{Z}_v^{m'} | M') \right).$$

The extended joint density $\bar{\mu}$ is constructed such that its marginal over \mathbf{M} coincides exactly with the *correct posterior distribution*. Indeed, using essentially the same argument as in Proposition 1, we have:

$$\begin{aligned}
 &\int_{\mathbb{R}_+^{V \times \mathcal{M}}} \bar{\mu}(\mathbf{M}, \bar{\mathbf{Z}}) d\bar{\mathbf{Z}} \\
 &= \frac{p(\mathbf{M})}{f(\mathbf{y})} \prod_{v \in V} \int_{\mathbb{R}_+^{\mathcal{M}}} \widehat{Z}_v^{M_v} \left(\prod_{m' \in \mathcal{M}} g(\widehat{Z}_v^{m'} | M') d\widehat{Z}_v^{m'} \right) \\
 &= \frac{p(\mathbf{M})}{f(\mathbf{y})} \prod_{v \in V} \int_{\mathbb{R}_+} \widehat{Z}_v^{M_v} g(\widehat{Z}_v^{M_v} | M_v) d\widehat{Z}_v^{M_v} \\
 &= \mu(\mathbf{M}).
 \end{aligned}$$

This further extended target distribution allows for some generalisations of the standard pseudo-marginal algorithm in the context of interest; it depends fundamentally on the fact that the variables associated with each node of the graph take values within a small finite set. It is possible to consider a variety of MH like moves applied to this extended target density. For simplicity we will consider only two types of proposal here: one which changes a single node’s associated model and another which refreshes the likelihood estimates for a single node.

Denote the proposal of the random vector $\bar{\mathbf{Z}} = (\widehat{Z}_v^M : v \in V, M \in \mathcal{M})$ by $\bar{\mathbf{Z}}^* := ((\widehat{Z}_v^M)^* : v \in V, M \in \mathcal{M})$. Intuitively, $\bar{\mathbf{Z}}^*$ denotes a re-computation of all the components of $\bar{\mathbf{Z}}$.

First consider a proposal Q_u^1 for the model associated with node u . Recalling that $M_{[u]}^*$ is equivalent to \mathbf{M} with component $M_u^* \neq M_u$, we have

$$\begin{aligned}
 &Q_u^1 \left((\mathbf{M}, \bar{\mathbf{Z}}), (\mathbf{M}^*, \bar{\mathbf{Z}}^*) \right) \\
 &= \left(q_u^1(M_u, M_u^*) \prod_{v \neq u} \delta_{M_v, M_v^*} \right) \prod_{v \in V} \prod_{M \in \mathcal{M}} \delta_{\widehat{Z}_v^M}((\widehat{Z}_v^M)^*)
 \end{aligned}$$

$$= \underbrace{\delta_{\mathbf{M}_{-u}, \mathbf{M}_{-u}^*} q_u^1(\mathbf{M}_u, \mathbf{M}_u^*)}_{\text{Gibbs-like state proposal}} \underbrace{\prod_{v \in V} \prod_{M \in \mathcal{M}} \delta_{\widehat{Z}_v^M}((\widehat{Z}_v^M)^*)}_{\text{Same marginal estimates}}.$$

Here $\delta_{x,y}$ is the Kronecker delta, taking value 1 if $x = y$ and 0 otherwise and $\delta_z(z^*)$ is, with the obvious abuse of notation, the singular measure concentrated at z evaluated over an infinitesimal neighbourhood of z^* . In this proposal, only a change of \mathbf{M}_u^* is proposed—in particular, $\bar{\mathbf{Z}}$ and \mathbf{M}_{-u} do not change. Subsequently, the usual Metropolis-Hastings acceptance probability for this move is:

$$\begin{aligned} & 1 \wedge \frac{\bar{\mu}(\mathbf{M}_{[u]}^*, \bar{\mathbf{Z}}) q_u^1(\mathbf{M}_u^*, \mathbf{M}_u)}{\bar{\mu}(\mathbf{M}, \bar{\mathbf{Z}}) q_u^1(\mathbf{M}_u, \mathbf{M}_u^*)} \\ &= 1 \wedge \frac{\widehat{Z}_u^{M_u^*} p(\mathbf{M}_{[u]}^*) q_u^1(\mathbf{M}_u^*, \mathbf{M}_u)}{\widehat{Z}_u^{M_u} p(\mathbf{M}) q_u^1(\mathbf{M}_u, \mathbf{M}_u^*)}. \end{aligned}$$

For simplicity we have assumed that $q_u^1(\mathbf{M}_u, \mathbf{M}_u^*)$ is independent of the remaining state variables, but this is not necessary for its correctness and proposals which depend upon their values could easily be implemented. Note that the augmentation of the state space ensures that $\widehat{Z}_u^{M_u^*}$ exists and is available.

Similarly, consider a proposal Q_u^2 which refreshes the augmenting variables associated with node u :

$$\begin{aligned} & Q_u^2((\mathbf{M}, \bar{\mathbf{Z}}), (\mathbf{M}^*, \bar{\mathbf{Z}}^*)) \\ &= \delta_{\mathbf{M}, \mathbf{M}^*} \prod_{M \in \mathcal{M}} \left(g((\widehat{Z}_u^M)^* | m) \prod_{v \neq u} \delta_{\widehat{Z}_v^M}((\widehat{Z}_v^M)^*) \right), \end{aligned}$$

for which the Metropolis-Hastings acceptance probability is simply

$$1 \wedge \frac{(\widehat{Z}_u^M)^*}{\widehat{Z}_u^{M_u}}$$

by exploiting the same cancellation as in the standard pseudo-marginal setting. Clearly moves which update only some of the augmenting variables associated with node u could be justified in the same way. The augmenting variable \widehat{Z}_u^M for any $M \in \mathcal{M}$ can be re-estimated, not be just the normalising constant associated with the current state of \mathbf{M}_u . This flexibility is the result of the extension of the state space.

Standard arguments allow the combination of moves of these types, and others, within mixtures or cycles to provide an irreducible chain allowing considerable flexibility. In particular, it is no longer necessary to sample a marginal likelihood estimate for every proposed move in the state space. In the numerical studies below, we evaluate the performance of using combinations of such proposals.

4.4 Approximations of the NWPM algorithm

We finish this section with a brief exploration of approximations of the exact pseudo-marginal algorithm described above with the aim of obtaining inference which is almost as good at a fraction of the computational cost by allowing for a small bias in those estimates.

Clearly, a significant portion of the computational load is used to compute the marginal likelihood; this is especially the case when using more sophisticated methods such as the SMC sampler. For the NWPM method, as presented in Algorithm 2, the pseudo-marginal MH Markov chain of length n requires $n|V|$ marginal likelihood estimates. This quickly becomes infeasible with large data sets such as PET images, which may contain up to 10^6 voxel time series to be analysed. It is, therefore, worthwhile to consider strategies that reduce the number of marginal likelihood estimates that are required.

The multiple-augmentation approach to the NWPM algorithm decouples the estimation of marginal likelihoods from moves within the state space. One can envisage, for example, updating the marginal likelihood estimates only every κ iterations for some integer κ . Taking the limiting case as $\kappa \rightarrow \infty$ and estimating the marginal likelihood once each model at each node suggests a cheap approximate scheme. Naturally, this approach leads to a Markov chain which is not irreducible on the extended space and which can no-longer be considered a pseudo-marginal algorithm.

Such an algorithm would simply sample each of the marginal likelihood of all the model orders only once, for each node, before starting the chain; then, follow Algorithm 2 otherwise, using only these initial single estimates. This amounts to making a stochastic approximation which will change the invariant distribution of the resulting Markov chain.

Doing leads to a (now marginal) MH Markov chain, that targets an approximation of the target density rather than the target density itself. As such, it would not be formally justified under the pseudo-marginal framework, and there would be a loss of accuracy with any subsequent results, since it is an approximation. Since we only do a single estimation of the marginal likelihoods we will refer to this method as the Node-wise Single-Estimation (NWSE) algorithm.

The biggest appeal of the NWSE is that, for a chain of the same length n , the sampler would only need to be used $|V||\mathcal{M}|$ times, reducing the costs by $n/|V||\mathcal{M}|$ times. Any reduction in accuracy and performance could be adjusted for by using some of the saved residual computational resources towards reducing the variance of the sampler marginal likelihood estimates.

This approximation would allow at least preliminary spatial analyses to be conducted with little additional computational cost beyond that required for the associated

mass-univariate analysis: if existing analysis has been done, and thus some estimates of the marginal likelihood have already been obtained then these can be readily used within the algorithm to incorporate spatial dependence.

The NWSE algorithm involves running an MCMC chain with the wrong invariant distribution: that in which the exact marginal likelihoods are replaced with the realisations obtained when they are sampled. It is of interest, therefore, to establish how different such a distribution is from the true posterior distribution. The following elementary proposition, although it certainly does not provide a formal justification of the NWSE algorithm itself, demonstrates that in the context of small discrete spaces the error introduced by approximating marginal likelihoods can be controlled under reasonable conditions.

Proposition 2 *For each model $M \in \mathcal{M}$, let $\widehat{Z}^M := \widehat{Z}(M)$ be a RV with expectation $Z^M := f(y|M)$, corresponding to the associated marginal likelihood, and variance $\sigma_M^2 < \sigma_*^2 < \infty$. If the $(\widehat{Z}^M)_{M \in \mathcal{M}}$ are mutually independent, then, letting $M^* = \arg \max\{Z^M\}_{M \in \mathcal{M}}$ which we assume to be unique the probability of selecting the correct model via maximisation of the marginal likelihood (equivalently, the posterior mode of the distribution over models given a uniform prior over \mathcal{M}) is at least:*

$$1 - (|\mathcal{M}| - 1) \left(1 + \frac{\Delta^2}{2\sigma_*^2}\right)^{-1},$$

where $\Delta = \min_{L \neq M^*} Z^{M^*} - Z^L$ and $\sigma_*^2 = \max_{L \in \mathcal{M}} \sigma_L^2$.

Proof Let $M^* = \arg \max\{Z^M\}_{M \in \mathcal{M}}$, then:

$$\begin{aligned} \mathbb{P}(\{\widehat{Z}^{M^*} > \widehat{Z}^L \forall L \neq M^*\}) &= 1 - \mathbb{P}(\cup_{L \neq M^*} \{\widehat{Z}^{M^*} \leq \widehat{Z}^L\}) \\ &\geq 1 - \sum_{L \neq M^*} \mathbb{P}(\{\widehat{Z}^{M^*} \leq \widehat{Z}^L\}) \\ &\geq 1 - (|\mathcal{M}| - 1) \left(1 - \min_{L \neq M^*} \mathbb{P}(\widehat{Z}^{M^*} > \widehat{Z}^L)\right), \end{aligned}$$

by the Union bound. Applying Cantelli’s inequality, noting that $Z^{M^*} - Z^L > 0$:

$$\begin{aligned} \mathbb{P}(\{\widehat{Z}^{M^*} > \widehat{Z}^L\}) &= \mathbb{P}(\{\widehat{Z}^{M^*} - \widehat{Z}^L - (Z^{M^*} - Z^L) > -(Z^{M^*} - Z^L)\}) \\ &\geq 1 - \frac{\sigma_{M^*}^2 + \sigma_L^2}{\sigma_{M^*}^2 + \sigma_L^2 + (Z^{M^*} - Z^L)^2} \\ &= 1 - \left(1 + \frac{(Z^{M^*} - Z^L)^2}{\sigma_{M^*}^2 + \sigma_L^2}\right)^{-1}. \end{aligned}$$

Combining these expressions yields:

$$\begin{aligned} \mathbb{P}(\{\widehat{Z}^{M^*} > \widehat{Z}^L \forall L \neq M^*\}) &\geq 1 - (|\mathcal{M}| - 1) \max_{L \neq M^*} \left[1 - 1 + \left(1 + \frac{(Z^{M^*} - Z^L)^2}{\sigma_{M^*}^2 + \sigma_L^2}\right)^{-1}\right] \\ &\geq 1 - (|\mathcal{M}| - 1) \left(1 + \frac{\min_{L \neq M^*} (Z^{M^*} - Z^L)^2}{2 \max_{L \in \mathcal{M}} \sigma_L^2}\right)^{-1} \\ &= 1 - (|\mathcal{M}| - 1) \left(1 + \frac{\Delta^2}{2\sigma_*^2}\right)^{-1}. \end{aligned}$$

□

This bound can be made arbitrarily close to 1 by choosing estimators with sufficiently small variance. As the variance of the normalising constant estimates needs to be assessed to allow pseudo-marginal algorithms to be tuned, it is reasonable to suppose that this information could be obtained in settings in which this type of method is used and hence that a reasonable degree of confidence can be obtained that the use of this approximation does not substantially influence the resulting inference. Naturally, this result suggests that if estimating each marginal likelihood once rather than for every algorithmic step as in a pseudo-marginal algorithm, a relatively small variance might be required to obtain good performance.

In the numerical study, presented below, we investigate, evaluate and compare the effects when using the NWSE algorithm in different settings.

5 Applications

In this section, the NWPM algorithm and some variants are applied to a simple toy model, in which the true marginal likelihood is known, and to a compartmental model for PET data with both simulated and real data.

The software implementation of these methods can be found in the R package `bayespetr`, available at: <https://github.com/dt448/bayespetr>. The proposed algorithms are somewhat computationally intensive and here we focus on rather short Markov chains, a situation which we expect to be most widely useful; in Appendix B we establish that running these chains for much longer does not materially change the resulting inference (as we have found to be the case in all settings which we have explored).

The 20×20 image displayed in Fig. 2 will be used to generate the ground truth Potts configuration for all simulated-data experiments. It was adapted from Bezener et al. (2018) and includes a range of spatial structures of varying complexity.




Fig. 2 Ground Truth Configuration: spatial structure used to produce ground truth Potts configuration for the simulated data studies. A 20×20 grid, with four homogeneous regions

The image is split into four regions, labelled R_0, \dots, R_3 ; all pixels within a given region have a common model order; the details vary between experiments and are given below.

Whenever the NWPM algorithms are used in the numerical studies below, M was initialised using the Gibbs a sampler targeting the prior (Potts) distribution. However, when analysing measured PET data in Sect. 5.3, the algorithms were initialised from the output of the spatially independent SMC sampler method.

In the empirical studies below, we investigated a simple proposal strategy based on the multiple augmentation space, as discussed in Sect. 4.3.

In brief, the proposal kernel Q^1 , together with Q^2 at every $\kappa \in \mathbb{N}$ graphical iteration, is used. Essentially, upon initialising the chain, only the change of the model at the node is proposed. After every κ -th graphical iterations, changes in the auxiliary variable or marginal likelihood estimates for every model order at every node is proposed to change. In this section, we will refer to this variant as the NWMA algorithm, with tuning parameter κ .

SMC Sampler: The algorithms considered require unbiased estimates of $\int f(y_v|\theta_v, M_v)p(\theta_v|M_v)d\theta_v$ at node $v \in V$, for each model order $M \in \mathcal{M}$. The SMC sampler can be used to compute these node-wise marginal likelihoods by targeting the parameter posterior distribution $\pi(\theta_v|y_v, M_v)$ using an annealing schedule.

More precisely, for each model order $M \in \mathcal{M}$, let the sequence of target distribution of the SMC sampler be $\{\pi_t\}_{t \geq 1}$; Here, we define

$$\pi_t(\theta_{v,t}|y_v, M_v) \propto p(\theta_{v,t}|M_v, y_v)f(y_v|\theta_{v,t}, M_v)^{\alpha(t/T)},$$

where the total number of intermediate distributions T and annealing scheme, $\alpha : [0, 1] \rightarrow [0, 1]$, may be different for each model order (Zhou et al. 2016, Algorithm 2). For example, the annealing scheme could be a simple fixed schedule such as $\alpha(t/T) : [0, 1] \rightarrow (t/T)^5$, which following Zhou et al. (2016), we refer to as the ‘‘Prior 5’’ annealing scheme. Alternatively, α_t can be determined adaptively based on metrics of the distance between intermediate distributions, such as the ESS (Jasra et al. 2011; Kong and Liu 1994) or CESS (Zhou et al. 2016).

For the numerical studies below, we use the SMC sampler. Pilot studies showed that Prior 5 and the CESS-adaptive annealing scheme produced similar results in these examples, thus the Prior 5 scheme is used for simplicity.

5.1 Toy model simulation studies

Toy Model: A simple toy model in which both the prior and model likelihood are normal, at every node, is used to validate the proposed methodology. The data-set comprises a 2-dimensional digital image, each pixel having a scalar intensity. The graph $G = (V, E)$ used to represent the spatial structure is a finite square lattice, i.e., it has vertex set $V \subset \mathbb{Z}^2$ with nodes $v = (v_1, v_2)^T$ and edge set $E = \{\{u, v\} : d(u, v) = 1\}$, where $d(u, v) = \sum_i |u_i - v_i|$ is the L^1 (Manhattan) distance.

In the notation of Sect. 4.1.2, given an image $Y = (y_v \in \mathbb{R} : v \in V)$, the model studied here is:

$$\begin{aligned} M &\sim \text{Potts}(J, G_Y), \\ \mu_v|M_v &\sim \mathcal{N}(\mu_0^{(M_v)}, \sigma_0^2) \text{ for all } v \in V, \\ Y_v|M_v, \mu_v &\sim \mathcal{N}(\mu_v, \sigma^2) \text{ for all } v \in V; \end{aligned}$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution of mean μ and variance $\sigma^2 > 0$.

At each pixel, v , we model the data point Y_v as a normal random variable with mean μ_v and variance σ^2 , which is assumed known and is common to all pixels. The prior over the latent variable μ_v is normal with mean $\mu_0^{(M_v)}$ determined entirely by the model order M_v and fixed variance σ_0^2 . As before, the model order configuration will be the random variable M with a Potts prior distribution.

The node-wise marginal likelihood can be straightforwardly evaluated:

$$\begin{aligned} f(y|M) &= \int_{\mathbb{R}} \mathcal{N}(\mu; \mu_0^{(M)}, \sigma_0^2)\mathcal{N}(y; \mu, \sigma^2)d\mu \\ &= \mathcal{N}(y; \mu_0^{(M)}, \sigma^2 + \sigma_0^2), \end{aligned}$$

for all $M \in \mathcal{M}$. This is trivial to compute since all these terms are known.

The spatial configuration of model order used to simulate the data was specified to represent simple spatial structures; the ground truth value of M was fixed to be the Potts configuration shown in Fig. 2.

Toy Model Ground Truth: Ground Truth Configuration, as shown in Fig. 2 was used to generate a ground truth model order image. For this toy model, this is a 20×20 Potts configuration with model order space $\mathcal{M} = \{A, B\}$. Region R_0 was fixed to model order $M_v = A$ and the remaining three regions R_1, R_2 and R_3 where fixed to model order $M_v = B$. At each node $v \in V$, the mean parameter $\mu_0^{(M_v)}$ will depend on M_v . In other words, the lighter pixels have nodes with hyper-parameter $\mu_0^{(A)} = +5$ and the darker pixels have nodes with hyper-parameter $\mu_0^{(B)} = -5$. Finally, we specified $\sigma_0^2 = 5^2$ and $\sigma^2 = 1^2$.


5.1.1 Simulation study 1: the coupling Constant

As discussed in Sect. 2.3, typically it is difficult to infer the coupling constant J of the Potts distribution (Møller et al. 2006; Moores et al. 2020). In this work, we make the assumption that J is known when using the NWPM algorithm — this is not too unreasonable as J is a parameter of the prior distribution. In this vein, here we evaluate the performance of the algorithm for various values of J .


The following design was used: An image data set of $20 \times 20 = 400$ pixels(nodes) from the toy model and model order ground truth as dictated by the Toy Model Ground Truth above, was generated. This simulated image was then analysed using the NWPM algorithm, for varying values of the coupling constant J . Coupling constants ranging from $J = 0$ to $J = 5$ were investigated. Based on pilot studies, an SMC sampler with $N = 50$ particles and $T = 80$ distributions was used within the NWPM algorithms. The pseudo-marginal Markov chain was ran for $n = 100$ iterations. As mentioned above, the pseudo-marginal MH Markov chain was initialised using a Gibbs sampler, with the respective coupling constant. This was done for 50 replicates of the NWPM Markov chain.

Model selection was carried out by selecting the modal model order marginally at each node $v \in V$. In other words, the modal state in which each $M_v^{(i)}$, for $i = 1, \dots, 100$ was selected for each pixel $v \in V$. The empirical averages, over the 50 replicates was calculated. The mean percentages of the nodes for which the correct model order was selected for the different variations of the NWPM algorithm is shown in Fig. 3a.

It is evident from looking at Fig. 3a the performance of the algorithm peaks for coupling constant $J = 0.4$. This is smaller than the critical threshold $J_{critical}$, as discussed in Sect. 2.3. This is noteworthy, as there is typically a sharp decrease in speed of mixing of the Markov chain above this



(a) Model selection performance for $20 \times 20=400$ pixel image.



(b) Model selection performance for $100 \times 100=10,000$ pixel image.

Fig. 3 Empirical average, and standard deviations (s.d.), of the percentages(%) of the total (400 or 10,000) pixels where the correct model order was selected, for different values of the coupling constant J . Model selection was done on simulated toy model data using NWPM with an SMC estimator (with $N = 50$, and $T = 80$) for the marginal likelihood. The configuration selected for estimation by looking at the percentage of time each node spent in each model order state. Both line plot shows the empirical average of the correct percentage of pixels. The lighter and darker region show the $2 \times$ s.d. and $1 \times$ s.d. error bands, respectively

critical value. In fact, the graph shows that the variance of the percentage of correctly selected model orders increases with the coupling constant; Importantly, there is a jump in the rate of increase between $J = 1$ and $J = 1.5$.

Given that the size of measured PET data will be considerably bigger, this experiment was repeated for a larger Toy Model image. Additionally, the model order space for compartmental models may be larger; For instance, in this paper we consider up to $M = 3$ -compartment models. Subsequently, a scaled version of the 20×20 image, with dimensions $100 \times 100 = 10,000$ pixels and model space $\mathcal{M} = \{C, D, E\}$. More specifically, the region corresponding to R_0 had mean parameter $\mu_0^{(C)} = +7$; Similarly, region R_1 had $\mu_0^{(D)} = 0$ and regions R_2 and R_3 had $\mu_0^{(E)} = -7$. The parameters σ_0^2 and σ were kept as above. The generated data

was the analysed using the same the methods and design. The results are shown in Fig. 3b.

Notably, we first see that the peak of the graph is flatter in comparison to Fig. 3a. Secondly, the performance of algorithm is better overall; For instance, the best performance is almost 100% compared to roughly 94% previously. Similarly, when looking at the performance corresponding to $J = 0.4$, it increases from 94% to 97%. This suggest that, in this context, a larger image size and model order space did not seem to impact the performance negatively. Similarly, there is evidence that a wider range of coupling constant values could be used in the setting with bigger images and larger model order space, to attain “optimal” model selection performance. A selection of model order outputs, for varying values of J , are given in Fig. 17.

When looking at these results, as presented in Fig. 3, an important question to consider here is how much of the performance difference is due to poor approximation of the posterior by the Markov chain, and how much to the posterior itself. It is reasonable to argue that below the critical value, reasonably good mixing is obtained and this probably reflects a good approximation of the posterior as indicated by the low variability. In contrast, above this critical value the results are largely dominated by the poor mixing of the Markov chain over this rather short period. We also note that the graph shows good performance for even this relatively short chain, for the appropriate coupling constant values.

Based on evaluating the performance for a range of coupling constant and due to the similar qualitative features of the different settings, in the interest of simplicity, for all the following studies we will fix $J = 0.4$. This lower value means that the risk of being too close to the critical value J_{critical} is avoided.

5.1.2 Simulation study 2: sample size vs chain length

Next, the computational trade-off between the number of particles used in the SMC sampler and the length of the pseudo-marginal MH Markov Chain was investigated. The same design as the study in Sect. 5.1.1 was used, with the exception of the Markov Chain length varying from $n = 50$ to $n = 200$, and the number of particles varying between $N = 50$ and $N = 200$. Here, we investigate the $20 \times 20 = 400$ pixel image.

Additionally, spatially independent model selection, using an SMC sampler to estimate the model evidence was also investigated. The NWSE approximation was also used to analyse the image. Similarly, the NWMA algorithm, as described above, with updates to the marginal likelihood estimates every $\kappa = 10$ graphical iterations was used. For this variant, an SMC sampler with $N = 200$, $T = 500$ was used. For these two, less resource-intensive, Markov chain meth-

ods (NWMA and NWSE), $n = 200$ graphical iterations were used.

Model selection was carried out in the same manner as described above (i.e. the modal model), and the percentage of correctly selected nodes was used as the metric of performance. This was done for 100 replicates, for each method used. Results are shown in Fig. 4.

These results suggest that, at least in this simple setting, the algorithm is not particularly sensitive to the allocation of computational resources. Importantly, we can see that the selected graph model/configuration stabilises fairly quickly for all the combination of tuning parameters. In particular, looking at the longer $n = 200$ chain in Fig. 4a, it is evident that the convergence is to around 92% of nodes having the correct model order selected when using the NWPM algorithm. This is considerably higher than the performance of the spatially independent SMC sampler, which was 79%; confirming that exploiting spatial structure can significantly improve estimation in this type of model. We also observe that the NWSE approximation method and the simple NWMA algorithm achieved almost identical results to the full NWPM algorithm: suggesting that, in this simple setting, there is enough signal to allow for the use of the approximation to achieve very similar results, with very little additional computational cost. Intermediate chain lengths not shown in the graph, i.e. $n = 75$ and $N = 134$, also produced similar results.

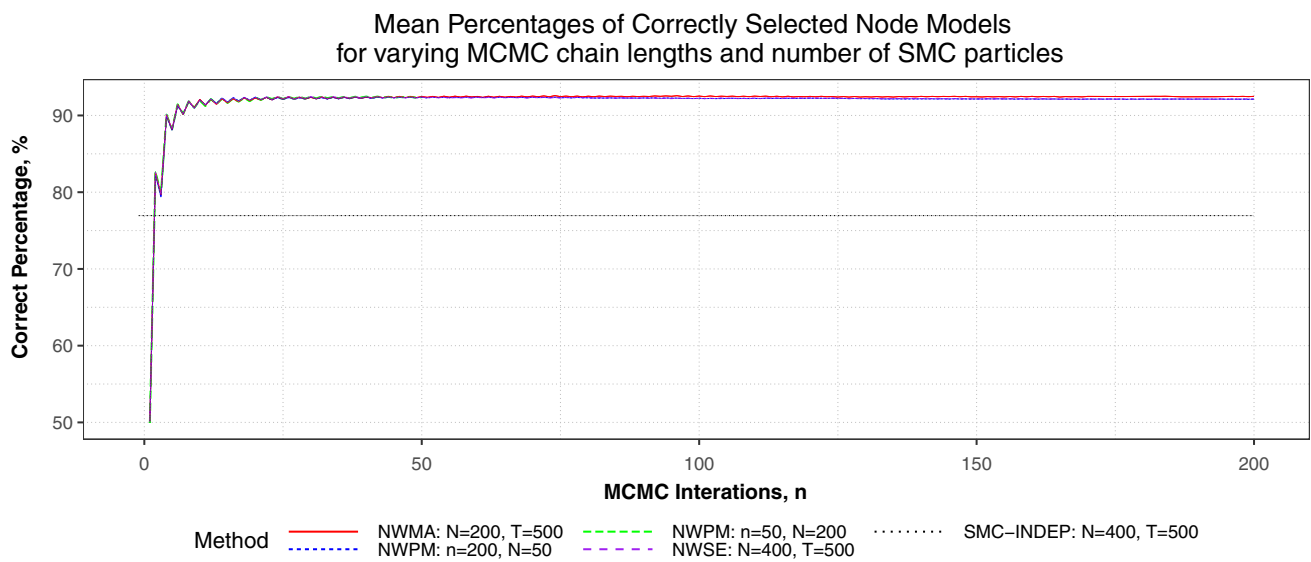
Figure 4b shows the average computational runtime for each method. The runtimes were computed by considering the average CPU time per iteration: That is, the total run time for each trail was recorded and then averaged over the total graphical iteration. We can draw similar conclusions as above, with the exception that the NWPM method with $N = 200$ particles seemed to take a longer time (in minutes) to reach stability.

The computational runtime for the NWMA method, for varying values of κ , was also studied: Results are shown in Fig. 15. As expected, the total runtime increase for smaller values of κ ; However, a higher refresh rate seems to slightly slow down (in minutes) how quickly the method attains its stable state.

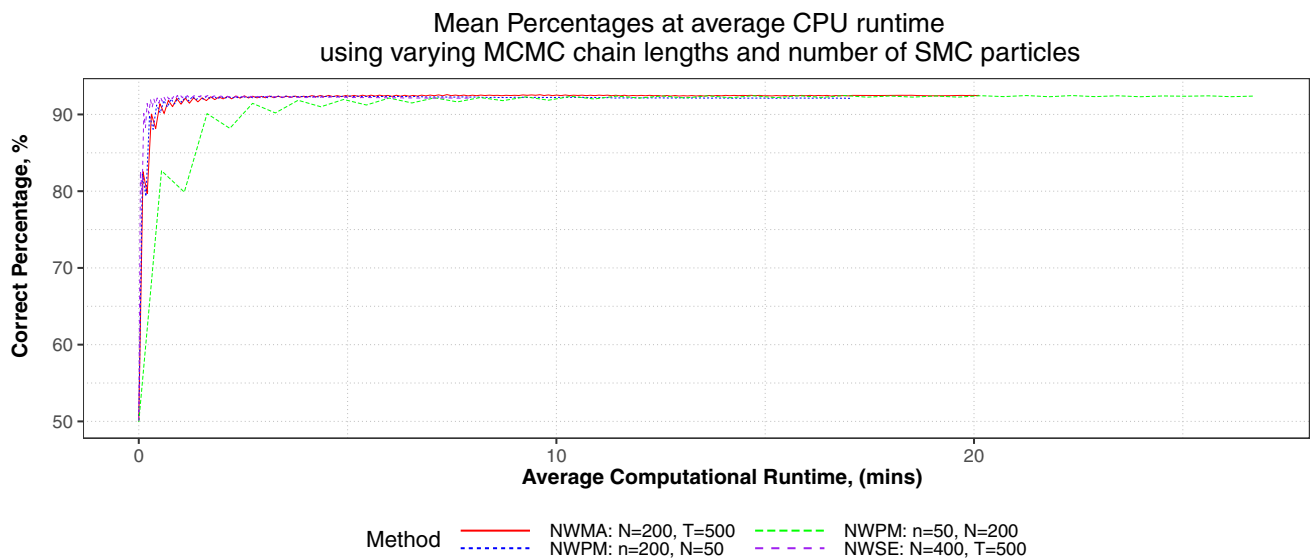
Finally, the long term behaviour for the NWSE and NWMA MCMC chains was also explored and results (i.e. that much longer chains lead to very similar conclusions) are shown in Appendix B, Fig. 13. Additionally, Fig. 11 shows an enhanced version of Fig. 4—allowing for better discernment of the small differences in performance.

5.2 Simulated PET data studies

Before we investigate the performance of the proposed algorithms on measured PET data, we verify and investigate tuning factors when applied to simulated PET data.



(a) Mean percentages at each graphical (MCMC) iteration.



(b) Mean computational runtime (averaged over each graphical iteration).

Fig. 4 Average percentages (%) of the whole image ($20 \times 20 = 400$ pixels, Toy Model) where the correct model order was selected at each iteration (shown in part a) of the MCMC chain, using NWPM, NWSE and NWMA for varying Markov Chain length, n , and number of par-

ticles, N , in the SMC sampler. The dashed line shows the average percentage when using spatially independent SMC model selection. The cumulative average over the graphical iterations $i = 1, \dots, n$ is shown. Average computational runtimes are shown in part b

5.2.1 Preliminaries

Consider a time series from a single voxel site of a PET image, denoted $y = (y_1, \dots, y_k)^T$. Recall the M -compartmental model for PET data (see (1) and (2), Sect. 2.2). In particular, note that the latent parameters $\phi_{1:M} = (\phi_1, \dots, \phi_M)$ and $\vartheta_{1:M} = (\vartheta_1, \dots, \vartheta_M)$ determined the dynamics of this model.

For the PET simulation study in this section, we simulate data from this collection of models with normally distributed errors as described in Sect. 2.2. For simplicity, we simulate data using a common parameter vector for every voxel associated with a given model order. These parameter values were taken from Jones et al. (1994) (and have subsequently been used for simulation studies by Gunn et al. (2002); Peng et al. (2008); Jiang et al. (2009) and Zhou et al. (2013)).

For the $M = 1$ and $M = 2$ compartmental models we use $\phi_1 = 4.9 \times 10^{-3} s^{-1}$, $\phi_2 = 1.8 \times 10^{-3} s^{-1}$, $\vartheta_1 = 5 \times 10^{-4} s^{-1}$ and $\vartheta_2 = 0.011 s^{-1}$; For the $M = 3$ -compartmental model we use $\phi_{1:3} = (4.4 \times 10^{-3} s^{-1}, 1 \times 10^{-4} s^{-1}, 1.4 \times 10^{-3} s^{-1})$ and $\vartheta_{1:3} = (4.5 \times 10^{-4} s^{-1}, 2.7 \times 10^{-3} s^{-1}, 1 \times 10^{-2} s^{-1})$. In each of these cases the volume of distributions is roughly $V_D \approx 10$.

Synthetic data was simulated using the same plasma input function and integration periods as in the measured PET data set studied in Sect. 5.3 and normally-distributed noise, with zero mean, was added. The variance was such that it was proportional to the true time activities divided by the length of the frame duration; and scaled such that the highest variance in the sequence is equal to the noise level. A noise level of 0.5, which lies within the representative range (between 0.01 and 1.28) of real data (Jiang et al. 2009), was used when simulating the 2-D PET image.

We employed vague uniform priors over ϕ and ϑ , which were constrained to lie within $[10^{-5}, 10^{-1}]$ and $[10^{-4}, 10^{-1}]$, respectively. These ranges are based on pilot and previous studies, and ensure that the parameters are physiologically meaningful (Cunningham and Jones 1993).

The precision parameter of the measure noise is denoted $\lambda = 1/\sigma^2$. A gamma distribution, with both parameters equal to 10^{-3} , was used as the prior for λ —a proper approximation to the improper Jeffrey's prior.

Details of the relevant distributions can be found in Appendix A.

PET Simulation Ground Truth As before, the ground truth of the model order configuration were based on spatial structure pattern of Ground Truth Configuration, shown in Fig. 2. In this PET simulation study, a 20×20 Potts configuration with model order space $\mathcal{M} = \{1, 2, 3\}$, representing the number of compartments was generated. Region R_0 corresponds to model order 2; region R_2 and R_3 to model order 1 and region R_1 to model order 3.

5.2.2 Simulation study 1: algorithm comparison

Pilot studies were initially carried out to evaluate the model selection performance and estimator variance of the SMC sampler. For the variance of the estimates, on agreement with numerical studies by Zhou et al. (2016): the numerical study showed that, there was a decrease linearly for higher values of N and similarly in T . However, beyond a certain threshold, the reduction were no longer useful. Similar results were seen in pilot studies for model selection performance, for 1-D data simulated from just the 2-compartment model with noise level 0.5, performance could not be increased beyond 65% regardless of any increase in tuning parameter values.

Next, the simulated 2-D PET image was analysed using the NWPM algorithm, and the effect of the pseudo-marginal MH chain was investigated. Based on the above pilot study,

an SMC sampler with the Prior 5 annealing scheme, $T = 400$ intermediate distributions, was used as the estimator of the normalising constants at all pixels for all models. These values, and the values of the number of particles N in the design below, were chosen as they allow for the variance of the likelihood estimates to be within a suitable range to allow for optimal scaling (Doucet et al. 2015). Chain lengths of $n = 50$, $n = 75$, $n = 100$ and $n = 200$ were investigated, with particles $N = 200$, $N = 134$, $N = 100$ and $N = 50$, respectively.

These, relatively smaller, graphical iteration lengths were chosen due to the computational overhead of this algorithm in this resource intensive setting. Note that, since model selection here uses the marginal modal state from a space of just 3 model orders—these shorter chains can still produce reliable results.

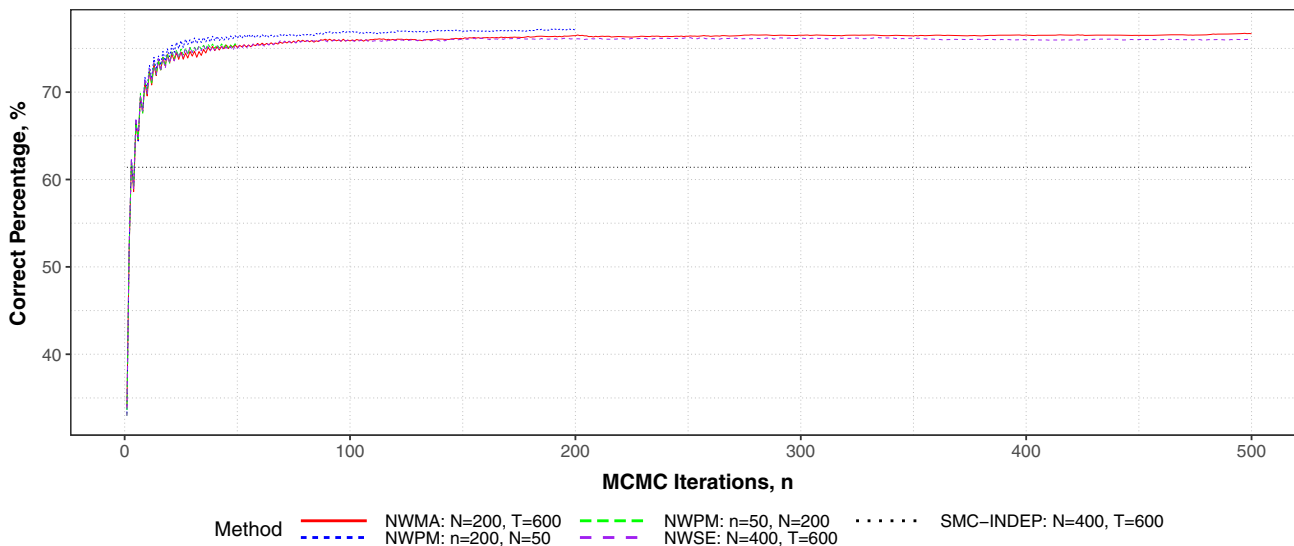
Additionally, the image was also analysed using spatially independent SMC sampler and the NWSE approximation; using an SMC sampler with $N = 400$ particles and $T = 600$. The NWMA algorithm, with SMC sampler using $N = 200$ and $T = 600$, at every $\kappa = 50$ iterations, was also evaluated. Given their significantly reduced computational costs, for the two pseudo-marginal methods, $n = 500$ iterations were used.

Following Algorithm 2, all the chains were initialised using the Gibbs sampler to target the prior distribution. 30 independent replicates of each algorithm were used to obtain the results shown in Fig. 5 and Table 1 (performance for the intermediate ranges of tuning parameters of the NWPM algorithm is shown in Appendix B, Fig. 16).

Looking at Fig. 5a, the results show notable improvements in model selection performance when using the different variants of the NWPM algorithms compared to spatially independent model selection using the SMC sampler. Analysing the simulated 2-D PET image using the spatially independent method gave 61.4% pixels correctly selected, compared to the NWPM algorithms all giving around 76%, even for as short as $n = 50$ graphical iterations. On average, this is a 23.4% increase in the number of pixels selected correctly when using the proposed NWPM approaches compared to the spatially independent SMC sampler method. On large data-sets, like PET images, this is a considerable number of pixels. Furthermore, since this is spatially dependent model selection, the improvement in correctly selected pixels when compared to SMC independent will be towards revealing the underlying spatial structures. These results suggest that incorporating spatial dependence does results in better model selection performance.

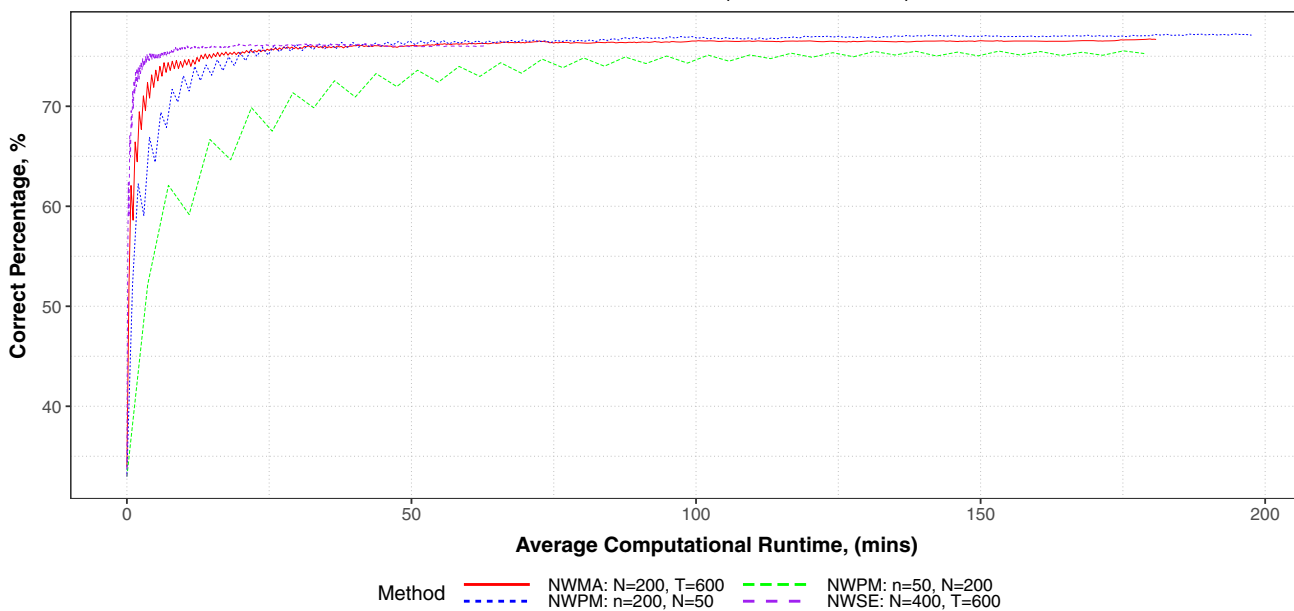
Figure 5 also shows that the performance of the NWSE approximation and the NWMA variant both provide similar performance to the full NWPM algorithm. Pragmatically, this is very promising as it suggests that even with an approximation current state of the art methods can be outperformed with very little additional costs. Alternatively, rather than using an

Mean Percentages of Nodes Correctly Selected for Simulated PET data(noise level = 0.5) using different methods.



(a) Mean percentages at each graphical (MCMC) iteration.

Mean Percentages at average computational time using varying MCMC chain lengths and number of SMC particles for Simulated PET data(noise level = 0.5)



(b) Mean computational runtime (averaged over each graphical iteration).

Fig. 5 Average percentages(%) of the whole image ($20 \times 20 = 400$ pixels, simulated PET data) where the correct model order was selected at each iteration of the pseudo-marginal MH chain, using: NWPM for varying Markov Chain length, n , and number of particles, N , in the SMC sampler; NWSE with $N = 400$ particles and $T = 600$ distribu-

tions and NWMA using SMC sampler with $N = 200$ and $T = 600$ at every $\kappa = 50$. The dashed line shows the average percentage when using spatially independent SMC($N = 400, T = 600$) model selection

Table 1 Average RMSE (and s.e.) of V_D estimates for simulated PET data

Method	Model Averaging V_D	Posterior V_D
NWPM $n = 50$	0.1319 (0.0061)	0.1321 (0.0061)
NWPM $n = 75$	0.1321 (0.0061)	0.1323 (0.0061)
NWPM $n = 100$	0.1320 (0.0061)	0.1323 (0.0061)
NWMA	0.1319 (0.0061)	0.1323 (0.0061)
NWSE	0.1322 (0.0061)	0.1324 (0.0061)
INDEP-SMC	0.1341 (0.0066)	0.1341 (0.0063)

approximation, algorithms based on the multiple augmented state space with significantly reduced computational overhead can be used. For example, the simple NWMA algorithm studied here. Similar to the toy model, there does not seem to be any notable difference when using higher chain length and lower particles. Additionally, there is only a small increase in performance when using more a computationally intensive algorithm, on average. That is, the full NWPM performs slightly better than the NWMA variant; which in turn also has a minute increase in performance compared to the NWSE approximation.

Additionally, for the chains where higher iterations were used, we can see that this increase in results was maintained. To verify long term behaviour of the algorithms: long-run MCMC chains, of the less expensive NWSE and NWMA algorithms, are shown in Appendix B, Fig. 14. As before, an enhanced version of Fig. 5a is shown in Fig. 12.

As before, Fig. 5b shows the average computational runtime for each method. The runtimes were computed in the same manner as described in Sect. 5.1.2. Similarly, as seen in the aforementioned section, the NWPM method with $n = 200$ seems to be the slowest in regards to mixing (in minutes).

Next, Table 1 shows the average mean squared error (MSE) of the V_D estimates for each variant of computational method. The quantity was calculated by taking the average MSE over all the pixels in the 2-D PET image, and then over all the replicates (the empirical variance of these quantities are shown in the parenthesis). Although the non-spatial SMC sampler method has the highest empirical MSE, the NWPM algorithms, including all chain lengths and the SE variant, produce a noticeably smaller MSE compared to the spatially independent SMC method, and there is a small improvement when using the full pseudo-marginal NWPM, compared to its approximation, these differences all fall within the range that could be plausibly explained by random fluctuations and it is difficult to draw any meaningful conclusion.

5.3 Measured [^{11}C]diprenorphine PET image study

The simulation studies above have verified the effectiveness of the proposed methodology and allowed for the identifica-

tion of tuning parameters which lead to promising results in settings close to that of interest. We now turn our attention to real data sets.

Analysis of the tracer kinetics of dynamic PET images, that use the [^{11}C]diprenorphine tracer, are generally aimed towards quantification of opioid receptor concentrations in the brain. Such studies involve investigation of neurological disease which often involve change in brain receptor density.

Investigation of similar PET data sets, using different statistical approaches, have been previously analysed quantitatively by Gunn et al. (2002), Peng et al. (2008) and Jiang et al. (2009); However, these works focused on parameter estimation rather than model selection. Zhou et al. (2013, 2016) are more recent works that investigated model selection; Albeit, in these works, spatial independence is assumed.

5.3.1 PET data

A test-retest pair of measured dynamic PET data of the concentration of the tracer [^{11}C]diprenorphine in a normal subject, for which an arterial input functions were available, was analysed using the proposed and other methods. This data was initially measured as part of the dynamic PET study in Hammers et al. (2007). In that study, the effects of using lower concentration of tracer was investigated. In addition, the impact of using different slow frequency cutoffs within the analysis of the data was also studied.

PET data acquisition The subject underwent 95 min dynamic [^{11}C]diprenorphine PET baseline scan. [^{11}C]diprenorphine is a tracer, of the non-selective antagonist, which binds to neural opioid receptors. The PET scans were acquired in 3D mode on a Siemens/CTI ECAT EXACT3D PET camera. After image reconstruction, the spatial resolution is approximately 5mm. Technical details including the data correction and normalisation methods of the ECAT EXACT3D PET camera can be found in Spinks et al. (2000). Thirty seconds after the start of the scan, [^{11}C]diprenorphine was injected and flushed through the cannula as a smooth bolus over a total of 30s. The subject was injected with 185 MBq (Mega Becquerel) of [^{11}C]diprenorphine. The PET data was reconstructed using the re-projection algorithm (Kinahan and Rogers 1989) with ramp and Colsher filters cutoff at the Nyquist frequency. Reconstructed voxel size were $2.096\text{mm} \times 2.096\text{mm} \times 2.43\text{mm}$. Acquisition was performed in listmode (event-by-event) and scans were rebinned into 32 time frames of increasing duration. The lengths of these periods, in seconds, are: (variable length background time, 3×10 , 7×30 , 12×120 , 6×300 , 75 , 11×120 , 210 , 5×300 , 450 and 2×600). Frame-by-frame movement correction was performed on the PET images. Overall this resulted in images of dimensions/size $128 \times 128 \times 95$ voxels. This gives a total of 1, 556, 480 separate times series respectively to be analysed. In this study, we will analyse a

cross-section of each of the sagittal, coronal and transversal planes of the brain—the total number of times series studied here is, therefore, roughly 11,000.

Details of further PET data pre-processing can be found in (Hammers et al. 2007, Material and methods).

Decay Correction The PET data is not decay corrected and this should be accounted for within the compartmental model. Gunn et al. (2002) (similarly, Cunningham and Jones (1993) for Spectral Analysis), suggest that the slow frequency cutoff be close to the decay constant of the tracer used. In other words, given that the half life of [¹¹C] is 20.4min, its decay constant is 0.0005663s^{-1} ; Thus, the slow frequency cutoff, denoted θ_{\min} , is fixed close to this constant. Hammers et al. (2007) showed that actual parametric values depended heavily on the cutoff slow frequencies(between 0.0008 s^{-1} and 0.00063 s^{-1}). Importantly, the prior ranges of $\theta_{1:M}$ should be based upon θ_{\min} , and subsequently this should be taken into consideration when computing the volume of distribution V_D . Following Zhou et al. (2016), in particular see Zhou (2014), we will use the cutoff 0.0007s^{-1} .

5.3.2 Data analysis and performance evaluation

Normally-distributed errors have been found to poorly model data of this sort (Zhou et al. 2013). Instead we will use the t -distribution to model the additive error variable ϵ . The same gamma prior as used for λ in the simulation studies above, will be used for the scale parameter, τ . A uniform distribution over the interval $[0, 0.5)$ will be used as the prior for $1/\nu$, allowing the likelihood to vary from having a very heavy tail to being arbitrarily close to normality.

The measured data was analysed using the following methods: Firstly, SMC sampler estimates of the marginal likelihood of each pixel for the three models were computed. This strategy assumes spatial independence, so, as before, we will refer to this method as the independent SMC method. The Prior 5 annealing scheme with $N = 300$ particles and $T = 500$ distributions was used, these values were based on the studies above and numerical results reported by Zhou et al. (2016). Three cross-sectional slices (coronal, transverse and sagittal planes) of the 3-D PET image were analysed, the model order and volume of distribution parametric images are shown in Fig. 6.

Next, the NWPM algorithm was used to analyse the same three cross-sections. An SMC sampler with $N = 300$ particles and $T = 500$ distributions was used as the marginal likelihood estimator. Here, the chain was initialised from the output of the SMC spatially independent model selection above (rather than from the prior distribution, as in the simulation studies above). A pseudo-marginal MH chain of length $n = 75$ was generated. This chain length was based on viability of computation cost for this relatively large data set, and the simulation studies above. In addition, the sagittal

cross-sectional image was analysed using a longer $n = 200$ iterations, as discussed below. The results are shown in Fig. 7.

Finally, Figs. 8 and 9 show the model configuration output from analysing the measured PET data using the NWSE approximation and the NWMA variant, respectively. In this case, an SMC sampler with $N = 400$ particles and $T = 600$ distributions was used to make the marginal likelihood estimates. The pseudo-marginal chain was initialised using the output of these estimates, and these initial estimates were re-used within the NWSE algorithm for $n = 500$ iterations. For the NWMA method, the chain was initialised in the exact same manner, but the marginal estimates were updated every $\kappa = 100$, for a total of $n = 500$ iterations.

For further comparison, the sagittal cross-section was also analysed using NLS (non-linear least squares; Zhou et al. (2013)). This is shown in Fig. 10. The spatially independent SMC sample model selection and NWPM method (with a longer chain length $n = 200$) model selection are included to show the difference in the final model order parametric image output.

The volume of distribution images, specifically Figs. 6 and 7, show very similar performance for all methods when considering inference of V_D . However, Fig. 10 shows clear difference in the model order image when using NWPM compared to spatially independent SMC and NLS. It is evident that analysis using NWPM reveals more spatial structure: Albeit, using the SMC independent approach does reveal some of the underlying spatial information when compared to NLS; NWPM has a de-noising effect and improves the clarity of the roughly formed clusters seen in the spatially independent SMC method. In particular, the $M = 1$ and $M = 2$ compartment models are selected to model most of the structural clusters. In contrast to this there is a decrease in the number of $M = 3$ compartment models selected when compared to the non-spatial SMC method.

Note also that, by looking at Fig. 10 we see that the NWPM algorithm converges to a noisy version of the final output image even at low iterations as $n = 25$. Importantly, there seem to be only small changes at the higher iterations, suggesting that the chain reaches a stable output relatively quickly. Similar behaviour is seen in the simulation studies above, albeit in a smaller, simpler settings.

Furthermore, similar model selection output is seen when using the SE estimate approximation and the multiple augmentation variant, NWMA, of the NWPM algorithm. Comparing both Figs. 8 and 9 with Fig. 7, we see that the model order outputs are not identical, but qualitatively show many similarities. In particular, we see that there is a decrease in the total number of 3-compartment models selected when using the SE approximation or the NWMA variant, compared to the full NWPM. A similar effect is seen when comparing the independent SMC model selection with the NWPM. Notably, the NWSE and NWMA model selection outputs

Independent SMC Sampler Method




Fig. 6 Model order(MO) and volume of distribution(V_D) parametric image of measured PET data (sagittal(left), coronal(right) and transverse(middle) cross-sections), using spatially independent SMC sampler($N = 300, T = 500$) model selection

NWPM Full Pseudo–marginal Method




Fig. 7 Model order(MO) and volume of distribution(V_D) parametric image of measured PET data (sagittal(left), coronal(middle) and transverse(right) cross-sections), using the NWPM($n = 75$) algorithm

for model selection with spatial dependence. An SMC sampler($N = 300, T = 500$) was used for the marginal likelihood estimator

NWSE Single Estimate Approximation




Fig. 8 Model order(MO) and volume of distribution(V_D) parametric image of measured PET data (sagittal(left), coronal(middle) and transverse(right) cross-sections), using the NWSE approximation

method($n = 500$) for model selection with spatial dependence. An SMC sampler($N = 400, T = 600$) was used for the marginal likelihood estimator

are almost identical, differing only at a very small number of pixels.

6 Conclusion

In this work we have illustrated a novel computational method for effectively incorporating spatial dependence

when performing model selection at each location within a graph. This approach extends the pseudo-marginal method, in a number of directions, within the context that it has been developed and evaluated in. By exploiting the structure of the problem at hand, the developed methodology allows for considerable flexibility in the updating of state variables and marginal likelihood estimates relative to generic pseudo-marginal algorithms. The empirical study suggests

NWMA Multiple Augmentation Algorithm


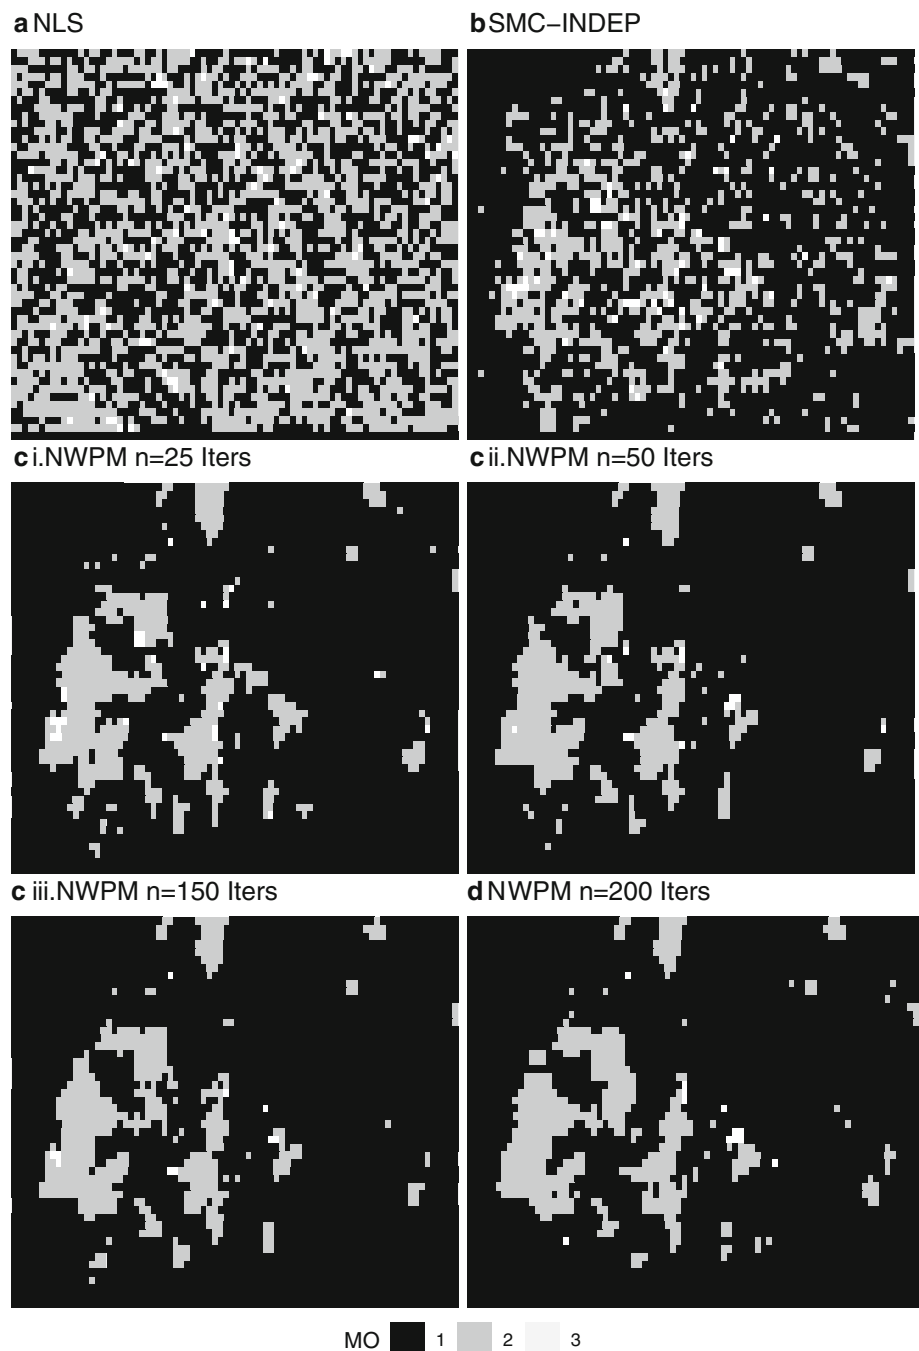


Fig. 9 Model order(MO) and volume of distribution(V_D) parametric image of measured PET data (sagittal(left), coronal(middle) and transverse(right) cross-sections), using the NWMA method($n =$

$500, \kappa = 100$) for model selection with spatial dependence. An SMC sampler($N = 400, T = 600$) was used for the marginal likelihood estimator

Fig. 10 Model selection for measured PET images, using: **a** spatially independent NLS and **b** SMC; **d** NWPM with spatial dependence incorporated using Potts model. Part **c** shows the progress of the NWPM chain. Model orders (MO) shown



that this approach can yield better inference than methods which impose assumptions of full spatial independence and, in particular, can reveal clearer spatial structure in the image of underlying model orders. As with pseudo-marginal algorithms, essentially any unbiased marginal likelihood estimators used in non-spatial analysis can be used within this algorithm.

This methodology was motivated by considering the problem of model selection in PET images. In such a setting both the amount of data and the complexity of the models

lead to substantial computational requirements. The images of model order provided by this approach supplement the information provided by volume of distribution and similar macroparameters in PET images. For instance, Fig. 10 (part d.) could be used to make the interpretation that the signal from some parts of the PET image require a higher model order compared to other parts. Specifically, this is apparent when comparing the posterior region of the brain image (where 2-compartmental models were selected) to the ante-

rior region of the brain (where 1-compartmental models were selected).

In order to reduce the computational cost, a multiple augmentation variant of the pseudo-marginal algorithm is introduced and a more extreme approximate form in which each marginal likelihood is estimated only once. The performance of these methods in the examples explored here is very encouraging: it suggests that where non-spatial analysis has been performed, spatial dependence can be incorporated with the existing estimates very easily and with little computational costs. Doing so carefully can result in similar performance to the full NWPM algorithm.

In summary, the NWPM algorithm, together with its approximations and extensions, enables a intuitive, natural approach to incorporating spatial dependence in realistic and complex settings; such as analysis of PET images. In particular, the methods extensions and approximations means that the subsequent performance and results of incorporating spatial information is very accessible and readily applicable.

Acknowledgements This work was supported by funding from the Engineering and Physical Sciences Research Council under grants EP/R034710/1 and EP/T004134/1, the Lloyd’s Register Foundation Programme on Data-centric Engineering at the Alan Turing Institute and EPSRC Studentship Project: 1943006. We thank Prof. Alexander Hammers and Prof. Sir John Aston for providing us with the PET data set used in this study.

Funding This work was supported by funding from the Engineering and Physical Sciences Research Council under Grants EP/R034710/1 and EP/T004134/1, the Lloyd’s Register Foundation Programme on Data-centric Engineering at the Alan Turing Institute and EPSRC Studentship Project: 1943006.

Availability of data and materials No novel data was produced. The PET data was obtained from Prof. John A. D. Aston.

Code availability The software implementation of the methods described in this paper can be found in the R package `bayespetr`, available at: <https://github.com/dt448/bayespetr>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A PET model

The posterior distribution for the PET compartmental models with normally-distributed errors is as follows: Let

$$t_i(\phi, \vartheta) := \frac{C_T(t_i; \phi, \vartheta)}{t_i - t_{i-1}},$$

the likelihood for data $y = (y_1, \dots, y_k)$ can be written:

$$\prod_{i=1}^k \sqrt{\frac{\lambda}{2\pi t_i(\phi, \vartheta)}} \exp \left\{ -\frac{\lambda}{2t_i(\phi, \vartheta)} (y_i - C_T(t_i; \phi, \vartheta))^2 \right\}.$$

The prior (joint) distributions over ϕ , ϑ and the precision parameter $\lambda = \frac{1}{\sigma^2}$ is given by:

$$\lambda^{\alpha-1} e^{-\beta\lambda} \prod_{j=1}^M I_{[\phi_j^a, \phi_j^b]} I_{[\vartheta_j^a, \vartheta_j^b]}$$

Here $\alpha = \beta = 10^{-3}$, the parameters of the prior distribution of λ . And ϕ_j^a and ϕ_j^b are the lower and upper bounds of the truncation interval of parameters ϕ_j and corresponding notation is used for ϑ_j .

For the t -distributed errors, the observation y_i has a t distribution with location $C_T(t_i)$, scale $\frac{t_i - t_{i-1}}{C_T(t_i)} \tau$, and degrees of freedom ν . The likelihood is,

$$\prod_{i=1}^k \left\{ \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})} \left(\frac{\tau}{t_i(\phi, \vartheta)\pi\nu} \right)^{\frac{1}{2}} \times \left(1 + \frac{\tau}{\nu t_i(\phi, \vartheta)} (y_i - C_T(t_i; \phi, \vartheta))^2 \right)^{-\frac{\nu+1}{2}} \right\}.$$

The prior density is given by:

$$\tau^{\alpha-1} e^{-\beta\tau} \times \frac{1}{\nu^2} \times I_{[a,b]} \left(\frac{1}{\nu} \right) \prod_{i=1}^M I_{[\phi_i^a, \phi_i^b]}(\phi_i) I_{[\vartheta_i^a, \vartheta_i^b]}(\vartheta_i),$$

where $\alpha = \beta = 10^{-3}$, the parameters of prior distribution of τ ; $a = 0$ and $b = 0.5$.

Appendix B MCMC Traces for studies and long-run chains

See Figs. 11, 12, 13, 14, 15 and 16.

Mean Percentages of Correctly Selected Node Models
for varying MCMC chain lengths and number of SMC particles
(iterations $n=10$ on wards)

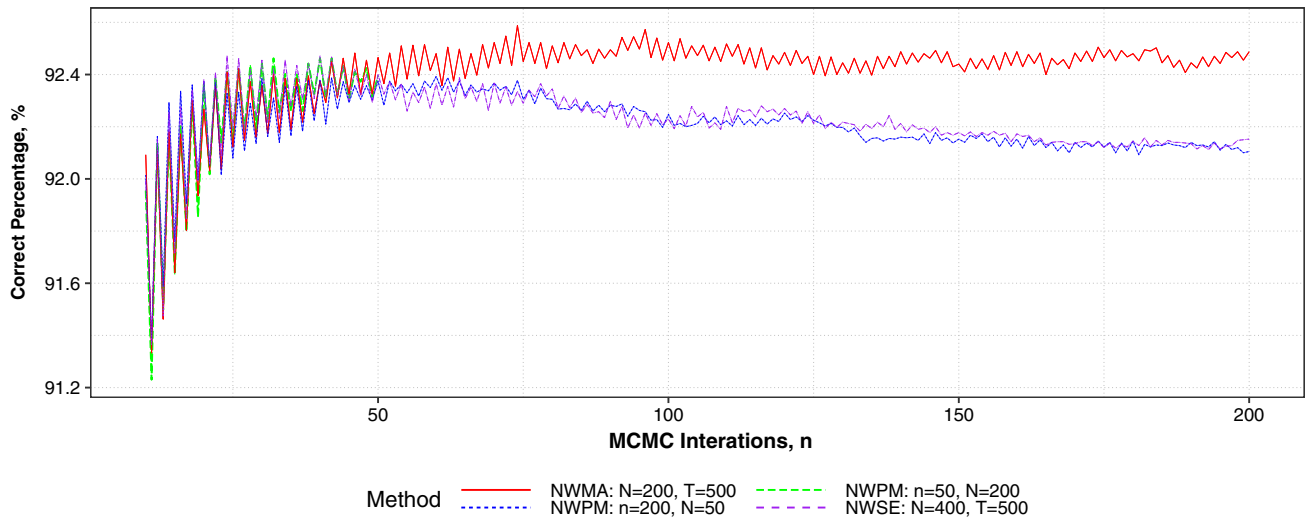


Fig. 11 Average percentages(%) of the whole toy model image ($20 \times 20 = 400$ pixels) where the correct model order was selected at each iteration of the MCMC chain. MCMC traces, of each method, for graphical iterations $n = 10$ on wards are shown

Mean Percentages of Nodes Correctly Selected for Simulated PET data(noise level = 0.5)
using different methods
(iterations $n=10$ on wards).

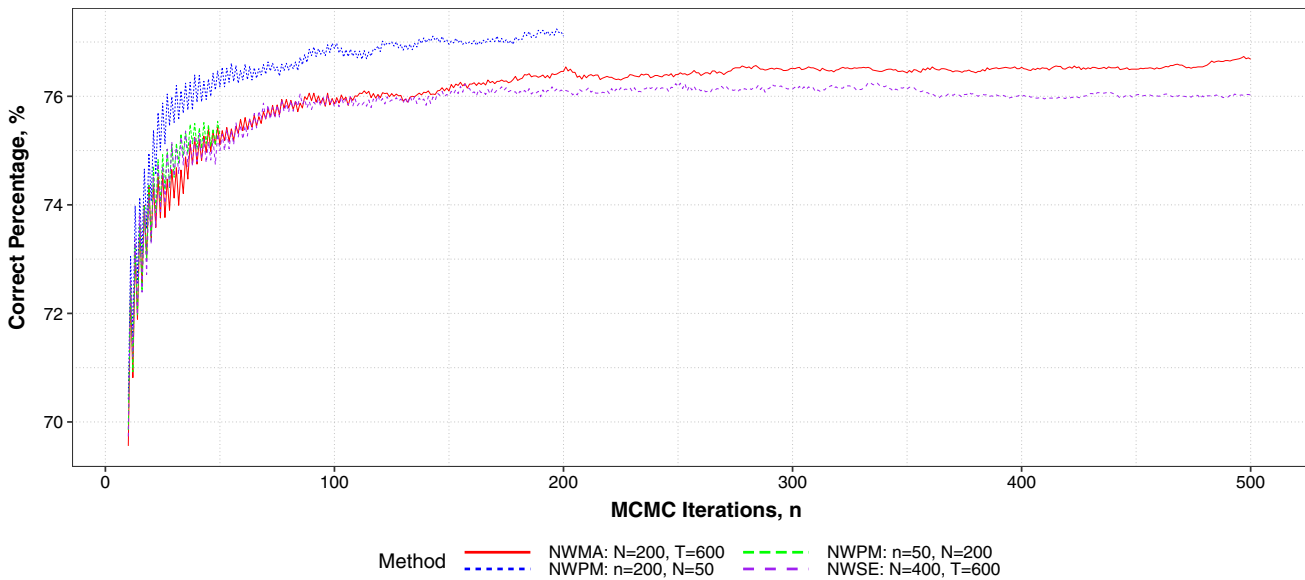


Fig. 12 Average percentages(%) of the whole image ($20 \times 20 = 400$ pixels) where the correct model order was selected at each iteration of the pseudo-marginal MH chain. MCMC traces, of each method, for graphical iterations $n = 10$ and on wards are shown




Fig. 13 MCMC traces of the percentage of correctly selected nodes, for the toy model. The NWSE and NWMA (refreshing marginal estimates at every $\kappa = 100$ graphical iterations). The SMC sampler with $N = 400$, $T = 500$ was used as the marginal likelihood estimator for both methods




Fig. 14 MCMC traces of the percentage of correctly selected nodes, for simulated 2-D PET image, using the NWMA (refreshing marginal estimates at every $\kappa = 1000$ graphical iterations) and NWSE methods. For both methods, the SMC sampler with $N = 200$, $T = 400$ was used

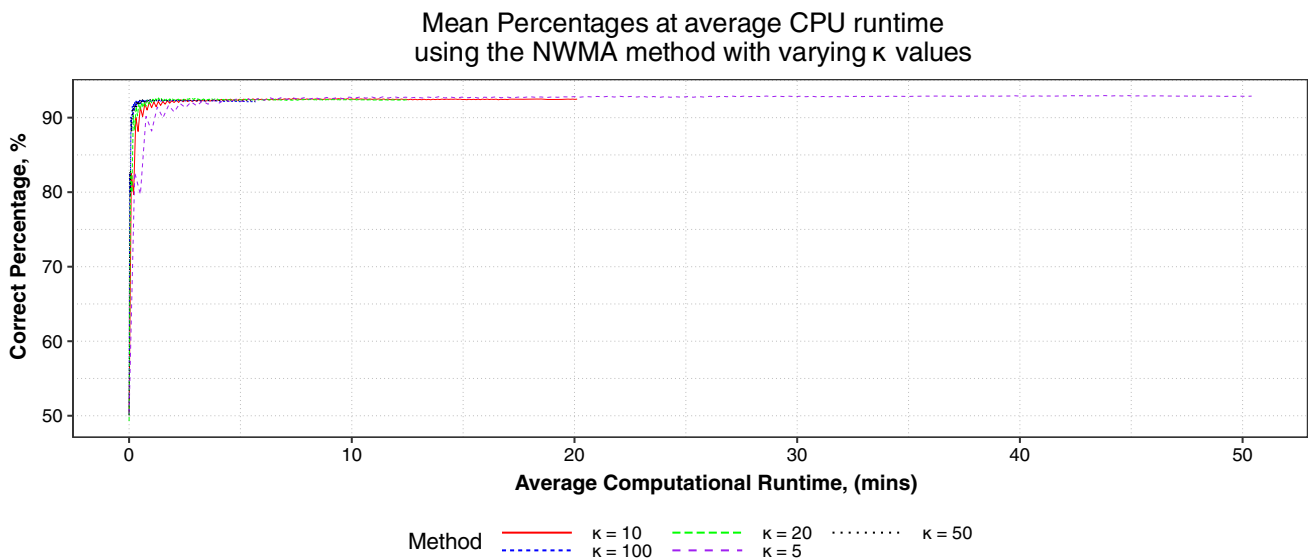


Fig. 15 Mean computational runtime (averaged over each graphical iteration), when analysing the Toy Model 20×20 image. The NWMA method (with parameters $n = 200$, $N = 200$, $T = 500$) was used, with κ varying from $\kappa = 5$ and $\kappa = 100$




Fig. 16 Average percentages(%) of the whole image where the correct model order was selected at each iteration of MH chain, using the NWPM algorithm for varying Markov Chain length, n , and number

of particles, N , in the SMC sampler. The dashed line shows the average percentage when using spatially independent SMC($N = 400, T = 600$) model selection

Appendix C Model order outputs for Toy model

See Fig. 17.




Fig. 17 Model order (MO) outputs for the 100×100 Toy Model simulated image, for varying coupling constant values J

References

- Akaike, H.: Information theory and an extension of the maximum likelihood principle. In: Parzen, E., Tanabe, K., Kitagawa, G. (eds.) Selected Papers of Hirotugu Akaike, pp. 199–213. Springer, New York (1973)
- Andrieu, C., Lee, A., Power, S., Wang, A.Q.: Comparison of Markov chains via weak Poincaré inequalities with application to pseudo-marginal MCMC. Technical Report 2112.05605, arXiv, (2021)
- Andrieu, C., Roberts, G.: The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.* **37**(2), 697–725 (2009)
- Andrieu, C., Vihola, M.: Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *Ann. Appl. Probab.* **25**(2), 1030–1077 (2015). <https://doi.org/10.1214/14-AAP1022>
- Andrieu, C., Vihola, M.: Establishing some order amongst exact approximations of MCMCs. *Ann. Appl. Probab.* (2016). <https://doi.org/10.1214/15-AAP1158>
- Andrieu, C., Doucet, A., Holenstein, R.: Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **72**(3), 269–342 (2010). <https://doi.org/10.1111/j.1467-9868.2009.00736.x>
- Beaumont, M.A.: Estimation of population growth or decline in genetically monitored populations. *Genetics* **164**(3), 1139–1160 (2003)
- Besag, J.: Spatial interaction and the statistical analysis of lattice systems. *J. R. Stat. Soc. Ser. B (Methodol.)* (1974). <https://doi.org/10.1111/j.2517-6161.1974.tb00999.x>
- Bezener, M., Hughes, J., Jones, G.: Bayesian spatiotemporal modeling using hierarchical spatial priors, with applications to functional magnetic resonance imaging (with discussion). *Bayesian Anal.* **13**(4), 1261–1313 (2018)
- Castellaro, M., Rizzo, G., Tonietto, M., Veronese, M., Turkheimer, F.E., Chappell, M.A., Bertoldo, A.: A Variational Bayesian inference method for parametric imaging of PET data. *NeuroImage* (2017). <https://doi.org/10.1016/j.neuroimage.2017.02.009>
- Chopin, N., Papaspiliopoulos, O.: An Introduction to Sequential Monte Carlo. Springer, Berlin (2020). <https://doi.org/10.1007/978-3-030-47845-2>
- Cunningham, V.J., Jones, T.: Spectral analysis of dynamic PET studies. *J. Cereb. Blood Flow Metab.* **13**(1), 15–23 (1993). <https://doi.org/10.1038/jcbfm.1993.5>
- Del Moral, P.: Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications, vol. 100. Springer, New York (2004)
- Del Moral, P., Doucet, A., Jasra, A.: Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B Stat. Methodol.* (2006). <https://doi.org/10.1111/j.1467-9868.2006.00553.x>
- Doucet, A., Johansen, A.M.: A tutorial on particle filtering and smoothing: fifteen years later. In: Crisan, D., Rozovsky, B. (eds.) Handbook of nonlinear filtering, pp. 656–704. Oxford University Press, (2011). ISBN 978-0199532902
- Doucet, A., Pitt, M.K., Deligiannidis, G., Kohn, R.: Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika* **102**(2), 295–313 (2015)
- Everitt, R.G.: Bayesian parameter estimation for latent Markov random fields and social networks. *J. Comput. Graph. Stat.* (2012). <https://doi.org/10.1080/10618600.2012.687493>
- Fan, Y., Emvalomenos, G., Grazian, C., Meikle, S.R.: PET-ABC: fully Bayesian likelihood-free inference for kinetic models. *Phys. Med. Biol.* (2021). <https://doi.org/10.1088/1361-1361/abfa37>
- Gelfand, A.E., Diggle, P., Guttorp, P., Fuentes, M.: Handbook of Spatial Statistics. CRC Press, Boca Raton (2010)
- Geman, D.: Random fields and inverse problems in imaging. In: Hennequin, P.-L. (ed.) *École d'Été de Probabilités de Saint-Flour XVIII*. Springer, Berlin, pp. 115–193, (1990). ISBN 978-3-540-46718-2. <https://doi.org/10.1007/bfb0103042>
- Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-6**(6), 721–741 (1984)
- Glauber, R.J.: Time-dependent statistics of the Ising model. *J. Math. Phys.* (1963). <https://doi.org/10.1063/1.1703954>
- Gunn, R.N., Gunn, S.R., Cunningham, V.J.: Positron emission tomography compartmental models. *J. Cereb. Blood Flow Metab.* **21**(6), 635–652 (2001)
- Gunn, R.N., Gunn, S.R., Turkheimer, F.E., Aston, J.A.D., Cunningham, V.J.: Positron emission tomography compartmental models: a basis pursuit strategy for kinetic modeling. *J. Cerebral Blood Flow Metab.* **22**, 1425–1439 (2002)
- Hammers, A., Asselin, M.C., Turkheimer, F.E., Hinz, R., Osman, S., Hotton, G., Brooks, D.J., Duncan, J.S., Koeppe, M.J.: Balancing bias, reliability, noise properties and the need for parametric maps in quantitative ligand PET: [¹¹C]diprenorphine test-retest data. *NeuroImage* **38**(1), 82–94 (2007). <https://doi.org/10.1016/j.neuroimage.2007.06.035>
- Hastings, W.: Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**(1), 97–109 (1970)
- Higdon, D.M.: Auxiliary variable methods for Markov chain Monte Carlo with applications. *J. Am. Stat. Assoc.* (1998). <https://doi.org/10.1080/01621459.1998.10473712>
- Hurn, M.A., Husby, O.K., Rue, H.: Spatial statistics and computational methods. In: *A Tutorial on Image Analysis*, pp. 87–141. Springer, New York (2003). https://doi.org/10.1007/978-0-387-21811-3_3
- Ising, E.: Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik* (1925). <https://doi.org/10.1007/BF02980577>
- Jasra, A., Stephens, D.A., Doucet, A., Tsagaris, T.: Inference for Lévy-Driven stochastic volatility models via adaptive sequential Monte Carlo. *Scand. J. Stat.* **38**(1), 1–22 (2011)
- Jiang, C.R., Aston, J.A., Wang, J.L.: Smoothing dynamic positron emission tomography time courses using functional principal components. *NeuroImage* **47**(1), 184–193 (2009). <https://doi.org/10.1016/j.neuroimage.2009.03.051>
- Jones, A.K., Cunningham, V.J., Ha-Kawa, S.-K., Fujiwara, T., Liyui, Q., Luthra, S.K., Ashburner, J., Osman, S., Jones, T.: Quantitation of [¹¹C]diprenorphine cerebral kinetics in man acquired by PET using presaturation, pulse-chase and tracer-only protocols. *J. Neurosci. Methods* **51**(2), 123–134 (1994)
- Kinahan, P.E., Rogers, J.G.: Analytic 3D image reconstruction using all detected events. *IEEE Trans. Nuclear Sci.* **36**(1), 964–968 (1989)
- Kong, A., Liu, J.S.: Sequential imputations and Bayesian missing data problems. *J. Am. Stat. Assoc.* **89**(425), 278–288 (1994). <https://doi.org/10.1080/01621459.1994.10476469>
- Matveev, V., Shrock, R.: Complex-temperature singularities in Potts models on the square lattice. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscipl. Top.* (1996). <https://doi.org/10.1103/PhysRevE.54.6174>
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A., Teller, E.J.: Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953)
- Møller, J., Pettitt, A.N., Reeves, R., Berthelsen, K.K.: An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika* (2006). <https://doi.org/10.1093/biomet/93.2.451>
- Moore, M., Nicholls, G.K., Pettitt, A.N., Mengersen, K.: Scalable Bayesian inference for the inverse temperature of a hidden Potts model. *Bayesian Anal.* (2020). <https://doi.org/10.1214/18-BA1130>
- Onsager, L.: Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Phys. Rev.* (1944). <https://doi.org/10.1103/PhysRev.65.117>
- Peng, J.Y., Aston, J.A., Gunn, R.N., Liou, C.Y., Ashburner, J.: Dynamic positron emission tomography data-driven analysis using sparse

- Bayesian learning. *IEEE Trans. Med. Imag.* (2008). <https://doi.org/10.1109/TMI.2008.922185>
- Phelps, M., Cherry, S., Dahlbom, M.: *PET: Physics, Instrumentation, and Scanners*. Springer, New York (2006)
- Potts, R.B.: Some generalized order-disorder transformations. *Math. Proc. Cambridge Philosophical Soc.* **48**(1), 106–109 (1952)
- Robert, C.P.: *The bayesian choice: from decision-theoretic foundations to computational implementation*. Springer Texts in Statistics, (2007). ISBN 9780387715988
- Sherlock, C.: Optimal scaling for the Pseudo-marginal random walk metropolis: insensitivity to the noise generating mechanism. *Methodol. Comput. Appl. Probab.* **18**(3), 869–884 (2016)
- Sherlock, C., Thiery, A.H., Roberts, G.O., Rosenthal, J.S.: On the efficiency of pseudo-marginal random walk metropolis algorithms. *Ann. Stat.* (2015). <https://doi.org/10.1214/14-AOS1278>
- Spinks, T.J., Jones, T., Bloomfield, P.M., Bailey, D.L., Miller, M., Hogg, D., Jones, W.F., Vaigneur, K., Reed, J., Young, J., Newport, D., Moyers, C., Casey, M.E., Nutt, R.: Physical characteristics of the ECAT EXACT3D positron tomograph. *Phys. Med. Biol.* **45**(9), 2601 (2000)
- Swendsen, R.H., Wang, J.-S.: Nonuniversal critical dynamics in Monte Carlo simulations. *Phys. Rev. Lett.* **58**, 86–88 (1987)
- Winkler, G.: *Image analysis, random fields and dynamic Monte Carlo methods*. Springer, Berlin (1995). <https://doi.org/10.1007/978-3-642-97522-6>
- Wu, F.Y.: The Potts model. *Rev. Modern Phys.* (1982). <https://doi.org/10.1103/RevModPhys.54.235>
- Zhou, Y.: vSMC: parallel sequential Monte Carlo in C++. *J. Stat. Softw.* **62**(9), 1–49 (2014)
- Zhou, Y., Huang, S.C., Bergsneider, M., Wong, D.F.: Improved parametric image generation using spatial-temporal analysis of dynamic PET studies. *NeuroImage* **15**(3), 697–707 (2002). <https://doi.org/10.1006/nimg.2001.1021>
- Zhou, Y., Aston, J.A., Johansen, A.M.: Bayesian model comparison for compartmental models with applications in positron emission tomography. *J. Appl. Stat.* **40**(5), 993–1016 (2013)
- Zhou, Y., Johansen, A.M., Aston, J.A.: Toward automatic model comparison: an adaptive sequential Monte Carlo approach. *J. Comput. Graph. Stat.* **25**(3), 701–726 (2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.