



<b>Title</b>	Model comparison for Gibbs random fields using noisy reversible jump Markov chain Monte Carlo
<b>Authors(s)</b>	Bouranis, Lampros, Friel, Nial, Maire, Florian
<b>Publication date</b>	2018-07-23
<b>Publication information</b>	Bouranis, Lampros, Nial Friel, and Florian Maire. "Model Comparison for Gibbs Random Fields Using Noisy Reversible Jump Markov Chain Monte Carlo." Elsevier, July 23, 2018. <a href="https://doi.org/10.1016/j.csda.2018.07.005">https://doi.org/10.1016/j.csda.2018.07.005</a> .
<b>Publisher</b>	Elsevier
<b>Item record/more information</b>	<a href="http://hdl.handle.net/10197/10402">http://hdl.handle.net/10197/10402</a>
<b>Publisher's statement</b>	This is the author's version of a work that was accepted for publication in Computational Statistics & Data Analysis. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in Computational Data Analysis (128, (2018)) <a href="https://doi.org/10.1016/j.csda.2018.07.005">https://doi.org/10.1016/j.csda.2018.07.005</a>
<b>Publisher's version (DOI)</b>	<a href="https://doi.org/10.1016/j.csda.2018.07.005">10.1016/j.csda.2018.07.005</a>

Downloaded 2026-05-02 01:15:10

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd\_oa)



© Some rights reserved. For more information

# Model comparison for Gibbs random fields using noisy reversible jump Markov chain Monte Carlo

Lampros Bouranis, Nial Friel\*, Florian Maire

*School of Mathematics and Statistics & Insight Centre for Data Analytics,  
University College Dublin, Ireland*

---

## Abstract

The reversible jump Markov chain Monte Carlo (RJMCMC) method offers an across-model simulation approach for Bayesian estimation and model comparison, by exploring the sampling space that consists of several models of possibly varying dimensions. A naive implementation of RJMCMC to models like Gibbs random fields suffers from computational difficulties: the posterior distribution for each model is termed doubly-intractable since computation of the likelihood function is rarely available. Consequently, it is simply impossible to simulate a transition of the Markov chain in the presence of likelihood intractability. A variant of RJMCMC is presented, called noisy RJMCMC, where the underlying transition kernel is replaced with an approximation based on unbiased estimators. Based on previous theoretical developments, convergence guarantees for the noisy RJMCMC algorithm are provided. The experiments show that the noisy RJMCMC algorithm can be much more efficient than other exact methods, provided that an estimator with controlled Monte Carlo variance is used, a fact which is in agreement with the theoretical analysis.

*Keywords:* Bayes factors, Intractable likelihoods, Markov random fields, Noisy MCMC.

---

## 1. Introduction

Model selection is a problem of great importance in statistical science. The aim is to choose which model among a set of possible ones best describes the data  $y \in \mathcal{Y}$ . From a Bayesian perspective, the prior beliefs for each model are reflected by a prior distribution and this information is then updated subjectively when data are observed. This step is typically carried out by calculating the marginal likelihood or *evidence* for each model, which is defined as the integrated likelihood with respect to the prior measure. In many cases, this quantity cannot be derived analytically and thus needs to be estimated.

This paper considers the problem of Bayesian model comparison of *doubly-intractable* distributions. The motivating application is Gibbs random fields (GRFs), which are discrete-valued Markov random fields where an intractable normalising constant that depends on the model parameters,  $z(\theta)$ , is present for the tractable un-normalised likelihood  $q(y | \theta)$ . The likelihood density, given a vector of parameters  $\theta \in \Theta \subseteq \mathbb{R}^d$  and a vector of statistics  $s(y) \in \mathcal{S} \subseteq \mathbb{R}_+^d$  that are

---

\*Corresponding author

*Email addresses:* lampros.bouranis@ucdconnect.ie (Lampros Bouranis), nial.friel@ucd.ie (Nial Friel), florian.maire@ucd.ie (Florian Maire)

sufficient for the likelihood, is

$$f(y | \theta) = \frac{q(y | \theta)}{z(\theta)} = \frac{\exp\{\theta^\top s(y)\}}{\sum_{y \in \mathcal{Y}} \exp\{\theta^\top s(y)\}}. \quad (1)$$

Posterior parameter estimation for GRFs, which have found applicability in areas like image analysis and disease mapping (Friel and Rue, 2007), genetic analysis (François et al., 2006) and social network analysis (Wasserman and Pattison, 1996), is termed a doubly-intractable problem because the normalisation terms of both the likelihood function and the posterior distribution  $\pi(\theta | y) \propto f(y | \theta)p(\theta)$  are intractable. Bayesian model comparison of such models has attracted the attention of researchers and several methods have been considered, relying on likelihood simulations (Friel, 2013; Caimo and Friel, 2013; Everitt et al., 2017a), approximations to the intractable likelihood (Bouranis et al., 2018) and likelihood-free simulation techniques like the Approximate Bayesian Computation (ABC) algorithm (Grelaud et al., 2009).

In this paper we explore trans-dimensional Markov chain Monte Carlo (MCMC) for GRFs, focusing on the direct approach of a single across-model Markov chain using the celebrated reversible jump MCMC (RJMCMC) technique (Green, 1995). The advantage of across-model approaches is that they avoid the need for computing the evidence for each competing model by treating the model indicator as a parameter, where the chain explores simultaneously the model set and the parameter space. In the context of GRFs, however, RJMCMC techniques simply cannot be implemented because the likelihood normalising constant  $z(\theta)$  cannot be computed point-wise. Below we present a summary of the main contributions of this paper.

A variant of the RJMCMC algorithm is developed, where the intractable ratio of normalising constants that the acceptance probability depends on is approximated by an unbiased estimator. The resulting algorithm falls in the *noisy MCMC* framework (Alquier et al., 2016) as it simulates a chain that is not invariant for the target distribution.

As pointed out in Alquier et al. (2016), noisy MCMC is connected to pseudo-marginal algorithms (Andrieu and Roberts, 2009), where an unbiased and positive estimate of the target density is required. In the presence of an intractable normalising constant, the pseudo-marginal approach requires an unbiased estimate of  $1/z(\theta)$ . However, the reciprocal  $1/\hat{z}(\theta)$  yields a biased approximation of  $1/z(\theta)$  and so is not directly applicable to inference using pseudo-marginal techniques. Recently, Lyne et al. (2015) addressed this bias with Russian Roulette sampling, which yields an asymptotically exact algorithm. The implementation of the algorithm is computationally expensive, however, creating difficulties when inferring the model parameters for GRFs. Non-negative unbiased estimators have also been studied in Jacob and Thiery (2015), which showed that finding such an estimate is very challenging.

Consequently, we do not pursue such a computational approach in this paper. Instead, motivated by the inefficiency of a standard RJMCMC algorithm, we develop a noisy RJMCMC sampler that targets an approximated posterior distribution, rather than the desired one. We extend the theoretical analysis of noisy MCMC algorithms proposed in Alquier et al. (2016) to trans-dimensional kernels, providing bounds on the total variation between the Markov chain of a noisy RJMCMC algorithm and a Markov chain with the desired target distribution under certain conditions.

We show that noisy RJMCMC algorithms are only useful when the estimator of the ratio of normalising constants has a small variance. Motivated by Gelman and Meng (1998), we propose a smoother transition path between different models and resort to an alternative estimator with lower

variance. We demonstrate empirically that this idea simultaneously: (i) improves the mixing of the RJ Markov chain and (ii) decreases the asymptotic bias between the exact RJ Markov chain and its noisy approximation. Finally, we construct efficient jump proposal distributions for random walk (RW) noisy RJMCMC, which could be useful in the context of a large number of nested competing models.

The outline of the article is as follows. Section 2 introduces the reader to basic concepts regarding Bayesian model comparison and to the reversible jump MCMC formulation. In Section 3 we discuss the extension of the reversible jump methodology to doubly-intractable posterior distributions and present theoretical properties and practical aspects of the RJMCMC samplers under such computational difficulties. Section 4 presents some proposal tuning strategies for noisy RJMCMC. In Section 5 we study the theoretical behavior of the noisy RJMCMC algorithm and derive convergence bounds. We investigate the performance of noisy RJMCMC with a detailed numerical study that focuses on social network analysis in Section 6. We conclude the paper in Section 7 with final remarks.

## 2. Preliminaries

Suppose a finite set of competing models  $\mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \dots\}$  are under consideration to describe the data  $y$ . In the Bayesian setting each model  $\mathcal{M}_m$ , where  $m \in \{1, 2, 3, \dots\}$ , is characterised by a likelihood function  $f(y | \theta_m, \mathcal{M}_m) = f_m(y | \theta_m) \propto q_m(y | \theta_m)$ , parameterised by an unknown parameter vector  $\theta_m \in \Theta_m$ . Each model is also associated with a prior distribution  $p(\theta_m | \mathcal{M}_m) = p_m(\theta_m)$ , used to express the beliefs about the parameter vector prior to observing the data  $y$ . The focus of interest in Bayesian inference for each competing model is the posterior distribution

$$\pi(\theta_m | y, \mathcal{M}_m) = \frac{f_m(y | \theta_m) p_m(\theta_m)}{\pi(y | \mathcal{M}_m)}. \quad (2)$$

The prior beliefs for each model are expressed through a prior distribution  $p(\mathcal{M}_m)$ , such that  $\sum_{m \in \mathcal{M}} p(\mathcal{M}_m) = 1$ .

### 2.1. Bayesian model comparison

The marginal likelihood or model *evidence* for model  $\mathcal{M}_m$  is

$$\pi(y | \mathcal{M}_m) = \int_{\Theta_m} f_m(y | \theta_m) p_m(\theta_m) d\theta_m$$

and is rarely analytically tractable. However, knowledge of the evidence is required for a quantitative discrimination between competing models with the posterior model probabilities. Using Bayes' theorem the posterior model probability for model  $\mathcal{M}_m$  is

$$\pi(\mathcal{M}_m | y) = \frac{\pi(y | \mathcal{M}_m) p(\mathcal{M}_m)}{\sum_{j=1}^{|\mathcal{M}|} \pi(y | \mathcal{M}_j) p(\mathcal{M}_j)}, \quad (3)$$

where  $|\mathcal{M}|$  is the cardinality of the model set. The probabilities  $\pi(\mathcal{M}_m | y)$  are treated as a measure of the uncertainty of model  $\mathcal{M}_m$ . Comparison of two competing models in the Bayesian setting is

performed through the Bayes factor

$$BF_{m,m'} = \frac{\pi(y | \mathcal{M}_m)}{\pi(y | \mathcal{M}_{m'})},$$

which provides evidence in favor of model  $\mathcal{M}_m$  compared with model  $\mathcal{M}_{m'}$ . Using (3), the Bayes factor can also be expressed as the ratio of the posterior model odds to the prior odds,

$$BF_{m,m'} = \frac{\pi(\mathcal{M}_m | y)}{\pi(\mathcal{M}_{m'} | y)} \frac{p(\mathcal{M}_m)}{p(\mathcal{M}_{m'})}.$$

A comprehensive review of Bayes factors is presented by Kass and Raftery (1995). There are at least two computational approaches to estimate Bayes factors: (i) within-model simulation, where the evidence is estimated separately for each model, see Friel and Wyse (2012) for a recent review on related methods; (ii) across-model simulation with trans-dimensional MCMC methods, which involves estimation of posterior model odds for chosen prior model odds. Reversible jump MCMC (Green, 1995) is a popular trans-dimensional MCMC method and is the focus of this work.

## 2.2. Reversible jump MCMC

Reversible jump MCMC (Green, 1995) generalises the Metropolis-Hastings (MH) algorithm (Metropolis et al., 1953; Hastings, 1970) to allow for sampling from a distribution on a union of spaces of possibly differing dimensions, permitting state-dependent choices of move types. Let  $\pi$  be a distribution over the general state space  $\mathcal{X}$ ,

$$\mathcal{X} = \bigcup_{m \in \mathcal{M}} (\{m\}, \Theta_m), \quad \mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \dots\}, \quad \Theta_m \subseteq \mathbb{R}^{d_m}.$$

The dimension of the parameter space  $\Theta_m$  of model  $\mathcal{M}_m$  is denoted by  $d_m$ . The target  $\pi$  is a joint posterior distribution on a model and a parameter,

$$\pi(\mathcal{M}_m, \theta_m | y) = \frac{f_m(y | \theta_m) p_m(\theta_m) p(\mathcal{M}_m)}{\sum_{k=1}^{|\mathcal{M}|} f_k(y | \theta_k) p_k(\theta_k) p(\mathcal{M}_k) d\theta_k}. \quad (4)$$

which can be factorised as the product of posterior model probabilities and model-specific parameter posteriors,

$$\pi(\mathcal{M}_m, \theta_m | y) = \pi(\mathcal{M}_m | y) \pi(\theta_m | y, \mathcal{M}_m). \quad (5)$$

To sample the model indicator and the model parameters jointly a Markov chain is constructed with state space  $\mathcal{X}$  and stationary distribution  $\pi(\mathcal{M}_m, \theta_m | y)$ . The state space  $\mathcal{X}$  is a finite union of subspaces of possibly varying dimensions. By marginalisation, we obtain the probability of being in subspace  $\Theta_m$ .

The reversible jump MCMC scheme allows for Metropolis-Hastings moves between states defined by  $x = (m, \theta)$  and  $x' = (m', \theta')$  that may have different dimensions  $d_m$  and  $d_{m'}$ , respectively. Since it is a MH-type algorithm it is  $\pi$ -reversible and thus  $\pi$ -invariant. As a consequence, the Markov chain simulated by RJMCMC produces samples  $(m, \theta_m) \sim \pi(\cdot | y)$ .

Below we formulate the RJMCMC algorithm in a way that generalises non trans-dimensional MCMC. This will ease the analysis of Section 5. Let  $h(\theta', m' | \theta, m) = \omega(m, m') T_{m,m'}(\theta, \theta')$  be the proposal distribution for the transition from  $(m, \theta_m)$  to  $(m', \theta_{m'})$ . We denote by  $\omega(m, m')$  the

probability of proposing a jump to model  $\mathcal{M}_{m'}$  when the chain is currently at model  $\mathcal{M}_m$  and by  $T_{m,m'}(\theta, \theta')$  the proposal for the parameter vector. For simplicity, we consider the case of RW updates when there are common parameters between the current and the proposed states. We define  $C_{m,m'}$  as the indicator of the common parameters between models  $\mathcal{M}_m, \mathcal{M}_{m'}$  so that  $\theta_{C_{m,m'}}$  is a sub-vector of model parameters and by  $\bar{C}_{m',m}$  the set of parameters that are in  $\mathcal{M}_{m'}$ , but not in  $\mathcal{M}_m$ . Then the proposed parameter vector is defined as  $\theta' := [\theta_{C_{m,m'}}, \theta'_{\bar{C}_{m',m}}] \in \mathbb{R}^{d_{m'}}$ . and the proposal distribution for the parameter vector takes the form

$$T_{m,m'}(\theta, \theta') = \begin{cases} g(\theta'_{\bar{C}_{m',m}} | \theta_{C_{m,m'}}) & \text{if } m \neq m', \\ w(\theta' | \theta) & \text{if } m = m', \end{cases} \quad (6)$$

where  $g(\cdot | \theta)$  and  $w(\cdot | \theta)$  are some proposal distributions. The RJMCMC scheme proceeds as in Algorithm 1. Green (1995, 2003) offers a representation in terms of random numbers and a Jacobian term to eliminate the apparent difficulty of moving between spaces of different dimensions.

The advantage of RJMCMC is that estimates of the posterior model probabilities are readily available along with parameter estimates for each competing model. Of course, if the Markov chain never visits a model because it is unlikely a-posteriori then the parameter estimates will not be available for that model. To assess support for the models under examination with the RJMCMC scheme, the output from the trans-dimensional chain is processed to calculate the Bayes factor. Assuming equal prior probabilities on models  $\mathcal{M}_m$  and  $\mathcal{M}_{m'}$ , this motivates the simple estimate of  $B_{m,m'}$  as  $F_m/F_{m'}$ , where  $F_m$  is the frequency of visits (out of a chain of length  $G$ ) to model  $\mathcal{M}_m$ .

---

**Algorithm 1** Reversible jump MCMC

---

- 1: Initialise  $(m_0, \theta_0)$ .
- 2: **for**  $n = 0, \dots, G - 1$  **do**
- 3:     Select a candidate model  $\mathcal{M}_{m'}$  with probability  $\omega(m_{(n)}, m')$ .
- 4:     Propose  $\theta' \sim T_{m_{(n)}, m'}(\theta_{(n)}, \theta')$ .
- 5:     Estimate the acceptance probability  $A_{m_{(n)}, m'}(\theta_{(n)}, \theta') = \min\{1, \rho_{m_{(n)}, m'}(\theta_{(n)}, \theta')\}$ ,

$$\rho_{m_{(n)}, m'}(\theta_{(n)}, \theta') = \frac{f_{m'}(y | \theta')}{f_{m_{(n)}}(y | \theta_{(n)})} \frac{p_{m'}(\theta')}{p_{m_{(n)}}(\theta_{(n)})} \frac{p(\mathcal{M}_{m'})}{p(\mathcal{M}_m)} \frac{\omega(m', m_{(n)})}{\omega(m_{(n)}, m')} \frac{T_{m', m_{(n)}}(\theta', \theta_{(n)})}{T_{m_{(n)}, m'}(\theta_{(n)}, \theta')}. \quad (7)$$

- 6:     Set  $(m_{(n+1)}, \theta_{(n+1)}) \leftarrow (m', \theta')$  with probability  $A$ , else  $(m_{(n+1)}, \theta_{(n+1)}) \leftarrow (m_{(n)}, \theta_{(n)})$ .
  - 7: **end for**
  - 8: **return**  $\{m_n, \theta_n\}_{n=1, \dots, G}$
- 

### 3. Reversible jump MCMC in the presence of likelihood intractability

The reversible jump algorithm presents an asymptotically (as the number of transitions  $G \rightarrow \infty$ ) exact MCMC method by targeting the posterior distribution of interest,  $\pi$ . When the posterior involves an intractable likelihood a naive implementation of RJMCMC for GRFs is challenging because it is simply not possible to simulate a transition of this exact chain.

Indeed, given the current state  $(m, \theta)$ , the proposed value of the next state of the chain,  $(m', \theta')$ , is accepted with probability  $A_{m,m'}(\theta, \theta') = \min\{1, \rho_{m,m'}(\theta, \theta')\}$ , where

$$\rho_{m,m'}(\theta, \theta') = \frac{q_{m'}(y | \theta') p_{m'}(\theta') p(\mathcal{M}_{m'}) \omega(m', m) T_{m',m}(\theta', \theta)}{q_m(y | \theta) p_m(\theta) p(\mathcal{M}_m) \omega(m, m') T_{m,m'}(\theta, \theta')} \times \frac{z_m(\theta)}{z_{m'}(\theta')}. \quad (8)$$

The acceptance probability is dependent on a ratio of intractable normalising constants and, thus, cannot be directly evaluated. GRF models belong to the exponential family and satisfy the identity

$$\frac{z_m(\theta)}{z_{m'}(\theta')} = \sum_{y \in \mathcal{Y}} \frac{q_m(y | \theta)}{q_{m'}(y | \theta')} \frac{q_{m'}(y | \theta')}{z_{m'}(\theta')} = \mathbb{E}_{y \sim f_{m'}(\cdot | \theta')} [\boldsymbol{w}(y)]. \quad (9)$$

This expectation can be estimated with importance sampling, where  $\boldsymbol{w}(y) = q_m(y | \theta) / q_{m'}(y | \theta')$  are the importance sampling weights.

Below we present variants of the RJMCMC algorithm that bypass the need to calculate the intractable ratio  $z_m(\theta) / z_{m'}(\theta')$  in (8), by replacing it with an unbiased estimator. It is important to note that using an estimator instead of the true ratio has significant consequences on the convergence of the Markov chain. In particular, the Markov chain is usually no longer asymptotically exact and in this context, such an algorithm is referred to as *noisy* MCMC. However, Alquier et al. (2016) have shown that the distribution of the noisy MCMC Markov chain can be made arbitrarily close to the desired stationary distribution, provided that an estimator with an arbitrarily small variance is available. In the following, we refer to as noisy RJMCMC the implementation of the RJMCMC algorithm where the ratio  $z_m(\theta) / z_{m'}(\theta')$  is estimated.

The efficiency of the noisy RJMCMC depends on the asymptotic approximation of  $\pi$ , an estimator of  $z_m(\theta) / z_{m'}(\theta')$  with small variance and the mixing property of the chain. The first and second points are discussed in the remainder of this Section through the presentation of different estimators of the ratio of normalising constants. The third point is considered at Section 4, where efficient proposal distributions are designed.

### 3.1. Reversible jump exchange algorithm

Murray et al. (2006) presented the exchange algorithm that allows inference for doubly-intractable distributions and circumvents the issue of intractability of (1). The exchange algorithm introduces an auxiliary variable  $y' \sim f_{m'}(\cdot | \theta')$  that is used to estimate the intractable ratio of normalising constants with a one-sample unbiased importance sampling estimator of (9).

Pseudo-marginal algorithms (Andrieu and Roberts, 2009) share the same principles with the exchange algorithm, replacing the intractable likelihood function with a positive and unbiased estimate that can be obtained using likelihood simulations or importance sampling. Like the exchange algorithm, pseudo-marginal MCMC defines an asymptotically exact MCMC algorithm.

Caimo and Friel (2014) devised a trans-dimensional extension of the exchange algorithm which is  $\pi$ -invariant and generalises the exchange algorithm of Murray et al. (2006) to trans-dimensional settings in the same way that RJMCMC generalises the Metropolis-Hastings algorithm. We refer to this instance as the *RJ exchange* when  $N = 1$  draws from the likelihood are used to estimate (9).

Despite yielding a Markov chain that is  $\pi$ -invariant, the one-sample unbiased importance sampling estimator usually has a large variance (see discussion below). Empirical results in Alquier et al. (2016) suggest that such an exchange algorithm is likely to be inefficient because

of slow mixing properties and propose taking multiple auxiliary draws  $N > 1$ . In the trans-dimensional setting, this motivates a noisy RJMCMC algorithm that uses  $N > 1$  likelihood draws or another estimator with small variance.

A graphical illustration that attempts to explain the inefficiency of the RJ exchange can be seen in Figure 1. This example involves transitions with a random walk proposal from ERG model  $\mathcal{M}_1$  to model  $\mathcal{M}_2$ , which we provide details of in Section 6.2. It is evident that the RJ exchange acceptance ratio will frequently underestimate the RJ MH acceptance ratio, making it more likely to reject the proposed move and affecting the convergence speed of the algorithm, particularly as the  $L^2$  norm  $\|\theta - \theta'\|_2$  increases.




Figure 1: In the context of a trans-dimensional move between 2 different exponential random graph (ERG) models (see Section 6), 10 samples  $(\theta, \theta')$  are drawn at random. For each of them we plot the MH acceptance rate (yellow spade) and the distribution of the exchange acceptance rate (violin plot). In practice the MH acceptance rate is not tractable but we use (10) with  $N = 5 \times 10^5$  to estimate it precisely. As the distance between  $\theta$  and  $\theta'$  increases, the acceptance ratio in the RJ exchange algorithm becomes more variable.

### 3.2. Noisy reversible jump MCMC

Estimating ratios of intractable normalising constants using (9) has been proposed in Alquier et al. (2016). A Monte Carlo estimator based on multiple auxiliary draws from the likelihood with respect to the proposed parameter,  $y'_1, y'_2, \dots, y'_N \sim f_{m'}(\cdot | \theta')$ , approximates (9) by

$$\frac{\widehat{z_m(\theta)}}{z_{m'}(\theta')} = \frac{1}{N} \sum_{i=1}^N \frac{q_m(y'_i | \theta)}{q_{m'}(y'_i | \theta')}. \quad (10)$$

For the remainder of the paper, we will refer to (10) as the standard importance sampling estimator (ISE). The special case where  $N = 1$  yields the RJ exchange algorithm. We now address the issue of estimator variance.

The work of Stoehr et al. (2017) that considered non trans-dimensional moves demonstrated that the quality of the ISE strongly decreases when the distance  $\|\theta - \theta'\|_2$  increases. This is even more involved when  $\theta$  and  $\theta'$  belong to different parameter spaces. Boland et al. (2017) proved that the variance of the estimator (10) is controlled when the distance between the parameters decreases. This result can be extended to jumps between models of different dimensions. For instance, in the case of an independence sampler the full parameter vector is updated. Consequently,  $\|\theta - \theta'\|_2$  can be large and so may be the variance of the estimator. This is a serious concern that may affect the performance of the algorithm, as mentioned at the beginning of Section 3.

To study the variance of the ISE in the trans-dimensional setting, we consider for simplicity the case of a random walk transition between two nested models where  $d_m < d_{m'}$  and define  $\dot{\theta} := [\theta_{C_{m,m'}}, 0_{d_{m'}-d_m}] \in \mathbb{R}^{d_{m'}}$ . A RW allows us to control the distance  $\|\dot{\theta} - \theta'\|_2$ , as it is simply equal to  $\left\| \frac{\theta'_{C_{m',m}}}{\bar{C}_{m',m}} \right\|_2$ , and to apply Proposition 1 below.

**Proposition 1** (Proposition 1 in Boland et al. (2017)). *For any GRF model and  $(\dot{\theta}, \theta') \in \Theta_{m'} \times \Theta_{m'}$ , the variance (denoted by the variance operator  $\mathbb{V}$ ) of the estimator in (10) decreases when  $\|\dot{\theta} - \theta'\|_2 \downarrow 0$ , such that*

$$\mathbb{V}_{y'_1, y'_2, \dots, y'_N \sim f_{m'}(\cdot | \theta')} \left[ \frac{\widehat{z_{m'}(\dot{\theta})}}{z_{m'}(\theta')} \right] = O(\|\dot{\theta} - \theta'\|_2), \quad (11)$$

To further decrease the variance in (11) we consider an alternative estimator that takes advantage of the auxiliary draws from the likelihood, and one which has also been considered in previous works (Friel, 2013; Bouranis et al., 2018; Stoehr et al., 2017; Everitt et al., 2017b). Let us introduce an auxiliary variable  $t \in [0, 1]$  and discretise  $[0, 1]$  as  $0 = t_1 < \dots < t_L = 1$ , where  $L \in \mathbb{Z}_+$ . We remain in the case of nested model jumps and define the vector-valued mapping

$$\theta'(t) := \begin{cases} (1-t)\dot{\theta} + t\theta' & \text{if } d_m < d_{m'}, \\ (1-t)\theta + t\theta' & \text{if } d_m = d_{m'}. \end{cases}$$

The idea is to replace the ratio of normalising constants by a telescopic product of ratios of normalising constants taken at two consecutive points of the path  $t \in [0, 1]$ , so that  $(\theta'(t_j), \theta'(t_{j+1}))$  are close to each other. More precisely, the ratio of normalising constants can be written as

$$\frac{z_{m'}(\theta'(t_1))}{z_{m'}(\theta'(t_L))} = \prod_{j=1}^{L-1} \frac{z_{m'}(\theta'(t_j))}{z_{m'}(\theta'(t_{j+1}))}. \quad (12)$$

Any factor in the RHS of this equation is estimated as follows. The un-normalised likelihood  $q_{m'}(y | \theta'(t_j))$  can be considered as an importance distribution for the "target" distribution  $f_{m'}(y |$

$\theta'(t_j)$ ), noting that

$$\frac{z_{m'}(\theta'(t_j))}{z_{m'}(\theta'(t_{j+1}))} = \mathbb{E}_{y^{(j+1)} \sim f_{m'}(\cdot | \theta'(t_{j+1}))} \left[ \frac{q_{m'}(y^{(j+1)} | \theta'(t_j))}{q_{m'}(y^{(j+1)} | \theta'(t_{j+1}))} \right].$$

An unbiased importance sampling estimator of this expectation can be obtained by simulating multiple draws  $y_1^{(j+1)'}, \dots, y_S^{(j+1)' } \sim f_{m'}(\cdot | \theta'(t_{j+1}))$ , yielding

$$\frac{\widehat{z_{m'}(\theta'(t_j))}}{\widehat{z_{m'}(\theta'(t_{j+1}))}} = \frac{1}{S} \sum_{b=1}^S \frac{q_{m'}(y_b^{(j+1)' } | \theta'(t_j))}{q_{m'}(y_b^{(j+1)' } | \theta'(t_{j+1}))}.$$

Since the expectation of a product of independent random variables is the product of the expectations, then

$$\frac{\widehat{z_{m'}(\theta'(t_1))}}{\widehat{z_{m'}(\theta'(t_L))}} := \prod_{j=1}^{L-1} \frac{\widehat{z_{m'}(\theta'(t_j))}}{\widehat{z_{m'}(\theta'(t_{j+1}))}} = \prod_{j=1}^{L-1} \left[ \frac{1}{S} \sum_{b=1}^S \frac{q_{m'}(y_b^{(j+1)' } | \theta'(t_j))}{q_{m'}(y_b^{(j+1)' } | \theta'(t_{j+1}))} \right] \quad (13)$$

gives an unbiased estimator of the intractable ratio of normalising constants.

This telescopic product estimator (TPE) will come at an extra computational cost, similarly to (10). The larger the number of path steps  $L$  and the number of draws  $S$  (where  $N = (L - 1) \times S$ ), the more precise the estimate of the ratio of normalising constants at the expense of computational time. Nevertheless, empirical results in Section 6 show that this estimator can lead to improved performance results compared to the ISE for the same number of likelihood draws in the estimator. Algorithm 2 contains pseudo-code for the noisy RJMCMC algorithm.

---

**Algorithm 2** Noisy reversible jump MCMC

---

- 1: Initialise  $(m_0, \theta_0)$ .
- 2: **for**  $n = 0, \dots, G - 1$  **do**
- 3:     Select a candidate model  $\mathcal{M}_{m'}$  with probability  $\omega(m(n), m')$ .
- 4:     Propose  $\theta' \sim T_{m(n), m'}(\theta(n), \theta')$ .
- 5:     Draw  $y'_1, \dots, y'_N \sim f_{m'}(\cdot | \theta')$  depending on the type of estimator of the ratio  $z_{m(n)}(\theta(n)) / z_{m'}(\theta')$ , using (10) or (13).
- 6:     Estimate the acceptance probability  $\hat{A}_{m(n), m'}(\theta(n), \theta', y') = \min\{1, \hat{p}_{m(n), m'}(\theta(n), \theta', y')\}$ ,

$$\begin{aligned} \hat{p}_{m(n), m'}(\theta(n), \theta', y') &= \frac{q_{m'}(y | \theta')}{q_{m(n)}(y | \theta(n))} \frac{p_{m'}(\theta')}{p_{m(n)}(\theta(n))} \frac{p(\mathcal{M}_{m'})}{p(\mathcal{M}_{m(n)})} \frac{\omega(m', m(n))}{\omega(m(n), m')} \frac{T_{m', m(n)}(\theta', \theta(n))}{T_{m(n), m'}(\theta(n), \theta')} \\ &\quad \times \frac{\widehat{z_{m(n)}(\theta(n))}}{z_{m'}(\theta')}. \end{aligned} \quad (14)$$

- 7:     Set  $(m_{(n+1)}, \theta_{(n+1)}) \leftarrow (m', \theta')$  with probability  $\hat{A}$ , else  $(m_{(n+1)}, \theta_{(n+1)}) \leftarrow (m(n), \theta(n))$ .
  - 8: **end for**
  - 9: **return**  $\{m_n, \theta_n\}_{n=1, \dots, G}$
-

It is worth noting that Karagiannis and Andrieu (2013) use a similar telescopic product, where the transition path is used to introduce a trans-dimensional Annealed importance sampling (AIS) estimator of the ratio  $\pi(m', \theta')/\pi(m, \theta)$  that appears in (7). Contrary to the noisy RJMCMC algorithm which uses an importance sampling estimator of the ratio of likelihood normalising constants, the proposed RJMCMC algorithm of Karagiannis and Andrieu (2013) does not target a similar posterior distribution as the noisy RJMCMC algorithm and the trans-dimensional AIS estimator is constructed in a different way.

#### 4. Choice of reversible jump proposal densities

All algorithms in this paper were run using a proposal distribution  $\omega(m, m')$  for the transition from  $(m, \theta_m)$  to  $(m', \theta_{m'})$  that was set to a discrete Uniform  $\mathcal{U}(1, |\mathcal{M}|)$ . Below we present the strategies that we followed in order to tune the proposal distribution for the parameter vector.

##### 4.1. Independent proposals

We adopted the auto-reversible jump (Auto-RJ) exchange algorithm that was proposed by Caimo and Friel (2013) to perform Bayesian model comparison of exponential random graph models in the area of Bayesian network analysis.

The Auto-RJ takes the form of an independence sampler, making use of a parametric approximation of the posterior distribution for each model in an attempt to bypass the issue of tuning the parameters of the jump proposal distributions and increase within-model acceptance rates. The Auto-RJ consists of pilot MCMC runs that are used to sample from the posterior distribution of each competing model with the exchange algorithm (Murray et al., 2006) and then to approximate the estimated posterior by Gaussian distributions determined by the first moments of each sample. The online step of the Auto-RJ sets  $T_{m,m'}(\theta, \theta')$  to this approximation of the posterior density for both the within- and between-model jumps.

The TPE in (13) is particularly suitable to the independence sampler, as it allows for "long" transitions (in terms of a larger  $L^2$  norm in this context) while having an estimator of the intractable ratio of normalising constants with small variance, compared to the ISE. On the other hand, the requirement of pilot MCMC runs reduces the appeal of this method when the number of models under investigation is large. Indeed, in such a case a within-model simulation approach may be as efficient. Additionally, the Auto-RJ is likely to fail when the posterior of at least one model is multi-modal.

##### 4.2. Random walk proposals

A natural alternative to the independence sampler comes from allowing the proposal  $(m', \theta')$  to depend on the current state  $(m, \theta)$ . Despite eliminating the need for pilot runs, an open question remains as to how optimally estimate the proposal variance parameter. An inefficient proposal mechanism will result in a Markov chain which slowly explores the state space and has high auto-correlation, increasing the asymptotic variance of Monte Carlo estimators. Such inefficiency can be caused by not proposing large moves away from the current state of the chain or by proposing moves with prohibitively small associated acceptance probabilities.

For the within-model Gaussian RW updates (see below), a proposal distribution with a variance-covariance matrix in the form of  $\Sigma_{m'} = \lambda_{m'}(\Omega_{m'} + C_{m'}^{-1})^{-1}$  was assumed, to account for possible correlations between the model parameters (Chib and Jeliazkov, 2001; Martin et al., 2011). The prior precision is denoted by  $\Omega_{m'}$  and  $\lambda_{m'} \in \mathbb{R}^+$  is a positive Metropolis tuning parameter. The

precision matrix  $C_{m'}^{-1}$  is set as the negative Hessian  $-\nabla_{\theta_{m'}}^2 \log f_{m'}(y | \theta_{m'})|_{\hat{\theta}_{m'}^{MLE}}$ . Estimation of the MLE for each ERG model in Section 6 was performed with the Monte Carlo Maximum Likelihood Estimation (MC-MLE) procedure proposed by Geyer and Thompson (1992). Alternative procedures exist, see Hunter and Handcock (2006). The gradient of the log-likelihood function can be written as

$$\begin{aligned} \nabla_{\theta_{m'}} \log f_{m'}(y | \theta_{m'}) &= s_{m'}(y) - \frac{\nabla_{\theta_{m'}} z_{m'}(\theta_{m'})}{z_{m'}(\theta_{m'})} \\ &= s_{m'}(y) - \frac{\sum_{y \in \mathcal{Y}} s_{m'}(y) \exp\{\theta_{m'}^\top s_{m'}(y)\}}{\sum_{y \in \mathcal{Y}} \exp\{\theta_{m'}^\top s_{m'}(y)\}} \\ &= s_{m'}(y) - \mathbb{E}_{y|\theta_{m'}}[s_{m'}(y)]. \end{aligned} \tag{15}$$

Then the Hessian matrix can be found using the identity

$$\begin{aligned} \nabla_{\theta_{m'}}^2 \log f_{m'}(y | \theta_{m'}) &= \nabla_{\theta_{m'}} \left[ -\frac{\nabla_{\theta_{m'}} z_{m'}(\theta_{m'})}{z_{m'}(\theta_{m'})} \right] \\ &= -\left\{ \mathbb{E}_{y|\theta_{m'}}[s_{m'}^2(y)] - \left[ \mathbb{E}_{y|\theta_{m'}}[s_{m'}(y)] \right]^2 \right\} \\ &= -\mathbb{V}_{y|\theta_{m'}}[s_{m'}(y)], \end{aligned} \tag{16}$$

where  $\mathbb{V}_{y|\theta_{m'}}[s_{m'}(y)]$  denotes the covariance matrix of the vector  $s_{m'}(y)$  with respect to  $f_{m'}(y | \theta_{m'})$ . For GRF models those Hessian matrices are intractable and so we resort to Monte Carlo sampling from  $f_{m'}(y | \theta_{m'})$  in order to estimate  $\mathbb{V}_{y|\theta_{m'}}[s_{m'}(y)]$ . This setting helped us reach a reasonable mixing rate within each model.

For the between-model moves, a popular choice in the implementation of the RJMCMC algorithm for nested cases is the second order method of Brooks et al. (2003). The second order method is based on a Taylor series expansion of the acceptance probability around certain canonical jumps, providing a fully automated framework for achieving local adaptation of the proposal density. It attempts to maximise the marginal acceptance probability  $A_{m,m'}(\theta, \theta)$  and not the extended acceptance probability  $A_{m,m'}(\theta, \theta, y')$  that uses the auxiliary data  $y'$  and so it is incompatible with the reversible jump exchange algorithm. While such a strategy would be sensible for noisy RJMCMC, it involves estimation of (15) and (16), which are unavailable in closed form for models like GRFs. The gradients can be numerically approximated with Monte Carlo simulation; this relies on repeated likelihood simulations, increasing considerably the computational cost per iteration of the noisy RJMCMC algorithm and so we do not pursue this strategy further.

Our implementation of the between-model moves is inspired by Ehlers and Brooks (2008), who proved that the full-conditional posterior distribution,  $\pi(\theta_{C_{m',m}}' | y, \theta_m, \mathcal{M}_{m'})$ , which is based upon the current state of the chain, is the optimal proposal distribution for the random vector  $\theta_{C_{m',m}}'$ . We tackled the intractability of the full-conditional by first considering a multivariate Gaussian distribution  $\mathcal{M}\mathcal{V}\mathcal{N}(\hat{\theta}_{m'}^{MLE}, \Sigma_{m'})$ , assuming diffuse prior distributions. Standard theory can be used to derive a mean  $\mu^*$  and covariance matrix  $\Sigma_{m'}^*$  and approximate the full-conditional distribution with a multivariate Gaussian distribution,  $\mathcal{M}\mathcal{V}\mathcal{N}\left(\theta_{C_{m',m}}'; \mu^*, \Sigma_{m'}^*\right)$ . To summarise, the

proposal distribution for our implementation of the noisy Gaussian RW RJMCMC is

$$T_{m,m'}(\theta, \theta') = \begin{cases} \mathcal{M}\mathcal{V}\mathcal{N}\left(\theta'_{\bar{c}_{m',m}}; \mu^*, \Sigma_{m'}^*\right) & \text{if } m \neq m', \\ \mathcal{M}\mathcal{V}\mathcal{N}\left(\theta'; \theta, \Sigma_{m'}\right) & \text{if } m = m'. \end{cases} \quad (17)$$

Finally, Karagiannis and Andrieu (2013) construct efficient RJMCMC transitions, relying on the assumption that the target distribution  $\pi(\mathcal{M}_m, \theta_m | y)$  is known up to a normalising constant. For the reasons explained at Section 3.2, we do not pursue that methodology further.

## 5. Theoretical guarantees for noisy reversible jump MCMC

The goal of this section is to extend the results of Alquier et al. (2016) to the trans-dimensional setting and then to apply these to the context of noisy RJMCMC with the different estimators developed earlier. To do so, we first present RJMCMC as a simple generalisation of the MH algorithm. Following Section 2.2, we consider again the general state space  $\mathcal{X}$  that defines the target distribution  $\pi$ . Our interest lies in constructing a Markov chain with transition kernel  $\mathcal{P}$ , with  $\pi$  as the invariant distribution.

Let us consider  $\pi$  as a probability measure on the compact set  $\mathcal{X}$ . Then  $\pi$  is an invariant distribution if

$$\int_{\mathcal{X}} \pi(dx) \mathcal{P}(x, dx') = \pi(dx).$$

A sufficient, but not necessary, condition is that the respective Markov chain is reversible, so that for all Borel sets  $\mathcal{B}, \mathcal{B}' \subset \mathcal{X}$

$$\int_{\mathcal{B}} \pi(dx) \mathcal{P}(x, \mathcal{B}') = \int_{\mathcal{B}'} \pi(dx') \mathcal{P}(x', \mathcal{B}). \quad (18)$$

Let  $\mathfrak{v}(x, dx')$  be a proposal measure for the move from the current state  $x = (m, \theta)$  to the proposed state  $x' = (m', \theta')$ , where  $x'$  is accepted with probability  $\alpha(x, x') = A_{m,m'}(\theta, \theta')$ . The transition kernel is given by

$$\mathcal{P}(x, \mathcal{B}') = \int_{\mathcal{B}'} \mathfrak{v}(x, dx') \alpha(x, x') + \mathcal{R}(x) \mathbb{1}(x \in \mathcal{B}'), \quad (19)$$

where

$$\mathcal{R}(x) = \int_{\mathcal{X}} \mathfrak{v}(x, dx') [1 - \alpha(x, x')]$$

is the probability of rejecting any proposed move, while at state  $x$ . By substituting (19) into (18), it is straightforward to show that

$$\int_{(x,x') \in \mathcal{B}, \mathcal{B}'} \pi(dx) \mathfrak{v}(x, dx') \alpha(x, x') = \int_{(x,x') \in \mathcal{B}, \mathcal{B}'} \pi(dx') \mathfrak{v}(x', dx) \alpha(x', x).$$

The above formulation encompasses standard MCMC.

We denote by  $\hat{\mathcal{P}}$  the transition kernel of the Markov chain resulting from the estimators (10) and (13), where the acceptance ratio  $\rho(x, x')$  in  $\alpha(x, x') = \min\{1, \rho(x, x')\}$  is replaced by an estimator  $\hat{\rho}(x, x', y')$ , yielding  $\hat{\alpha}(x, x', y') = \min\{1, \hat{\rho}(x, x', y')\}$ .

**Theorem 2** (Corollary 3.1 in Mitrophanov (2005)). *We assume that*

**(H1)** *The reversible jump MH Markov chain with transition kernel  $\mathcal{P}$  and acceptance probability  $\alpha$  is uniformly ergodic, for which it holds that*

$$\sup_{x_0 \in \mathcal{X}} \|\delta_{x_0} \mathcal{P}^n - \pi\| \leq Q \xi^n$$

for some constants  $Q < \infty$  and  $\xi < 1$ , where  $\|\cdot\|$  is the total variation distance and  $\delta$  is the Dirac delta function.

Then for any transition  $n \in \mathbb{N}$  and for any starting point  $x_0 \in \mathcal{X}$  it holds that

$$\|\delta_{x_0} \mathcal{P}^n - \delta_{x_0} \hat{\mathcal{P}}^n\| \leq \left( \phi + \frac{Q \xi^\phi}{1 - \xi} \right) \|\mathcal{P} - \hat{\mathcal{P}}\|,$$

where  $\|\mathcal{P} - \hat{\mathcal{P}}\| = \sup_{x \in \mathcal{X}} \|\delta_x \mathcal{P} - \delta_x \hat{\mathcal{P}}\|$  is the total variation measure between the two kernels and  $\phi = \left\lceil \frac{\log(1/Q)}{\log(\xi)} \right\rceil \leq n - 1$ .

An application of Theorem 2 is now provided when an approximation to the true transition kernel arises from Algorithm 2.

**Corollary 3.** *Let us assume that*

**(H1)** *The Markov chain with transition kernel  $P$  is uniformly ergodic holds.*

**(H2)** *There exists a function  $\gamma: \mathcal{X}^2 \rightarrow \mathbb{R}^+$ , such that for all  $(x, x') \in \mathcal{X}^2$ ,*

$$\mathbb{E}_{y' \sim f(\cdot|x')} |\hat{\rho}(x, x', y') - \rho(x, x')| \leq \gamma(x, x'). \quad (20)$$

Then for any  $n \in \mathbb{N}$  and for any starting point  $x_0 \in \mathcal{X}$  it holds that

$$\|\delta_{x_0} \mathcal{P}^n - \delta_{x_0} \hat{\mathcal{P}}^n\| \leq \left( \phi + \frac{Q \xi^\phi}{1 - \xi} \right) \sup_x \int dx' v(x, x') \gamma(x, x'),$$

where  $\phi = \left\lceil \frac{\log(1/Q)}{\log(\xi)} \right\rceil$ .

When the upper bound in (20) is uniformly bounded such that for all  $(x, x') \in \mathcal{X}^2$  it holds that  $\gamma(x, x') \leq \gamma < +\infty$ , then

$$\|\delta_{x_0} \mathcal{P}^n - \delta_{x_0} \hat{\mathcal{P}}^n\| \leq \gamma \left( \phi + \frac{Q \xi^\phi}{1 - \xi} \right).$$

Consequently, letting  $n \rightarrow \infty$  yields

$$\limsup_{n \rightarrow \infty} \|\delta_{x_0} \hat{\mathcal{P}}^n - \pi\| \leq \gamma \left( \phi + \frac{Q \xi^\phi}{1 - \xi} \right).$$

Below we make a series of assumptions that will help us show that the noisy RJMCMC algorithm will yield a Markov chain which will converge to the target posterior density as: (i)  $N \rightarrow \infty$ , if the ratio  $z_m(\theta)/z_{m'}(\theta')$  is estimated with (10) or (ii)  $L, S \rightarrow \infty$ , if the TPE in (13) is considered, instead.

(A1) For each prior over the model-specific parameters, there is a constant  $c_{p_m}$  such that

$$c_{p_m}^{-1} \leq p_m(\theta) \leq c_{p_m}, \forall m.$$

(A2) There is a constant  $c_{v(x,x')}$  such that  $c_{v(x,x')}^{-1} \leq v(x, x') = h(\theta', m' | \theta, m) \leq c_{v(x,x')}, \forall m, m'$ .

(A3) A prior probability  $0 < p(\mathcal{M}_m) \leq c_{\mathcal{M}_m} \leq 1, \forall m$  is assigned to each model.

(A4) For any  $\theta \in \Theta_m$  and  $\theta' \in \Theta_{m'}$ ,  $\mathbb{V}_{y' \sim f(\cdot | x')} \left[ \frac{q_m(y' | \theta)}{q_{m'}(y' | \theta')} \right] < \infty, \forall m, m'$ .

These assumptions are met when  $\Theta_m$  is a bounded set for each and every model in the set, such that

$$R_m := \sup_{\theta_m \in \Theta_m} \|\theta_m\|, \quad S_m := \sup_{y \in \mathcal{Y}} \|s_m(y)\|, \quad K_m := \exp\{R_m S_m\}$$

are finite, which gives  $0 < \exp\{-R_m S_m\} \leq q_m(y | \theta_m) \leq \exp\{R_m S_m\}$  for any  $\theta_m = \theta$  and  $y \in \mathcal{Y}$ , using the Cauchy-Schwartz inequality. Hence, we can set  $\mathbb{V}_{y' \sim f(\cdot | x')} \left[ \frac{q_m(y' | \theta)}{q_{m'}(y' | \theta')} \right] \leq K_m K_{m'}$ .

We acknowledge that Assumptions (A1) to (A4) are, in general, quite strong. If  $\Theta_m$  is not a bounded set then they will be unrealistic. However, for the specific case of exponential random graph models, these assumptions can be deemed realistic for the following reason. These models are known to feature degenerate regions, i.e. parts of the parameter space where the model generates only full or empty graphs. Hence, the model parameters will effectively lie on a compact space and so a Bayesian analysis of those models may be carried out, as a matter of course, with priors whose support is included in the non-degenerate region. In this sense, we feel that Assumptions (A1) to (A4) become rather mild so that the theoretical results are then relevant.

**Lemma 4.** *Under assumptions (A1)-(A4), for any  $x, x' \in \mathcal{X}$ ,  $\hat{\rho}(x, x', y')$  estimated with (10) satisfies*

$$\begin{aligned} \mathbb{E}_{y' \sim f(\cdot | x')} \left| \hat{\rho}(x, x', y') - \rho(x, x') \right| &\leq \gamma_{IS}(x, x') \\ &= \frac{1}{\sqrt{N}} \frac{q_{m'}(y | \theta')}{q_m(y | \theta)} \frac{p_{m'}(\theta')}{p_m(\theta)} \frac{p(\mathcal{M}_{m'})}{p(\mathcal{M}_m)} \frac{h(\theta, m | \theta', m')}{h(\theta', m' | \theta, m)} \\ &\quad \times \sqrt{\mathbb{V}_{y' \sim f(\cdot | x')} \left[ \frac{q_m(y' | \theta)}{q_{m'}(y' | \theta')} \right]}, \end{aligned}$$

where the upper bound  $\gamma_{IS}(x, x')$  refers to the case when the ISE is used.

Theorem 5 places a bound on the total variation between the Markov chain of a noisy RJMCMC algorithm and a Markov chain with the desired target distribution, when the intractable ratio of normalising constants is estimated by the ISE in (10). Theorem 3.1 in Alquier et al. (2016) is the special case of Theorem 5 when a within-model transition is attempted.

**Theorem 5.** *Under assumptions (A1)-(A3) then (H2) in Corollary 3 is satisfied with*

$$\gamma_{IS}(x, x') \leq \frac{c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{v(x,x')} c_{v(x',x)} K_m^2 K_{m'}^2}{\sqrt{N}}$$



$$\times \sqrt{\mathbb{V}_{y'} \left[ \prod_{j=1}^{L-1} \frac{1}{S} \sum_{b=1}^S \frac{q_m(y_b^{(j+1)'}) | \theta'(t_j))}{q_{m'}(y_b^{(j+1)'}) | \theta'(t_{j+1}))} \right]},$$

where the upper bound  $\gamma_{TP}(x, x')$  refers to the case when the TPE is used.

We conclude by an application of Corollary 3 that allows to assess the convergence of the noisy RJMCMC scheme with the TPE for  $L, S \gg 1$ .

**Theorem 8.** *Under assumptions (A1)-(A3) then (H2) in Corollary 3 is satisfied with*

$$\gamma_{TP}(x, x') \leq c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathcal{V}(x, x')} c_{\mathcal{V}(x', x)} [K_m K_{m'}]^{\frac{L+1}{2}} \left\{ \left(1 + \frac{1}{S}\right)^{L-1} - 1 \right\}^{1/2}$$

and

$$\|\delta_{x_0} \mathcal{P}^n - \delta_{x_0} \hat{\mathcal{P}}^n\| \leq \mathcal{D}_{TP} [K_m K_{m'}]^{\frac{L-1}{2}} \left\{ \left(1 + \frac{1}{S}\right)^{L-1} - 1 \right\}^{1/2},$$

where  $\mathcal{D}_{TP} = \mathcal{D}_{TP}(c_p, c_{\mathcal{M}}, c_h, K)$  is explicitly known, see (A.7) of the Appendix.

## 6. Applications to exponential random graph models

We apply the reversible jump methodology to some challenging models for the analysis of network data. All computations in this paper were carried out with the statistical environment R (R Core Team, 2017).

The two ERGM examples below aim at providing insights regarding the efficiency of the noisy RJMCMC algorithm relative to the RJ exchange algorithm. Any RJMCMC algorithm allows for Bayesian multi-model inference as well as estimation of posterior model probabilities and so we are also interested in the performance of the RJMCMC algorithm for each model separately. We assess this performance in terms of:

- The estimated within-model acceptance rate and effective sample size ( $\text{ESS}_W$ ) of the sampler for the model with the highest estimated posterior model probability. For an MCMC run of length  $G$  with lag  $\ell$  auto-correlation  $\rho_\ell$  the effective sample size is defined as  $\text{ESS}_W = G / (1 + 2 \sum_{\ell=1}^{\infty} \rho_\ell)$  (Liu, 2001). This  $\text{ESS}_W$  (the larger the better) gives an estimate of the equivalent number of independent iterations that the chain represents. In our implementation of the RJMCMC algorithms,  $\text{ESS}_W$  is estimated separately for each parameter and so we report the smallest estimate among these.
- The estimated Bayes factor, between-model acceptance rate and total computation time.
- $\text{ESS}_B$ , the estimated effective sample size of the Markov chain that targets the discrete stationary distribution with probabilities  $(\pi(\mathcal{M}_1 | y), \pi(\mathcal{M}_2 | y), \pi(\mathcal{M}_3 | y), \dots)$ . We implemented the computational approach proposed in Heck et al. (2018) to estimate  $\text{ESS}_B$  and define the efficiency (EFF) of each RJMCMC algorithm as  $\text{ESS}_B$  per second.

Throughout the analysis we chose a Multivariate Gaussian prior distribution for the model-specific parameters,  $\mathcal{MVN}(0_{d_m}, 10I_{d_m})$ , where  $0_{d_m}$  is the null vector and  $I_{d_m}$  is the identity matrix of size equal to the number of model dimensions  $d_m$ . In further simulation experiments not shown here, changing the prior distribution to  $\mathcal{MVN}(0_{d_m}, 100I_{d_m})$  did not lead to very different results. The parameter spaces  $\Theta_m, m \in \{1, 2, 3, \dots\}$  were defined on a bounded set for both examples, so that Assumptions (A1) to (A4) hold. We assumed that the models are equally probable a-priori. Sampling from the likelihood at each RJMCMC iteration was performed with an auxiliary Markov chain of length 3,000 as a proxy for an exact sampler. Each main Markov chain was run for 500,000 iterations discarding the first 50,000 as burn-in. This allowed for adequate exploration of the posterior distribution of each model.

The Auto-RJ sampler performed a pilot MCMC run for each model using the population MCMC approach of Caimo and Friel (2011), with  $2000 \times d_m$  main iterations (discarding the first  $500 \times d_m$  iterations as burn-in). The Metropolis tuning parameters involved in the RW RJMCMC algorithms were appropriately selected to give the desired within-model acceptance rates.

Finally, the TPE was obtained using a linear path (equal spacing) in  $t \in [0, 1]$  that consisted of  $L$  steps, including the starting and finishing point. At each path point the auxiliary step to draw  $y'$  was followed by an extra number of iterations thinned (by a factor of 50, yielding  $S$  graphs). In Section 6.2 the TPE was implemented using serial computation on a single core for illustration purposes. Of course, this adds to the CPU burden per iteration of the noisy RJMCMC algorithm if a more refined transition path is considered. Assuming that parallel computations occur no additional cost, the  $L - 1$  forward-simulations required in (13) can be performed in parallel, taking advantage of the inherently parallel nature of the TPE while yielding a useful noisy RJMCMC algorithm. We implemented a parallel computation of the TPE in Section 6.3 to emphasise the improvements offered by a parallel implementation of the TPE relative to the RJ exchange algorithm.

### 6.1. Exponential random graph models

Networks are relational data represented as mathematical graphs of nodes and edges. A  $n \times n$  random adjacency matrix  $Y$  on  $n$  nodes and a set of edges (relationships) describes the connectivity pattern of a graph, considering the set of all possible graphs on  $n$  nodes (actors),  $\mathcal{Y}$ . A realisation of  $Y$  is denoted with  $y$  and the presence or absence of an edge (directed or undirected) between the pair of nodes  $(i, j)$  is coded as

$$y_{ij} = \begin{cases} 1, & \text{if } (i, j) \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases}$$

An edge connecting a node to itself is not permitted so  $y_{ii} = 0$ .

Exponential random graph models play an important role in network analysis since they can represent transitivity and other structural features in network data that define complicated dependence patterns not easily modeled by more basic probability models (Wasserman and Pattison, 1996). The reader is also referred to Robins et al. (2007) for a detailed review and the references therein for more details. ERGMs belong to the exponential family of distributions with natural parameter  $\theta$  and sufficient statistics  $s(y)$  and are Gibbs random fields that are defined on the edge space of networks.

The distribution of  $Y$  assuming model  $\mathcal{M}_m$  is formulated as

$$f_m(y | \theta_m) = \frac{q_m(y | \theta_m)}{z_m(\theta_m)} = \frac{\exp\{\theta_m^\top s_m(y)\}}{\sum_{y \in \mathcal{Y}} \exp\{\theta_m^\top s_m(y)\}}, \quad \theta_m^\top s_m(y) = \sum_{j=1}^{d_m} \theta_{mj} s_{mj}(y), \quad (21)$$

where  $q_m(y | \theta_m)$  is the un-normalised likelihood,  $s_m(y) \in \mathcal{S}_m \subseteq \mathbb{R}_+^{d_m}$  is a known vector of overlapping sub-graph configurations/ sufficient statistics (e.g. the number of edges, degree statistics, triangles, etc.) and  $\theta_m \in \Theta_m \subseteq \mathbb{R}^{d_m}$  is the vector of model parameters. More examples can be found at Snijders et al. (2006) and Hunter and Handcock (2006). Additionally, ERGMs allow for inclusion of covariate information  $X$ , eg. the number of edges on the graph within the same attribute category as a measure of the homophily effect of that attribute. The respective sufficient statistics are denoted by  $s_m(y, x)$ .

The evaluation of  $z_m(\theta_m)$  is feasible for only trivially small graphs as this sum involves  $2^{\binom{n}{2}}$  terms for undirected edge graphs and  $2^{n(n-1)}$  terms for directed edge graphs. Recent works on the inference of ERGMs with the Bayesian approach have been proposed by Koskinen (2004), Caimo and Friel (2011), Wang and Atchadé (2014), Caimo and Mira (2015), Thiemichen et al. (2016) and Bouranis et al. (2017).

**Remark 1.** *The estimators (10) and (13) require perfect sampling from the likelihood to generate the auxiliary data. This is feasible for GRFs such as the Ising and Potts models (Propp and Wilson, 1996; Huber, 2004), but this is not the case for ERG models. Caimo and Friel (2011) considered the pragmatic alternative of replacing exact simulation with Gibbs updates, where a long auxiliary Markov chain implemented by the tie-no-tie (TNT) sampler is used to return  $N = 1$  draw that is approximately distributed under the true likelihood, in place of exact simulation (Hunter et al., 2008). It is possible to carry out inference for graphs of larger size (eg. 1000 nodes), but at the cost of an increased computational time. For convergence results of the Markov chain that uses approximate draws from the likelihood, see Everitt (2012) and Wang and Atchadé (2014).*

Let  $aux$  be the number of auxiliary iterations used in the TNT sampler and let  $k$  be the lag between consecutive draws from the target distribution of the sampler. The computational cost per iteration of the noisy RJMCMC algorithm for the estimation of the ratio  $z_m(\theta)/z_{m'}(\theta')$  is  $O(aux + kN)$  when the ISE is used and  $O((L-1)(aux + kS))$  when the TPE is used. For the special case of the RJ exchange algorithm, the computational cost per iteration is  $O(aux + 1)$ . For models like GRFs, these estimators depend on the simulated data only through the collection of sufficient statistics  $s_{m'}(y)$  which need to be stored in the computer memory.

## 6.2. Zachary's karate club network

Zachary's Karate Club (Figure 2) is a social network of friendships between 34 members of a karate club at a US university in 1970. For this network we are interested in the effect of triad closure, therefore it is natural to consider  $k$ -star statistics for  $k \geq 2$  and triangle counts, which express the level of transitivity.

For a fixed network density however,  $k$ -stars become more prevalent as heterogeneity increases at the cost of often inducing degeneracy. The *degeneracy problem* causes the generated graphs to be either very sparse or very dense and only in rare cases does a generated graph have a density close to that of the data network. By considering some commonly used network statistics like the heterogeneity of degree statistics and high-order transitivity statistics (Snijders et al., 2006;

Hunter, 2007), we try to remedy the degeneracy of  $k$ -stars and triangle counts. Two competing models are assessed:

$$\begin{aligned}\mathcal{M}_1: q_1(y | \theta_1) &= \exp \{ \theta_{11} s_1(y) + \theta_{12} v(y, \phi_v) \} \\ \mathcal{M}_2: q_2(y | \theta_2) &= \exp \{ \theta_{21} s_1(y) + \theta_{22} v(y, \phi_v) + \theta_{23} u(y, \phi_u) \},\end{aligned}$$

where  $s_1(y) = \sum_{i < j} y_{ij}$  is the number of edges. The other model terms are defined below.




Figure 2: Zachary's Karate Club graph.

**Shared Partnership:** Let  $EP_k(y)$ , called the edgewise shared partnership statistic, denote the number of connected pairs with exactly  $k$  common neighbors.  $EP_k(y)$  is a function of the triangle counts and as such, it is equivalent to modeling the high-order transivities. The distribution of edgewise shared partnership can be modeled as a function of a single parameter by placing decreasing weights on the higher transivities, leading to the geometrically weighted edgewise shared partnership (GWESP) statistic. GWESP is defined by:

$$v(y, \phi_v) = e^{\phi_v} \sum_{k=1}^{n-2} \left\{ 1 - \left( 1 - e^{-\phi_v} \right)^k \right\} EP_k(y).$$

**Geometrically Weighted Degree:** Let the degree count,  $D_k(y)$ , denote the number of pairs that have exactly  $k$  common neighbors. The number of stars is a function of the degrees, therefore  $D_k(y)$  is equivalent to modeling the  $k$ -star statistic. The geometrically weighted degree (GWD) statistic enables to model all degree distributions as a function of single parameter by placing decreasing weights on the higher degrees. GWD is defined by:

$$u(y, \phi_u) = e^{\phi_u} \sum_{k=1}^{n-1} \left\{ 1 - \left( 1 - e^{-\phi_u} \right)^k \right\} D_k(y).$$

The scale parameters  $(\phi_v, \phi_u)$  specify the decreasing rates of weights placed on the higher order terms, are treated as constants and are set to  $(\phi_v, \phi_u) = (0.2, 0.8)$ , following Caimo and Friel (2014). Tables 1 points towards positive evidence in favor of  $\mathcal{M}_1$  over  $\mathcal{M}_2$ . This shows that

the effect captured by the geometrically weighted degree network statistic does not enhance the observed network.

Some general comments can be made about the performance of the two algorithms. In this example all algorithms suffer from low acceptance rates of the trans-dimensional move, since the probability mass is concentrated on one model. It would be expected that the proposed jumps would be accepted more frequently if the probability mass was spread over different models. The Auto-RJ sampler for this example yields a Markov chain with a better asymptotic efficiency than the RW RJMCMC algorithm, however this algorithm will become costly when the number of models is large as it can be computationally expensive to tune the proposal distributions for each competing model. In terms of asymptotic efficiency the Auto-RJ sampler, in general, may become way worse than the noisy RW RJMCMC sampler simply because of the fact that the independent Gaussian proposal may not be a good match to the posterior distribution for some of the competing models.

Table 1: Zachary karate club - Bayes factor (standard deviation), acceptance rate of the trans-dimensional move and CPU time in hours based on thirty independent noisy RJMCMC runs. We also report the effective sample size of the Markov chain on the posterior model probabilities ( $ESS_B$ ) and the efficiency (EFF) for each algorithm.

Method		ISE ( $N$ )			TPE ( $L, S$ )	
		1	5	$10^2$	(6,1)	(11,10)
RW-RJ	$BF_{12}$	12.81 (0.37)	10.67 (0.32)	13.53 (0.33)	13.01 (0.42)	13.07 (0.45)
	% accepted	6.3	9.8	6.9	7.2	7.5
	$ESS_B$	69,479	78,575	85,389	89,612	93,090
	CPU	0.92	0.98	2.04	4.08	8.95
	EFF	18.94	22.27	10.69	5.15	2.65
Auto-RJ	$BF_{12}$	13.06 (0.21)	15.19 (0.24)	15.21 (0.26)	13.08 (0.17)	13.09 (0.15)
	% accepted	5.6	6.6	8.3	9.6	10.6
	$ESS_B$	53,760	75,251	98,852	100,101	113,297
	CPU	1.76	1.83	2.22	4.49	9.88
	EFF	8.48	18.41	12.37	6.19	1.66

Table 1 shows how all of the noisy RJMCMC algorithms displayed better mixing in terms of acceptance rate of the trans-dimensional move when compared to the RJ exchange algorithm. This improvement, though, comes with a cost: an ISE based on  $N > 1$  appears to have an effect on the frequency of visits to each model, hence biasing the Bayes factor estimate. All samplers are faced with a biased Bayes factor estimate when  $N = 5$  and this bias appears to decrease with larger  $N$ , when the noisy acceptance ratio mimics the MH acceptance ratio.

The noisy RJMCMC samplers also benefit from improved performance in terms of mixing when the TPE is used for the estimation of the intractable acceptance ratio. Table 1 suggests that even a small number of draws from the auxiliary distribution can improve the acceptance rate of the trans-dimensional move, resulting in a 40% increase on average for the noisy Auto-RJ sampler ( $L = 6, S = 1$ ) relative to the Auto-RJ exchange algorithm ( $N = 1$ ), at the cost of some additional computational expense. Increasing the overall number of auxiliary draws further improves the mixing of the noisy RJMCMC algorithms, but then their implementation becomes impractical even for this small model set. This is an indicator that the TPE will have small variance, making it particularly useful for RJMCMC.

Table 2 shows an improvement in the ESS for the most probable model over the RJ exchange algorithm ( $N = 1$ ), which translates into a decrease in the auto-correlation of the chain for each model. Both the ISE and the TPE improved the mixing within each model, which is in line with the results of Alquier et al. (2016). Taken together with the results of Table 1, we can observe a trade-off between bias and efficiency, as the noisy RJMCMC algorithm is targeting an approximation of the true posterior (4).

A natural question that arises is how big should  $N$  or  $L, S$  be in order to have a trade-off between the bias of the distribution and the computational time. From this experiment, it appears that the bias of the Bayes factor estimate is connected to the bias in the target distribution of the noisy RJMCMC chain. Knowledge of the variance of the estimator of the ratio of intractable ratio of normalising constants could give a rough idea about choosing appropriate values for the factors  $N$  or  $L, S$ .

Finally, the results in Table 1 point towards decreased efficiency of the noisy RJMCMC algorithm when a larger number of likelihood draws ( $N = 100$ ) is taken. As discussed in Alquier et al. (2016), the rationale behind the introduction of a noisy (RJ)MCMC algorithm is that the combined statistical and computational efficiency is sufficiently improved to outweigh the effect of any bias that is introduced. In the case of an ISE with  $N = 100$ , however, the increased CPU cost decreases the efficiency of the algorithm. A serial implementation of the TPE decreases the efficiency of the noisy RJMCMC algorithm relative to the RJ exchange sampler, as expected. In Section 6.3 we show that a parallel implementation of the TPE helps to improve the efficiency.

Table 2: Zachary karate club - Acceptance rate and effective sample size of the Markov chain on the posterior parameter estimates ( $ESS_W$ ) for the most probable model,  $\mathcal{M}_1$ , based on thirty independent noisy RJMCMC runs.

Method		ISE ( $N$ )			TPE ( $L, S$ )	
		1	5	$10^2$	(6,1)	(11,10)
RW-RJ	AR	22%	21%	22%	21%	22%
	$ESS_W$	12,498	16,094	20,297	20,664	24,199
Auto-RJ	AR	24%	28%	38%	40%	46%
	$ESS_W$	18,855	25,744	43,950	47,153	62,795

### 6.3. Collaboration between Lazega’s lawyers

The Lazega network dataset (Figure 3) originates from a network study of corporate law partnership that was carried out in a Northeastern US corporate law firm in New England (Lazega, 2001). The dataset consists of 36 nodes (partners) and the presence of an edge between two nodes indicates a collaboration between the two partners. Information about nodal attributes is also available; here we are interested in the attribute variables of gender (1=male; 2=female) and practice (1=litigation; 2=corporate).

We compare the following three models as in Caimo and Friel (2014):

$$\begin{aligned} \mathcal{M}_1: q_1(y | \theta_1) &= \exp \{ \theta_{11}s_1(y) + \theta_{12}v(y, \phi_v) + \theta_{13}s_1(y, x) + \theta_{14}s_2(y, x) + \theta_{15}s_3(y, x) \} \\ \mathcal{M}_2: q_2(y | \theta_2) &= \exp \{ \theta_{21}s_1(y) + \theta_{22}v(y, \phi_v) + \theta_{23}s_1(y, x) + \theta_{24}s_2(y, x) \} \\ \mathcal{M}_3: q_3(y | \theta_3) &= \exp \{ \theta_{31}s_1(y) + \theta_{32}v(y, \phi_v) \} \end{aligned}$$

where  $s_1(y)$  and  $v(y, \phi_v)$  are the same terms as in Section 6.2 and  $\phi_v = \log(2)$ . The covariate

statistics are defined as

$$s(y, x) = \sum_{i \neq j} y_{ij} \times \begin{cases} x_i + x_j & : \text{"main effect"}, \\ \mathbb{1}_{\{x_i=x_j\}} & : \text{"homophily effect"}. \end{cases}$$

The covariate statistic  $s_1(y, x)$  represents the homophily effect of practice,  $s_2(y, x)$  is the homophily effect of gender and  $s_3(y, x)$  is the main effect of practice.




Figure 3: Lazega's network of corporate law partnership.

All algorithms point towards  $\mathcal{M}_3$  having the highest estimated posterior model probability. For illustration purposes, below we report the estimated Bayes factor for the comparison between  $\mathcal{M}_3$  against  $\mathcal{M}_2$ . The results in Table 3 show that there is positive evidence in favor of  $\mathcal{M}_3$  against  $\mathcal{M}_2$ .

Table 3: Lazega's lawyers - Bayes factor, acceptance rate of the trans-dimensional move, effective sample size of the Markov chain on the posterior model probabilities ( $ESS_{pm}$ ) and efficiency (EFF) for each algorithm, based on a long RJMCMC run.

Method		ISE ( $N$ )			TPE ( $L, S$ )	
		1	5	$10^2$	(6,1)	(11,10)
RW-RJ	$BF_{32}$	13.12	23.12	13.01	13.35	13.76
	% accepted	1.0	1.2	1.3	2.0	3.0
	$ESS_{pm}$	1,007	4,131	5,382	5,104	6,422
	EFF	0.23	0.87	1.25	1.16	1.46
Auto-RJ	$BF_{32}$	13.46	46.31	13.72	14.02	13.89
	% accepted	1.0	1.0	2.0	3.0	3.0
	$ESS_{pm}$	11,211	59,200	58,863	41,432	54,146
	EFF	1.26	6.55	4.65	4.66	6.05

The higher-dimensionality of the competing models allows us to investigate the performance of noisy RJMCMC in a situation that is more likely to be faced in practice. Table 3 shows that a small number of draws on the ISE ( $N = 5$ ) leads to a bias in the Bayes factor estimate that is pronounced in this example for both the random walk and the independence samplers. At the

same time, the mixing of the RJ Markov chain improves, but many more draws ( $N > 100$ ) will be needed to obtain a better approximation to the target posterior distribution and consequently reduce the bias in the estimated posterior model probabilities and the Bayes factor estimate.

On the contrary, the TPE proves itself useful even with a small number in the factors  $L, S$ , eg. ( $L = 6, S = 1$ ). The increased efficiency of the noisy RJMCMC algorithms with the parallel implementation of the TPE depicted in Table 3 shows more emphatically that the new approaches offer large improvements over running the RJ exchange algorithm, while simultaneously improving the statistical ( $ESS_{pm}$ ) and computational (EFF) efficiency of the sampler.

The same comments made for Table 2 can be also made for the results in Table 4: the noisy RJMCMC algorithms had less auto-correlation within each model compared to the RJ exchange algorithm, similarly to Zachary’s Karate Club network example.

In both ERGM examples, we observed that the noisy independence sampler does better than the noisy RW RJMCMC algorithm. This is a result of having a posterior distribution for each model that is close to Gaussian and so the automated design of the independence proposal is efficient. However, this may not generally be the case and in this instance we expect that the noisy RW RJMCMC algorithm may give better results.

Table 4: Lazega’s lawyers - Acceptance rate and effective sample size of the Markov chain on the posterior parameter estimates ( $ESS_W$ ) for the most probable model,  $\mathcal{M}_3$ , based on a long RJMCMC run.

Method		ISE ( $N$ )			TPE ( $L, S$ )	
		1	5	$10^2$	(6,1)	(11,10)
RW-RJ	AR	22%	19%	25%	24%	25%
	$ESS_W$	4,441	4,915	5,881	9,117	10,063
Auto-RJ	AR	17%	18%	19%	25%	28%
	$ESS_W$	8,948	10,763	15,318	20,526	25,032

## 7. Discussion

The present paper contributes to the growing literature of approximate MCMC methods for Bayesian analysis of doubly-intractable distributions (Alquier et al., 2016; Everitt et al., 2017a; Boland et al., 2017; Everitt et al., 2017b) by introducing a variant of RJMCMC for Bayesian model comparison of Gibbs random fields. The resulting algorithm generalises noisy MCMC to trans-dimensional settings, where the transition kernel of the exact RJMCMC algorithm is approximated. Drawing from the study of the stability of Markov chains (Mitrophanov, 2005; Alquier et al., 2016), we have given bounds on the total variation between the Markov chain of a noisy RJMCMC algorithm and a Markov chain with the desired target distribution in the case where the chain is uniformly ergodic. We acknowledge, though, that this is a strong assumption that may not be met in practice.

We have illustrated that the noisy RJMCMC algorithm can suffer from considerable bias (reflected by the Bayes factor estimate) when the variance of the estimator of the ratio of intractable likelihood normalising constants is large. In particular, the unbiased estimator proposed in Alquier et al. (2016) is not useful for RJMCMC. Aiming to overcome these inefficiencies and to decrease the variance of this estimator, we used a variation based on a telescopic product of unbiased importance sampling estimators. The telescopic product estimator (TPE) helped reduce the bias of the noisy RJMCMC algorithm, which is then of the same order of magnitude as that

of the RJ exchange algorithm, while being more asymptotically efficient than the RJ exchange algorithm.

Despite the fact that this framework also applies to more general situations apart from Gibbs random field models, there are limitations to the implementation of the noisy RJMCMC algorithm. A crucial aspect of the computational performance of the algorithm is its requirement to perform forward simulations from the likelihood at each iteration. The computational cost is expected to increase when models on large lattices or networks with thousands of nodes are analysed. The estimators that we investigated in this paper rely on importance sampling and might not be practical in high-dimensional parameter spaces, a point which has been raised by Everitt (2012) and Everitt et al. (2017b). In real-life GRF applications, the parameter spaces do not usually exceed 10-15 dimensions.

## Acknowledgements

The authors thank the co-editor, the associate editor and the anonymous referees for their constructive comments that helped improve the article. The Insight Centre for Data Analytics is supported by Science Foundation Ireland under Grant Number SFI/12/RC/2289. Nial Friel's research was also supported by a Science Foundation Ireland grant: 12/IP/1424.

## References

- Alquier, P., Friel, N., Everitt, R., and Boland, A. (2016). Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, 26(1):29–47.
- Andrieu, C. and Roberts, G. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725.
- Boland, A., Friel, N., and Maire, F. (2017). Efficient MCMC for Gibbs random fields using pre-computation. *arXiv:1710.04093*.
- Bouranis, L., Friel, N., and Maire, F. (2017). Efficient Bayesian inference for exponential random graph models by correcting the pseudo-posterior distribution. *Social Networks*, 50:98–108.
- Bouranis, L., Friel, N., and Maire, F. (2018). Bayesian model selection for exponential random graph models via adjusted pseudolikelihoods. *Journal of Computational and Graphical Statistics*. <https://doi.org/10.1080/10618600.2018.1448832>.
- Brooks, S. P., Giudici, P., and Roberts, G. O. (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions (with discussion). *Journal of the Royal Statistical Society, Series B*, 65(1):3–55.
- Caimo, A. and Friel, N. (2011). Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55.
- Caimo, A. and Friel, N. (2013). Bayesian model selection for exponential random graph models. *Social Networks*, 35(1):11–24.
- Caimo, A. and Friel, N. (2014). Bergm: Bayesian exponential random graphs in R. *Journal of Statistical Software*, 61(2):1–25.

- Caimo, A. and Mira, A. (2015). Efficient computational strategies for doubly intractable problems with applications to Bayesian social networks. *Statistics and Computing*, 25(1):113–125.
- Chib, S. and Jeliazkov, I. (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association*, 96:270–281.
- Ehlers, R. S. and Brooks, S. P. (2008). Adaptive proposal construction for reversible jump MCMC. *Scandinavian Journal of Statistics*, 35:677–690.
- Everitt, R. (2012). Bayesian parameter estimation for latent Markov random fields and social networks. *Journal of Computational and Graphical Statistics*, 24(4):940–960.
- Everitt, R., Johansen, A., Rowing, E., and Hogan, M. (2017a). Bayesian model selection with un-normalised likelihoods. *Statistics and Computing*, 27(2):403–422.
- Everitt, R., Prangle, D., Maybank, P., and Bell, M. (2017b). Marginal sequential Monte Carlo for doubly intractable models. *arXiv:1710.04382*.
- François, O., Ancelet, S., and Guillot, G. (2006). Bayesian clustering using hidden Markov random fields in spatial population genetics. *Genetics*, 174(2):805–816.
- Friel, N. (2013). Evidence and Bayes factor estimation for Gibbs random fields. *Journal of Computational and Graphical Statistics*, 22(3):518–532.
- Friel, N. and Rue, H. (2007). Recursive computing and simulation-free inference for general factorizable models. *Biometrika*, 93(3):661–672.
- Friel, N. and Wyse, J. (2012). Estimating the evidence - a review. *Statistica Neerlandica*, 66(3):288–308.
- Gelman, A. and Meng, X. L. (1998). Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185.
- Geyer, C. and Thompson, E. (1992). Constrained Monte Carlo maximum likelihood for dependent data (with discussion). *Journal of the Royal Statistical Society, Series B*, 54:657–99.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82:711–732.
- Green, P. J. (2003). Trans-dimensional Markov chain Monte Carlo. In Green, P. J., Hjort, N. L., and Richardson, S., editors, *Highly Structured Stochastic Systems, Oxford Statistical Science Series, No. 27*, page 179–198. Oxford University Press, Oxford.
- Grelaud, A., Robert, C., Marin, J., Rodolphe, F., and Taly, J. (2009). ABC likelihood-free methods for model choice in Gibbs random fields. *Bayesian Analysis*, 4(2):317–335.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109.

- Heck, D., Overstall, A., Gronau, Q., and Wagenmakers, E. (2018). Quantifying uncertainty in trans-dimensional Markov chain Monte Carlo Using discrete Markov models. *arXiv:1703.10364v2*.
- Huber, M. (2004). Perfect sampling using bounding chains. *The Annals of Applied Probability*, 14(2):734–753.
- Hunter, D. (2007). Curved exponential family models for social networks. *Social Networks*, 29(2):216–230.
- Hunter, D. and Handcock, M. (2006). Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, 15(3):565–583.
- Hunter, D., Handcock, M., Butts, C., Goodreau, S., and Morris, M. (2008). ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3):1–29.
- Jacob, P. and Thiery, A. (2015). On nonnegative unbiased estimators. *The Annals of Statistics*, 43(2):769–784.
- Karagiannis, G. and Andrieu, C. (2013). Annealed Importance Sampling Reversible Jump MCMC Algorithms. *Journal of Computational and Graphical Statistics*, 22(3):623–648.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90:773–795.
- Koskinen, J. (2004). Bayesian analysis of exponential random graphs - Estimation of parameters and model selection. Technical Report 2, Department of Statistics, Stockholm University.
- Lazega, E. (2001). *The collegial phenomenon : the social mechanisms of cooperation among peers in a corporate law partnership*. Oxford University Press.
- Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York.
- Lyne, A., Girolami, M., Atchadé, Y., Strathmann, H., and Simpson, D. (2015). On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statistical science*, 30(4):443–467.
- Martin, A., Quinn, K., and Park, J. (2011). MCMCpack: Markov chain Monte Carlo in R. *Journal of Statistical Software*, 42(9):1–21.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.
- Mitrophanov, A. (2005). Sensitivity and convergence of uniformly ergodic Markov chains. *Journal of Applied Probability*, 42:1003–1014.
- Murray, I., Ghahramani, Z., and MacKay, C. (2006). MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, pages 359–366.

- Propp, J. G. and Wilson, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, 9(1-2):223–252.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Robins, G., Snijders, T., Wang, P., Handcock, M., and Pattison, P. (2007). Recent developments in exponential random graph ( $p^*$ ) models for social networks. *Social Networks*, 29:192–215.
- Snijders, T., Pattison, P., Robins, G., and Handcock, M. (2006). New specifications for exponential random graph models. *Sociological Methodology*, 36(1):99–153.
- Stoehr, J., Benson, A., and Friel, N. (2017). Noisy Hamiltonian Monte Carlo for doubly-intractable distributions. *arXiv:1706.10096*.
- Thiemichen, S., Friel, N., Caimo, A., and Kauermann, G. (2016). Bayesian exponential random graph models with nodal random effects. *Social Networks*, 46:11–28.
- Wang, J. and Atchadé, Y. F. (2014). Approximate Bayesian computation for exponential random graph models for large social networks. *Communications in Statistics - Simulation and Computation*, 43(2):359–377.
- Wasserman, S. and Pattison, P. (1996). Logit models and logistic regression for social networks: I. An introduction to Markov graphs and  $p^*$ . *Psychometrika*, 61:401–425.

## Appendix A. Proofs

*Proof of Proposition 1.* By definition of  $\hat{\theta}$ , it holds that  $z_m(\theta) = z_{m'}(\hat{\theta})$ . Then we can write

$$\begin{aligned} \mathbb{V}_{y'_1, y'_2, \dots, y'_N \sim f_{m'}(\cdot | \theta')} \left[ \frac{\widehat{z_{m'}(\hat{\theta})}}{z_{m'}(\theta')} \right] &= \mathbb{E}_{\theta'} \left[ \exp\{2(\hat{\theta} - \theta')^\top s_{m'}(y)\} \right] - \left[ \frac{z_{m'}(\hat{\theta})}{z_{m'}(\theta')} \right]^2 \\ &= \mathbb{E}_{\theta'} \left[ \exp\{2\zeta^\top s_{m'}(y)\} \right] - \left[ \frac{z_{m'}(\hat{\theta})}{z_{m'}(\theta')} \right]^2, \end{aligned}$$

where  $\zeta = \hat{\theta} - \theta'$ . We introduce the notation  $\psi = O(\|\zeta\|_2)$ , where  $\|\cdot\|_2$  is the  $L^2$  norm. A Taylor expansion of the variance term around  $\theta'$  yields

$$\begin{aligned} \mathbb{V}_{y'_1, y'_2, \dots, y'_N \sim f_{m'}(\cdot | \theta')} \left[ \frac{\widehat{z_{m'}(\hat{\theta})}}{z_{m'}(\theta')} \right] &= 1 + 2\zeta^\top \mathbb{E}_{\theta'} [s_{m'}(y)] + \mathbb{E}_{\theta'} [\psi(\zeta, y)] \\ &\quad - \frac{1}{z_{m'}^2(\theta')} \left[ z_{m'}^2(\theta') + 2z_{m'}(\theta')\zeta^\top \nabla_{\theta'} z_{m'}(\theta') + O(\|\zeta\|_2^2) \right]. \end{aligned}$$

We note that for models that belong to the exponential family it holds that  $\nabla_{\theta} z(\theta)/z(\theta) = \mathbb{E}_{\theta} [s(y)]$   $\forall \theta \in \Theta$ . Additionally,  $\mathbb{E}_{\theta'} [\psi(\zeta, y)] = O(\|\zeta\|_2)$ , which gives

$$\mathbb{V}_{y'_1, y'_2, \dots, y'_N \sim f_{m'}(\cdot | \theta')} \left[ \frac{\widehat{z_{m'}(\hat{\theta})}}{z_{m'}(\theta')} \right] = O(\|\zeta\|_2).$$

This concludes the proof.  $\square$

*Proof of Corollary 3.* Let  $x = (m, \theta)$  and  $x' = (m', \theta')$ . We assume that there exists a symmetric measure  $\mu$  on  $\mathcal{X} \times \mathcal{X}$  that dominates  $\pi(dx)\nu(x, dx')$ . Then  $\pi(dx)\nu(x, dx')$  has density  $\nu(x, x')$  (the Radon-Nikodym derivative) with respect to  $\mu$ . For all Borel sets  $\mathcal{B}, \mathcal{B}' \subset \mathcal{X}$  such that  $\mathcal{B}' = dx' = \{X \in \mathcal{X}; \|X - x'\| < \varepsilon\}$ , the following holds for  $\mathcal{P}(x, dx')$  and  $\hat{\mathcal{P}}(x, dx')$  in order to apply Theorem 2:

$$\mathcal{P}(x, dx') = \nu(x, dx') \min(1, \rho(x, x')) + \delta_x(dx') \int dt \nu(x, t) [1 - \min(1, \rho(x, t))]$$

and

$$\begin{aligned} \hat{\mathcal{P}}(x, dx') &= \int dy' f(y' | x') \left[ \nu(x, x') \min(1, \hat{\rho}(x, x', y')) \right] \\ &\quad + \delta_x(dx') \iint dt dy' \nu(x, t) f(y' | t) [1 - \min(1, \hat{\rho}(x, t, y'))]. \end{aligned}$$

We write

$$\begin{aligned} (\mathcal{P} - \hat{\mathcal{P}})(x, dx') &= \int dy' f(y' | x') \nu(x, dx') \left[ \min(1, \rho(x, x')) - \min(1, \hat{\rho}(x, x', y')) \right] \\ &\quad + \delta_x(dx') \iint dt dy' \nu(x, t) f(y' | t) \left[ \min(1, \hat{\rho}(x, t, y')) - \min(1, \rho(x, t)) \right], \end{aligned}$$

which gives

$$\begin{aligned} \|\mathcal{P} - \hat{\mathcal{P}}\| &= \sup_{\substack{x \in \mathcal{X} \\ \mu = \delta_x}} \|\mu(\mathcal{P} - \hat{\mathcal{P}})\| \\ &= \sup_{x \in \mathcal{X}} \|\delta_x \mathcal{P} - \delta_x \hat{\mathcal{P}}\| \\ &= \frac{1}{2} \sup_{x \in \mathcal{X}} \int |\mathcal{P} - \hat{\mathcal{P}}|(x, dx') \\ &= \frac{1}{2} \sup_{x \in \mathcal{X}} \left\{ \left| \iint dt dy' f(y' | t) \nu(x, t) \left[ \min(1, \hat{\rho}(x, t, y')) - \min(1, \rho(x, t)) \right] \right| \right. \\ &\quad \left. + \left| \iint dy' dx' f(y' | x') \left[ \nu(x, x') \min(1, \rho(x, x')) - \nu(x, x') \min(1, \hat{\rho}(x, x', y')) \right] \right| \right\} \\ &= \sup_{x \in \mathcal{X}} \left\{ \left| \iint dt dy' f(y' | t) \nu(x, t) \left[ \min(1, \hat{\rho}(x, t, y')) - \min(1, \rho(x, t)) \right] \right| \right\} \\ &\leq \sup_{x \in \mathcal{X}} \iint dy' dx' f(y' | x') \nu(x, x') \left| \min(1, \rho(x, x')) - \min(1, \hat{\rho}(x, x', y')) \right| \\ &= \sup_{x \in \mathcal{X}} \int dx' \nu(x, x') \int dy' f(y' | x') \left| \min(1, \rho(x, x')) - \min(1, \hat{\rho}(x, x', y')) \right| \\ &\leq \sup_{x \in \mathcal{X}} \int dx' \nu(x, x') \gamma(x, x'). \end{aligned}$$

$\square$

*Proof of Lemma 4.* Let  $x = (m, \theta)$  and  $x' = (m', \theta')$ . We check that

$$\begin{aligned}
\mathbb{E}_{y' \sim f(\cdot | x')} |\hat{\rho}(x, x', y') - \rho(x, x')| &\leq \int f(y' | x') |\hat{\rho}(x, x', y') - \rho(x, x')| dy' \\
&= \frac{q_{m'}(y | \theta')}{q_m(y | \theta)} \frac{p_{m'}(\theta')}{p_m(\theta)} \frac{p(\mathcal{M}_{m'})}{p(\mathcal{M}_m)} \frac{h(\theta, m | \theta', m')}{h(\theta', m' | \theta, m)} \\
&\quad \times \mathbb{E}_{y'_1, \dots, y'_N \sim f(\cdot | x')} \left| \frac{1}{N} \sum_{i=1}^N \frac{q_m(y'_i | \theta)}{q_{m'}(y'_i | \theta')} - \frac{z_m(\theta)}{z_{m'}(\theta')} \right| \\
&\leq \frac{1}{\sqrt{N}} \frac{q_{m'}(y | \theta')}{q_m(y | \theta)} \frac{p_{m'}(\theta')}{p_m(\theta)} \frac{p(\mathcal{M}_{m'})}{p(\mathcal{M}_m)} \frac{h(\theta, m | \theta', m')}{h(\theta', m' | \theta, m)} \\
&\quad \times \sqrt{\mathbb{V}_{y'_1 \sim f(\cdot | x')} \left[ \frac{q_m(y'_1 | \theta)}{q_{m'}(y'_1 | \theta')} \right]} \\
&= \gamma_{IS}(x, x').
\end{aligned}$$

Then

$$\begin{aligned}
\gamma_{IS}(x, x') &\leq \frac{1}{\sqrt{N}} c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\nu(x, x')} c_{\nu(x', x)} \frac{q_{m'}(y | \theta')}{q_m(y | \theta)} \sqrt{\mathbb{E}_{y' \sim f(\cdot | x')} \left[ \left( \frac{q_m(y' | \theta)}{q_{m'}(y' | \theta')} \right)^2 \right]} \\
&\leq \frac{1}{\sqrt{N}} c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\nu(x, x')} c_{\nu(x', x)} K_m^2 K_{m'}^2, \tag{A.1}
\end{aligned}$$

which concludes the proof.  $\square$

*Proof of Theorem 5 (see also Theorem 3.1 of Alquier et al. (2016)).* Under the assumptions of Theorem 5, note that (7) leads to

$$\begin{aligned}
\rho(x, x') &= \frac{f_{m'}(y | \theta')}{f_m(y | \theta)} \frac{p_{m'}(\theta')}{p_m(\theta)} \frac{p(\mathcal{M}_{m'})}{p(\mathcal{M}_m)} \frac{\omega(m', m)}{\omega(m, m')} \frac{T_{m', m}(\theta', \theta)}{T_{m, m'}(\theta, \theta')} \\
&\geq [c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\nu(x, x')} c_{\nu(x', x)} K_m^2 K_{m'}^2]^{-1}. \tag{A.2}
\end{aligned}$$

Let us consider any measurable subset  $\mathcal{B}'$  of  $\mathcal{X}$  and  $x' \in \mathcal{X}$ . We have

$$\begin{aligned}
\mathcal{P}(x, \mathcal{B}') &= \int_{\mathcal{B}'} \delta_x(dx') \int dt \nu(x, t) [1 - \min(1, \rho(x, t))] + \int_{\mathcal{B}'} dx' \nu(x, x') \min(1, \rho(x, x')) \\
&\geq \int_{\mathcal{B}'} dx' \nu(x, x') \min(1, \rho(x, x')) \\
&\geq [c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\nu(x, x')} c_{\nu(x', x)} K_m^2 K_{m'}^2]^{-1} \int_{\mathcal{B}'} dx' \nu(x, x') \text{ thanks to (A.2)} \\
&\geq [c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\nu(x, x')}^2 c_{\nu(x', x)} K_m^2 K_{m'}^2]^{-1} \int_{\mathcal{B}'} dx'.
\end{aligned}$$

This proves that  $\mathcal{B}'$  is a small set for the Lebesgue measure (multiplied by a constant) on  $\mathcal{X}$ . Following the proof of Theorem 3.1 of Alquier et al. (2016) and the reference therein, this proves

that:

$$\sup_{x_0 \in \mathcal{X}} \|\delta_{x_0} \mathcal{P}^n - \pi(\cdot | y)\| \leq Q\xi^n,$$

where

$$Q = 2 \text{ and } \xi = 1 - [c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathbf{v}(x,x')}^2 c_{\mathbf{v}(x',x)} K_m^2 K_{m'}^2]^{-1}.$$

By definition,  $\mathcal{K}, c_p, c_v > 1$  and so  $\xi \in (0, 1)$ . This satisfies condition **(H1)** in Corollary 3. Moreover, (A.1) satisfies condition **(H2)** in Corollary 3. We can apply this Corollary to give

$$\sup_{x_0 \in \mathcal{X}} \|\delta_{x_0} \mathcal{P}^n - \delta_{x_0} \hat{\mathcal{P}}^n\| \leq \frac{\mathcal{D}_{IS}}{\sqrt{N}}$$

with

$$\mathcal{D}_{IS} = c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathbf{v}(x,x')} c_{\mathbf{v}(x',x)} K_m^2 K_{m'}^2 \left( \phi + \frac{Q\xi^\phi}{1-\xi} \right), \quad (\text{A.3})$$

where  $\phi = \left\lceil \frac{\log(1/Q)}{\log(\xi)} \right\rceil$ . □

*Proof of Lemma 6.* The proof is for trans-dimensional moves and follows from Boland et al. (2017). We will use the fact that for any one-dimensional random variable  $X$ ,

$$\exists M \in \mathbb{R} \quad \text{s.t.} \quad X \leq M \quad \Rightarrow \quad \mathbb{V}[X] \leq \mathbb{E}[X^2] \leq M^2. \quad (\text{A.4})$$

By definition of (13),  $\mathbb{V} \left[ \prod_{j=1}^{L-1} \bar{X}_S^j \right]$  is a collection of  $2^r - 1$  products of  $r = L - 1$  positive factors. Each factor is either a squared expectation,  $\mathbb{E}^2[\bar{X}_S^j]$ , or a variance,  $\mathbb{V}[\bar{X}_S^j]$ , so that one of the  $2^r - 1$  products that contains  $k > 0$  variances and  $r - k$  squared expectations is

$$p_k := \prod_{t=1}^k \mathbb{V}[\bar{X}_S^t] \times \prod_{q=k+1}^r \mathbb{E}^2[\bar{X}_S^q],$$

which can be re-expressed as

$$p_k := \frac{1}{S^k} \prod_{t=1}^k \mathbb{V}[X^t] \times \prod_{q=k+1}^r \mathbb{E}^2[X^q].$$

A uniform bound in  $k$  can be placed on  $p_k$ , such that

$$p_k \leq \frac{1}{S^k} \prod_{t=1}^k \mathbb{E}[(X^t)^2] \times \prod_{q=k+1}^r \mathbb{E}^2[X^q] \leq \frac{(K_m K_{m'})^r}{S^k}. \quad (\text{A.5})$$

There are  $\binom{r}{k}$  terms that have  $k$  variances and  $r - k$  squared expectations. Therefore, their sum  $\bar{p}_k$  can be bounded by the uniform bound in (A.5) so that

$$\bar{p}_k \leq \frac{(K_m K_{m'})^r}{S^k}.$$

Rearrangement of the  $2^r - 1$  products and aggregation of the products with the same number of



$$\begin{aligned}
&\leq c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathfrak{V}(x,x')} c_{\mathfrak{V}(x',x)} K_m K_{m'} [K_m K_{m'}]^{\frac{L-1}{2}} \left\{ \left(1 + \frac{1}{S}\right)^{L-1} - 1 \right\}^{1/2} \\
&= c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathfrak{V}(x,x')} c_{\mathfrak{V}(x',x)} [K_m K_{m'}]^{\frac{L+1}{2}} \left\{ \left(1 + \frac{1}{S}\right)^{L-1} - 1 \right\}^{1/2}. \tag{A.6}
\end{aligned}$$

□

*Proof of Theorem 8.* Following the proof of Theorem 5, it is straightforward to show that condition **(H1)** in Corollary 3 is satisfied, with

$$Q = 2 \text{ and } \xi = 1 - \left[ c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathfrak{V}(x,x')}^2 c_{\mathfrak{V}(x',x)} (K_m K_{m'})^{\frac{L+1}{2}} \right]^{-1}.$$

Moreover, (A.6) satisfies condition **(H2)** in Corollary 3. Applying this Corollary gives

$$\sup_{x_0 \in \mathcal{X}} \|\delta_{x_0} \mathcal{P}^n - \delta_{x_0} \hat{\mathcal{P}}^n\| \leq \mathcal{D}_{TP} \left\{ \left(1 + \frac{1}{S}\right)^{L-1} - 1 \right\}^{1/2},$$

with

$$\mathcal{D}_{TP} = c_{p_m} c_{p_{m'}} c_{\mathcal{M}_m} c_{\mathcal{M}_{m'}} c_{\mathfrak{V}(x,x')} c_{\mathfrak{V}(x',x)} [K_m K_{m'}]^{\frac{L+1}{2}} \left( \phi + \frac{Q \xi^\phi}{1 - \xi} \right) \tag{A.7}$$

and  $\phi = \left\lceil \frac{\log(1/Q)}{\log(\xi)} \right\rceil$ .

□