

# Trans-dimensional Markov chains: A decade of progress and future perspectives

Scott A. Sisson  
Department of Statistics, School of Mathematics,  
University of New South Wales.

October 28, 2004

Keywords: Automatic sampling; Convergence assessment; Efficient chains; Reference priors; Reversible jump; Software.

## Abstract

The last ten years have witnessed the development of sampling frameworks that permit the construction of Markov chains which simultaneously traverse both parameter and model space. In this time substantial methodological progress has been made. In this article we present a survey of the current state of the art and evaluate some of the most recent advances in this field. We also discuss future research perspectives in the context of the drive to develop sampling mechanisms with high degrees of both efficiency and automation.

## 1 Introduction

Simultaneous inference on both model and parameter space is an issue that is fundamental to modern statistical practice. In general, for observed data  $\mathbf{x}$  we might consider a countable set of models  $\mathcal{M} = \{M_1, M_2, \dots\}$  indexed by a parameter  $k \in \mathcal{K}$ , each with a parameter vector defined on  $\theta_k \in \Theta_k$  of length  $n_k$ . Under a Bayesian framework we would relate each model to a posterior distribution

$$M_k : \quad \tilde{\pi}_k(\theta_k | \mathbf{x}) = \pi_k(\theta_k | \mathbf{x}) / m_k(\mathbf{x}) \\ \propto L_k(\mathbf{x} | \theta_k) p_k(\theta_k),$$

generally only known up to a constant of proportionality,  $m_k(\mathbf{x})^{-1}$ , where  $L_k$  and  $p_k$  denote the likelihood and parameter prior under model  $M_k$ . Explicitly expressing  $m_k(\mathbf{x}) = \int_{\Theta_k} L_k(\mathbf{x} | \theta_k) p_k(\theta_k) d\theta_k$  as the marginal or predictive densities of  $\mathbf{x}$  under model  $M_k$ , the normalised posterior probability of model  $M_k$  is given by

$$M_k(\mathbf{x}) = \frac{\rho_k m_k(\mathbf{x})}{\sum_{i=1}^{|\mathcal{K}|} \rho_i m_i(\mathbf{x})} = \left( 1 + \sum_{i \neq k} \frac{\rho_i}{\rho_k} B_{ik} \right)^{-1}, \quad (1)$$

where  $B_{ik} = m_i(\mathbf{x})/m_k(\mathbf{x})$  is the Bayes Factor of model  $M_i$  to  $M_k$ , and  $\rho_k$  is the prior probability of model  $k$ . See, for example, Chipman et al. (2001), Berger and Pericchi (2001), Kass and Raftery (1995), Ghosh and Samanta (2001), Berger and Pericchi (2004), Barbieri and Berger (2004), Robert (2001), George and McCulloch (1996), Madigan and Raftery (1994) for discussion of Bayesian model selection techniques. As an alternative to the selection of a single model, a common approach within the Bayesian framework is that of model averaging which incorporates model uncertainty in addition to parameter uncertainty. Here interest would be in some predictive density

$$\pi(y|\mathbf{x}) = \int_{\mathcal{K}} \int_{\Theta_k} \pi(y|\theta_k) \tilde{\pi}_k(\theta_k|\mathbf{x}) d\theta_k dk$$

with the integration over both model and parameter space. There is a fantastic literature on the application of Bayesian methods for model uncertainty. General review articles such as Clyde and George (2004), Chipman et al. (2001), Clyde (1999a), Hoeting et al. (1999) contain a wealth of information and references. Similarly Andrieu et al. (2004) review the recent computational and technological advances relating to Bayesian analyses, and Müller and Quintana (2004), Heikkinen (2003) summarise the current state of nonparametric Bayesian inference.

A typical Bayesian analysis based on the above will encounter two related problems. Firstly, the density  $m_k(\mathbf{x})$  will in general be unavailable due to analytic intractability. Secondly the number of candidate models,  $|\mathcal{K}|$ , will often be very large, prohibiting a brute-force calculation of  $M_k(\mathbf{x})$  via (1). One of the more flexible and popular techniques employed to overcome each of these problems are Markov chain Monte Carlo (MCMC) methods. For the approximation of  $m_k(\mathbf{x})$  in particular, these include various forms of Metropolis-Hastings samplers – see, for example, Robert and Casella (2004), Cappé and Robert (2000), Gilks et al. (1996) for a discussion of standard MCMC methods and implementation issues. However, techniques capable of simultaneously considering a large number of candidate models did not become available until the mid-1990’s.

Almost exactly a decade ago Green (1995) recast the terms and definitions involved in the Metropolis-Hastings algorithm in a more rigorous manner — in particular, the idea of using a time reversibility condition for the transition kernel of a Markov chain in order to ensure convergence to the desired stationary distribution was extended to more general state spaces. In integral form, the detailed balance condition for a general transition kernel  $P$  and its invariant distribution  $\pi$  can be written as

$$\int_{(\mathbf{x}, \mathbf{x}') \in \mathcal{A} \times \mathcal{B}} \pi(d\mathbf{x}) P(\mathbf{x}, d\mathbf{x}') = \int_{(\mathbf{x}, \mathbf{x}') \in \mathcal{A} \times \mathcal{B}} \pi(d\mathbf{x}') P(\mathbf{x}', d\mathbf{x}) \quad (2)$$

for all Borel sets  $\mathcal{A} \times \mathcal{B} \subset \Theta$  for a general state space  $\Theta$  – see, for example, Green (2001); Tierney (1998). The novelty of this re-expression was that the generality of the state space under consideration now included formulations that could encompass multiple models. One often considered instance is  $\Theta = \bigcup_{k \in \mathcal{K}} \Theta_k \times \{k\}$ ; that is, a countable union of sub-spaces of possibly varying dimensionality. Via standard Metropolis-Hastings updates, this development enabled the implementation of Markov chains simultaneously spanning both parameter and model space,  $\Theta$ , with stationary distribution  $\pi$  which is absolutely continuous in  $\Theta_k$  for each  $k \in \mathcal{K}$ , with respect to the  $n_k$ -dimensional Lebesgue measure. As a result, the estimation of posterior model probabilities and other marginal densities of interest is in theory

easily obtainable, irrespective of the order of  $\mathcal{K}$ . The general class of Markov chains that admit transitions between states of differing dimension have since been developed further to achieve a broad family of inter-related sampling frameworks. Given their model spanning nature, they have recently been termed *trans-dimensional Markov chains*.

This article aims to achieve three objectives. Firstly, to provide an accessible evaluation of the current state of the art in terms of the practical implementation of trans-dimensional sampling technologies; secondly to overview the most recent and most important developments in multi-model sampling frameworks, in terms of ongoing theoretical and methodological development, and finally, to suggest some perspective of residual open problems, and highlighting those facets which would strongly benefit from further research.

While trans-dimensional sampling algorithms are now well documented in the literature, in the first part of this article, Section 2, we present both a brief overview of the main frameworks that have been developed in the last ten years, and an assessment of their impact and implementation including an evaluation of the freely available software for this purpose. We also consider non-Bayesian applications, the efficient estimation of Bayes Factors and the highly important, although frequently overlooked issue of convergence assessment.

In the second part of this article, Section 3, we examine one facet of trans-dimensional sampling schemes that holds considerable potential for future research; the development of algorithms with increasing degrees of efficiency and automation. Achieving this — one of the fundamental goals of modern sampling frameworks — would permit routine implementation of trans-dimensional samplers by non-expert practitioners, perhaps via stand-alone software packages such as the popular *WinBUGS* suite (Gilks et al. 1992; Spiegelhalter et al. 1996b). Recent developments have made great strides in this direction, providing advances in the areas of between model transitions, both in terms of efficiency and constructing generic mappings, the extension of perfect sampling schemes to the trans-dimensional case and progress in default prior specifications over joint model and parameter spaces.

## 2 Trans-dimensional Markov Chains and their Implementation

A number of frameworks have been proposed since the mid-1990s which supplement or extend the existing fixed-dimension Monte carlo sampling schemes to encompass across model stochastic simulation. Each scheme may be related to others in a conceptually straightforward manner, facilitating natural settings for sampler comparison.

### 2.1 Sampling Frameworks

Introducing the  $\Theta = \bigcup_{k \in \mathcal{K}} \Theta_k \times \{k\}$  formulation of model space, Grenander and Miller (1994) proposed a sampling strategy based on continuous time jump-diffusion dynamics. Such a Markov process essentially jumps between parameter spaces (and therefore models) at random times, and between the jumps follows a diffusion process according to a Langevin stochastic differential equation indexed by time,  $t$ , satisfying

$$d\theta_k^t = dB_k^t + \frac{1}{2} \nabla \log \pi(\theta_k^t) dt \quad (3)$$

where  $dB_k^t$  denotes an increment of Brownian motion, and  $\nabla$  the vector of partial derivatives. In practice (3) is approximated by a discrete-time version with a Metropolis-Hastings step to preserve the stationary distribution  $\pi$  (Roberts and Tweedie 1996). This method has found some application in signal processing and other Bayesian analyses (Miller et al. 1995; Phillips and Smith 1996 for example), but has in general been superceded by the more accessible reversible jump sampler (Green 1995). In fact, correcting for the time-discretisation approximation via the Metropolis-Hastings acceptance probability template, the dump-diffusion sampler can be shown to result in an implementation of the reversible jump algorithm (Besag 1994).

The widely implemented reversible jump sampler was introduced by Green (1995) in a Bayesian model determination setting. Tierney (1998) and Green (2003a) both provide interesting expositions on this theme. One reason for the popularity of this algorithm in particular is conceptual – the framework is a natural generalisation of the standard Markov chain theory, lending it certain appeal. In general, assuming that detailed balance (2) is satisfied we may denote the acceptance probability of a proposed between-model move from  $(\theta_k, k)$  in model  $M_k$  to the state  $(\theta_{k'}, k')$  in model  $M_{k'}$  to be  $\min\{1, A[(\theta_k, k) \rightarrow (\theta_{k'}, k')]\}$ . Here

$$A[(\theta_k, k) \rightarrow (\theta_{k'}, k')] = \frac{L_{k'}(\mathbf{x}|\theta_{k'})p_{k'}(\theta_{k'})\rho_{k'}q(k' \rightarrow k)q_{k'}(u_{k'})}{L_k(\mathbf{x}|\theta_k)p_k(\theta_k)\rho_kq(k \rightarrow k')q_k(u_k)} \left| \frac{\partial g_{k \rightarrow k'}(\theta_k, u_k)}{\partial(\theta_k, u_k)} \right|$$

where  $\theta_{k'} = g_{k \rightarrow k'}(\theta_k, u_k)$  for a random vector  $u_k \sim q_k(u; \psi_k)$  with parameter vector  $\psi_k$ , and  $q(k \rightarrow k')$  is the probability of proposing to move from model  $M_k$  to  $M_{k'}$ . Here  $g_{k \rightarrow k'} : \Theta_k \times \Theta_k^q \rightarrow \Theta_{k'}$  denotes a mapping of the state  $(\theta_k, k)$  together with the vector  $u_k$  to the state  $(\theta_{k'}, k')$ . The mapping satisfies  $g_{k' \rightarrow k}(g_{k \rightarrow k'}(\theta_k, u_k), u_{k'}) = \theta_k$ , and requires that  $n_k + d_k = n_{k'} + d_{k'}$ , where  $d_k$  is the dimension of  $\Theta_k^q$  (known as ‘dimension matching’).

It must be noted that the reversible jump algorithm is not limited to the countable set of models  $\mathcal{M}$ , although it is frequently presented in this context. In fact, one may implement the sampler without knowing the size of the model space beforehand, which may contribute a certain amount to the popularity of the algorithm, although at least some knowledge of the model space is recommended for the construction of an efficient chain. The most common setting involving an uncountable number of models is in Bayesian non-parametrics where both the number of basis functions, and the functions themselves are free to vary — via fractional polynomial regression (Royston and Altman 1994) or free-knot splines for example. See Denison et al. (2002), DiMatteo et al. (2001), Denison et al. (1998) and Smith and Kohn (1996), for useful instances of Bayesian non-parametric curve fitting.

Providing an alternative to samplers designed for implementation on unions of model spaces, a number of approaches have developed conventional Markov chain technology on a product super-model space,  $\Theta^* = \mathcal{K} \times \otimes_{k \in \mathcal{K}} \Theta_k$ , that encompasses all model spaces jointly, thereby circumventing the necessity of between-model transitions. Defining a composite parameter vector,  $\theta^*$ , consisting of a concatenation of all parameters under all models, Carlin and Chib (1995) proposed a formulation in which the posterior distribution for the composite model space was given by

$$\pi(k, \theta^* | \mathbf{x}) \propto L_k(\mathbf{x}|\theta_{\mathcal{I}_k}^*)p_k(\theta_{\mathcal{I}_k}^*)p_k(\theta_{\mathcal{I}_{-k}}^*|\theta_{\mathcal{I}_k}^*)\rho_k,$$

where  $\mathcal{I}_k$  and  $\mathcal{I}_{-k}$  are index sets respectively identifying and excluding the parameters  $\theta_k$  from  $\theta^*$ . Here  $\mathcal{I}_k \cap \mathcal{I}_{k'} = \emptyset$  for all  $k \neq k'$ , so that the parameters for each model are

distinct. In practice, this setting requires that the size of the model space,  $|\mathcal{K}|$ , is finite, thereby restricting its range of application relative to (say) the reversible jump algorithm. Even when this setting is appropriate, a number of impracticalities are apparent for even a moderately sized model space. Although some of the computation in sampling the full parameter vector,  $\theta^*$ , may be avoided (Godsill 2003; Dellaportas et al. 2002; Green and O’Hagan 1998) this approach requires the definition of  $p_k(\theta_{\mathcal{I}_{-k}}^* | \theta_{\mathcal{I}_k}^*)$ , termed pseudo-priors. While their specification is essentially arbitrary in terms of obtaining the desired marginal distributions, sampler performance is sensitive to their specification, introducing practical problems in terms of efficiency and tractability (see Godsill 2003; Green 2003a; Godsill 2001 for a discussion). However, it is believed that, in contrast to the lack of memory of previously visited states inherent in the reversible jump sampler, in the product space formulations (which contain a perfect memory) the information contained within  $\theta_{\mathcal{I}_{-k}}^*$  may be useful in generating efficient between-model transitions when in model  $M_k$ . This idea is exploited by Brooks et al. (2003) which we consider in more detail in Section 3.1.

In a comparison between the methods of Green (1995) and Carlin and Chib (1995), Godsill (2001) proposed a further generalisation of the above which achieves an enhanced flexibility by permitting individual model parameter vectors to overlap arbitrarily. That is, the restriction that  $\mathcal{I}_k \cap \mathcal{I}_{k'} = \emptyset$  for all  $k \neq k'$  is relaxed — this may seem intuitive for, say, nested models. The approach of Godsill (2001) may additionally be shown to encompass the reversible jump sampler, thereby providing a general platform to facilitate comparisons between all previously introduced algorithms.

An alternative approach to the above formulations is based on spatial birth-and-death processes, originally investigated by Preston (1977) and Ripley (1977). Stephens (2000) proposed observing particular trans-dimensional statistical problems in the guise of continuous time abstract marked point processes (see also Geyer and Møller 1994). Finite mixture modelling is one such setting with obvious interpretations for the birth and death of mixture components. Recent work by Cappé et al. (2003) has shown that the sampler of Stephens (2000) may be considered as a particular continuous time limiting version of a sequence of reversible jump algorithms.

A number of illustrative comparisons of the above frameworks can be found in the literature. Andrieu et al. (2001) and Dellaportas et al. (2002) both contrast reversible jump with the pseudo-prior approach of Carlin and Chib (1995), with the former analysis also providing a brief exposition on jump diffusion methodology. Godsill (2001) and Godsill (2003) provide insight into the associations between composite union and product state spaces formulations, and as mentioned above Cappé et al. (2003) examine the relationship between the reversible jump algorithm, the sampler of Stephens (2000) and more general birth and death samplers.

Trans-dimensional sampling algorithms have had an undoubted influence in both statistical and mainstream research literatures. Perhaps unsurprisingly, given the nature of certain technological advances over recent years, is the number of genetical applications relative to other subject areas. Overall, one in every five citations of Green (1995) can be broadly classified as genetics-based research. In general, the large majority of areas in which trans-dimensional Markov chains have strongly benefited to date have tended to be computationally or biologically related. Accordingly a high number of developmental and application studies are to be found in the signal processing literature and the related fields of computer vision and image analysis. Epidemiological and medical studies also feature strongly, and Markov chain Monte Carlo and model selection methodological advances are other obvious

additional high inclusion groupings. With a very few exceptions the overwhelming use of the reversible jump algorithm — and by implication, trans-dimensional samplers in general — has been concerned with the generic problem of model selection.

## 2.2 Finessing Trans-dimensionality

Implementation of trans-dimensional Markov chains typically involves simultaneous exploration of both model and parameter space. However, depending on the aim or the complexity of a multi-model analysis, it may be that using trans-dimensional simulation techniques would be somewhat heavy-handed as occasional reductions to fixed-dimensional simulations may be attained. In some Bayesian model selection settings trans-dimensionality can be avoided if one is prepared to make certain assumptions regarding prior choice, such as conjugacy or objective prior specifications (Berger and Pericchi 2001). Under these settings, explicit expressions for posterior model probabilities,  $m_k(\mathbf{x})$ , may be available (Casella and Moreno 2002, for example). Similarly, it is not uncommon to find situations where approximations to  $m_k(\mathbf{x})$  may be acceptable. In these cases the only “parameter” of interest is the model indicator  $k$ .

A second framework occurs when the full (normalised) conditional distributions,  $\tilde{\pi}_k(\theta_k|\mathbf{x})$ , for each model  $M_k$  are known in closed form. This again is not uncommon in many conjugate models such as linear regression, regression and classification trees, decomposable Gaussian and discrete graphical models, and even exponential family models where simulation from block full conditional distributions is feasible. If the random vector  $u_k \sim q(u; \psi_k) = \tilde{\pi}_{k'}(\theta'_{k'}|\mathbf{x})$  is a draw directly from this conditional distribution, and the proposal state  $\theta'_{k'}$  determined through the mapping  $\theta'_{k'} = g_{k \rightarrow k'}(\theta_k, u_k) = u_k$ , then the reversible jump acceptance probability (2.2) reduces to

$$A[(\theta_k, k) \rightarrow (\theta'_{k'}, k')] = \frac{\rho_{k'} q(k' \rightarrow k) m_{k'}(\mathbf{x})}{\rho_k q(k \rightarrow k') m_k(\mathbf{x})},$$

which is independent of both current and proposed parameter states. The algorithm thereby becomes a fixed dimensional sampler over the space of models (see Clyde 1999b, for example). It is (or it shortly will be) possible to implement both of these simplified simulation frameworks in the popular *WinBUGS* suite — see Section 2.3. When full model conditionals are not available, trans-dimensionality may still be avoided by adopting any of the product space formulations (Brooks et al. 2003; Godsill 2001; Carlin and Chib 1995)

## 2.3 Implementation

While the majority of Bayesian analyses are likely to be novel in some aspect, thereby raising the likelihood that custom-written code is required for their implementation, at the same time it is easily appreciable that certain model types recur frequently enough that new analyses may be based loosely upon them. Depending on the nature of an application it may even be the case that software is already freely available on the web, thereby permitting routine implementation of such algorithms via the code of their authors.

Table 1 details a list of some of the resources currently available from the trans-dimensional sampling community (the majority of fixed-dimensional analyses may be performed using

the *WinBuGS* suite). The implementations most well supported are Gaussian mixture algorithms of varying forms (Cappé et al. 2003; Richardson and Green 1997; Sisson and Fan 2004a) and a number of methods which finesse the trans-dimensional nature of variable selection analyses by integrating out the within-model parameters,  $\theta_k$ , prior to the analysis. The resulting simulations may then be performed in *WinBuGS* (Katsis and Ntzoufras 2003; Ntzoufras 2002, among others). Other useful algorithms implement multiple changepoint analyses Green (1995), the automatic sampler of (Green 2003a) (discussed further in Section 3.3), and an *R* (available from <http://www.r-project.org>) implementation of the trans-dimensional algorithm proposed by Petris and Tardella (2003).

Currently in development for *WinBuGS* is a new ‘*Jump*’ component for the implementation of reversible jump sampling for models in which the full conditional distribution for the within-model parameters is available in closed form (Lunn et al. 2004). As discussed in Section 2.2, utilising these full conditionals as proposal distributions, the resulting simplification in the acceptance probability is then independent of the proposed parameters (which accordingly only require generation *after* the proposed move has been accepted). Such an approach thereby permits an efficient trans-dimensional implementation within the *WinBuGS* framework that retains the full parameter vectors, which is useful in terms of (say) predictive inference.

Even when generic or model-specific software is unavailable, there are many articles in the literature that provide illuminating details on implementations of the various samplers. Table 2 provides a brief selection of just some of these in a range of modelling scenarios, and the interested reader is directed to the details therein. In practice however, it is apparent that the majority of routine data analyses are performed using the reversible jump algorithm given its more accessible nature. This algorithm is certainly the most understood and developed of all trans-dimensional sampling algorithms.

Despite the general dominance of the reversible jump sampler, there are a number of situations in which adoption of one of the alternative algorithms may provide a simpler or more intuitive implementation. For example, when the total number of models is relatively small, so that the length of the composite parameter vector  $\theta^*$  is moderate, then adopting one of the product space formulations may be useful (Ntzoufras 2002, for example). Similarly, the birth-and-death approach has found some application in model settings that may be more naturally expressed in the point-process setting (Stephens 2000; Cappé et al. 2003), although these tend to be problem specific (see Table 2 for references to specific illustrations). Jump-diffusion methods, however, are more easily conceived in the discrete time setting. As a consequence, they have tended to be superseded in their application by the reversible jump framework. Despite occasional implementation in applied settings, in general product-space approaches have found their greatest utility for sampler developmental purposes given the insights they provide into the relationships between the different sampling algorithms. We examine some of the resulting technological improvements in Section 3.1.

With the theory of model-spanning sampling frameworks now well established in the statistical literature, the main driving force of research in this field concerns the nature of difficulties encountered in their implementation. Issues of efficiency and mixing translate over from fixed-dimensional sampling schemes, as do problems in assessing convergence, with the additional obstacle of a substantial increase in the complexity of the problem. Between-model transitions play a key role in this setting. A related matter is the necessary specification of tuning parameters and the form of between model mappings  $g_{k \rightarrow k'}$  — features we examine

in closer detail in Section 3.

## 2.4 Exploring Model Space

When the number of candidate models is large, model selection methods are generally concerned with the maximisation (or minimization) of model ranking functionals according to a non-deterministic optimisation process (George and McCulloch 1993; Chipman et al. 2001, for example). As a means to automate model selection Brooks, Friel, and King (2003) (see also Andrieu et al. 2000) proposed an extension to the standard simulated annealing framework by constructing a trans-dimensional Markov chain with stationary distribution proportional to the Boltzmann distribution

$$\mathcal{B}_T(\theta_k, k) \propto \exp\{-f(\theta_k, k)/T\}, \quad (4)$$

where  $T \geq 0$  and  $f(\theta_k, k)$ ,  $(\theta_k, k) \in \Theta$ , is a model ranking function to be minimised. A stochastic annealing framework may then be defined by periodically decreasing the value of  $T$  according to some schedule while using the Markov chain to explore functional space. As  $T \rightarrow 0$  the distribution (4) converges to a point mass at  $(\theta_{k^*}^*, k^*) = \arg \max f(\theta_k, k)$ . Assuming adequate chain mixing the algorithm will thereby identify the model determined by  $f$ . Applied to ecological capture-recapture analyses this methodology has achieved success in facilitating the classical model selection procedure according to the Akaike information criterion (Sisson and Fan 2004b; King and Brooks 2004) by setting  $f(\theta_k, k) = -2 \log L_k(\mathbf{x}|\theta_k) + 2n_k$ . In Bayesian analyses one natural choice would be to consider  $f(\theta_k, k) = M_k(\theta_k)$  — the posterior model probability (Clyde 1999a, for example). Stochastic optimisation techniques such as the above are not limited to the classical nor model selection frameworks. Under Bayesian decision theoretic settings where loss functions defined on variable dimensional space take non-standard form (Sisson and Hurn 2004 for example), there is an obvious benefit in adopting flexible optimisation methods for the derivation of Bayes rules.

Unfortunately, selection of the model with the highest scoring model ranking functional need not necessarily be the most useful in terms of a range of criteria. In the Bayesian normal linear model framework, Barbieri and Berger (2004) discuss optimality conditions when a single model must be chosen for predictive purposes. In particular, they were unable to identify general conditions under which the optimal predictive model coincided with that of the highest posterior probability model (although this does occur in some cases), but instead concluded that such optimality theorems in fact existed for the median probability model, defined to consist of those variables whose posterior probability of inclusion is at least  $1/2$ . In the classical analysis setting, Sisson and Fan (2004b) derive a sequence of profile models  $\mathcal{M}^p \subset \mathcal{M}$  such that  $M_p^p \in \mathcal{M}^p$  is defined to consist of those variables whose “probability” of inclusion is at least  $p$ , for all  $p \in [0, 1]$ , and then perform a reduced model search on this set. Here their aim was to avoid the determination of over-fitted models, resulting from use of the AIC, in order to retain (in this case) biological interpretability. These two examples each illustrate subtle concerns with the process of model selection. The latter of these is indicative of the problems faced in the realistic specification of the model ranking functional  $f$  when the number of candidate models is large, in analogy with meaningful prior specification in the Bayesian setting. The former example underlines that the procedure is not necessarily straight-forward, even when natural model-ranking measures are available.

## 2.5 Assessing Convergence

Perfect sampling schemes aside (see Section 3.2), under the assumption that an acceptably efficient method of constructing a trans-dimensional sampler is available, one obvious prerequisite to inference is that the Markov chain converges to its equilibrium state. This is a contentious issue even for fixed dimensional cases. *A priori* convergence bounds are in general difficult or impossible to determine; *a posteriori* convergence diagnostics assess necessary rather than sufficient indicators of chain convergence (see, for example, Mengersen et al. 1999; Cowles and Carlin 1996 for comparative reviews). The trans-dimensional setting generates additional concerns — in particular, how might one assess convergence not only within each of a potentially large number of models, but also across models with respect to posterior model probabilities?

One natural approach would be the implementation of independent sub-chain assessments, both within models and for the model indicator. However, in isolation this would erroneously associate convergence of the full chain with those of multiple sub-chains, thereby generating potential for the underestimation of convergence time for the full density. For example, with focus purely on model selection, Brooks, Giudici, and Philippe (2003) propose various ideas based on the sample-path of the model indicator,  $k$ , under the assumption that replicated chains that have converged will generate similar posterior model probability estimates. The focus on the model indicator permits the application of a number of non-parametric hypothesis tests — the chi-squared and Kolmogorov-Smirnov tests are discussed in detail. In this manner, distributional assumptions of the models (but not the statistics) are circumvented at the price of associating marginal with full density convergence.

Additional sampler performance issues arise as sparsely realised models of low posterior probability will be poorly represented (if at all), presenting an obstacle to even marginal assessment (Brooks 1997). One further problem is that of credibly assessing chain convergence of complex equilibrium distributions defined over high-dimensional state spaces by a univariate statistic, however well chosen.

Given the involved difficulties, it is hardly notable that to date there are relatively few diagnostics specifically designed for trans-dimensional samplers. One strategy, proposed by Brooks and Giudici (2000), requires the determination of a model reparameterisation such that as many model parameters as possible,  $\theta^+$ , retain their interpretation for all models under consideration. These parameters may then be monitored to provide an indication of chain performance. Specifically, Brooks and Giudici (2000) suggest a two-way ANOVA decomposition of the variance of a functional,  $h(\theta^+)$ , over multiple chain replications. A similar approach was advocated by Castelloe and Zimmerman (2002) who address the observed sensitivity of the Brooks and Giudici (2000) method to imprecise sample means from rarely visited models. They argue that a single visit to a rare model in one chain should not overwhelmingly dominate the diagnostic, and accordingly develop an unbalanced two-way ANOVA, with weights constructed in proportion to the frequency of model visits. Castelloe and Zimmerman (2002) also extend their methodology to the multivariate (MANOVA) setting on the observation that monitoring several functionals of marginal parameter subsets is more robust than monitoring a single statistic.

While both groups of authors identify useful statistics to monitor, and Castelloe and Zimmerman (2002) offer some innovation in surmounting a problem of label-switching, one immediate problem with this approach in general is the difficulty identifying the requisite

parameter set. A lesser issue regards the extent of approximation induced by violations of the ANOVA assumptions of independence and normality. Even ignoring the underlying implication of marginal assessment, the issue of parameter selection is magnified when considering that even common parameters may change meaning from one model to another (Berger and Pericchi 1996, for example). This leads naturally to statistics,  $h$ , based on fitted and predicted values of observations as the obvious choice in many cases, reducing the problem to the fixed dimension setting (Green 2003a).

Sisson and Fan (2004a) suggest a method that circumvents the trans-dimensional nature of the problem when the underlying model can be formulated in the marked point process framework of Stephens (2000). Specifically, the differences in intensity functions between chain replicates are determined by statistics based on the distributions of point-to-nearest-neighbour distances, thereby permitting the direct comparison of parameter vectors of varying dimension and, as a result, naturally incorporating a measure of across model convergence. Due to the manner of their construction, Sisson and Fan (2004a) are able to monitor an arbitrarily large number of such functionals. While this approach may have some appeal, it is limited by the need to construct the model in the marked point process setting.

Given the spectrum of difficulties involved in the performance monitoring of trans-dimensional samplers, and in the obligingly small suite of tried-and-tested diagnostic tools available for this task, the ever increasing numbers of articles containing analyses based on these sampling methods is perhaps some cause for concern. The most common approach adopted in the various literatures using the reversible jump sampler (Jackson and Sharples 2004; Bottolo et al. 2003; Salmenkivi et al. 2002; Suchard et al. 2001; Kasuya and Takagawa 2001, for example) rests on the monitoring of those parameters which “retain their interpretation as the sampler moves between topologies and [which] may be used effectively to monitor how well the MCMC sampler is performing” (Suchard et al. 2001), despite the aforementioned difficulties in selecting such parameters, and the questionable “effectiveness” which marginal monitoring may provide. These parameters, generally a small subset of the full parameter set, are then monitored using popular fixed-dimensional performance measures (Smith 2001; Brooks and Gelman 1998; Gelman and Ruben 1992; Geweke 1992, are typical), although in many cases this analysis is limited to a single diagnostic. Of course, there is an obvious danger in monitoring only a single diagnostic to evaluate sampler performance — as much so as relying on marginal assessment. For instance, the finite mixture of normals (reversible jump) algorithm of Richardson and Green (1997) is perhaps the most studied in terms of its performance via a number of diagnostics (Sisson and Fan 2004a; Brooks et al. 2003; Brooks and Giudici 2000; Richardson and Green 1997). However, in spite of the general perception that the sampler is fairly efficient, there is less unanimity on exactly “when” convergence may have been attained, depending on the diagnostic implemented.

While it is undeniable that the benefits for the practitioner in implementing trans-dimensional sampling schemes are immense, it would seem arguable that the practical importance of ensuring chain convergence is often overlooked. More charitably perhaps, it is more likely the case that current technology is insufficiently advanced to permit a more rigorous default assessment of sampler convergence, and until this shortcoming is resolved conscientious practitioners will be obliged to manage with the best that is currently available.

## 2.6 Estimating Bayes Factors

When the number of candidate models,  $|\mathcal{K}|$ , is large, the use of trans-dimensional sampling algorithms to evaluate Bayes Factors between competing models raises issues of efficiency. In moving from  $(\theta_k, k)$  in model  $M_k$  to  $(\theta_{k'}, k')$  in model  $M_{k'}$ , Bartolucci and Scaccia (2003) demonstrate that the expected value of the reversible jump acceptance probability under the distribution  $f_k(\theta_k, u_k) = \pi_k(\theta_k|\mathbf{x})q_k(u_k)$  is bounded above by  $B_{k'k}\rho_{k'}/\rho_k$ . Accordingly, if model  $M_k$  accounted for a large portion of posterior mass, the reversible jump algorithm will tend to persist in model  $M_k$  and visit others models rarely and, as a consequence, resulting estimates of Bayes Factors based on model-visit proportions will tend to be inefficient (Han and Carlin 2001 for example). That is, the auxiliary random process adopted for transitions between models increases the variability of the estimator. In contrast, individually estimating the marginal model probabilities  $m_k(\mathbf{x})$  and  $m_{k'}(\mathbf{x})$  (Chib and Jeliazkov 2001; Chib 1995) or their ratio (Mira and Nicholls 2004; Meng and Schilling 2002; Chen and Shao 1997) via independent fixed-dimension simulations is more precise, although impracticalities emerge when the model space is large.

In an interesting recent development, Bartolucci and Mira (2003) (see also Bartolucci and Scaccia 2003), propose an extension to the Bridge estimator for the estimation of the ratio of normalising constants of two distributions (Meng and Wong 1996). The authors augment the state spaces of the two distributions in the exact manner that is implicit in the specification of the auxiliary variables  $u_k, u'_{k'}$  in the reversible jump algorithm, so that the distributions of interest,  $f(\theta_k, u_k)$  and  $f(\theta_{k'}, u'_{k'})$ , are defined upon the same, but not necessarily nested, space. Accordingly using realisations from these two distributions, possibly directly from the reversible jump sampler, the estimator of Bartolucci and Mira (2003) essentially integrates out the auxiliary random process, which thereby depends on a property of the acceptance probabilities and consequently provides more efficient estimates.

## 3 Future Perspectives - Towards Automation

Since their inception there has been a concerted drive to design sampling algorithms, both fixed- and trans-dimensional, that require the minimal initialisation overheads but that achieve the maximum in efficiency. This effort has manifest itself in a number of different ways, but to a greater extent current research is focused on embedding varying degrees of automation into the sampling process. Whether through default prior specification and one-size-fits-all generic algorithms, via methods that adapt chain performance to enhance efficiency during implementation, or by attempting to circumvent issues of convergence entirely, it is clear that the benefits in achieving fully automatic trans-dimensional algorithms are immense for the statistical community as a whole, and for those in other disciplines who wish to use these methodologies.

### 3.1 Efficient Chain Construction

The popularity, conceptual simplicity and obvious utility of trans-dimensional samplers notwithstanding, they have gained a reputation of an inefficiency in implementation and a poor performance in general. This is perhaps a somewhat unfair assessment – it would seem credible that at least some techniques of chain construction could be determined that

would generate acceptable performance given that, for example, the reversible jump algorithm encompasses all Metropolis-Hastings methods for between-model state transitions (Green 2003a). Failure to achieve acceptable performance could be considered merely a result of poorly constructed between-model transitions or inappropriate tuning of proposal distributions. Some methods targeted specifically at improving the acceptance rate of between-model transitions include the multi-step candidate generating scheme of Al-Awadhi et al. (2004), and the delayed rejection sampling scheme of Green and Mira (2001). Mengersen and Robert (2003) describe a self-avoiding population Monte Carlo scheme aimed at increasing exploration of the state space.

It should perhaps even be anticipated that implementation of a trans-dimensional sampling scheme may result in enhanced mixing, even when applied in a fixed dimensional setting. In this case, sampling from a single model distribution with a more sophisticated machinery might be considered loosely analogous with the extra power gained with augmented state space sampling methods. For example, in the case of a strongly multi-modal posterior, Richardson and Green (1997) report markedly superior mixing in examining the output of a reversible jump algorithm conditioned upon there being exactly three mixture components, in contrast to the output generated by the respective fixed-dimension sampler. Similarly, George et al. (1999) achieved improved chain performance in an analysis concerning the ordering of genetic markers – a fixed-dimensional model – although in this case the ‘birth’ and ‘death’ moves were preformed simultaneously so that the dimension of the model remained constant. In a short study, Green (2003a) presents a discussion on which inferential circumstances may determine whether the adoption of a trans-dimensional sampler may be beneficial (although see Han and Carlin 2001 for an argument to suggest that trans-dimensional sampling may have a detrimental effect on efficiency).

One complicating factor arises when the number of candidate models is considerable. Here, the state space of the between-model structure may become difficult to visualise, causing problems in aligning regions of high probability and thereby in the construction of efficient proposal templates – although this factor may be of less concern under specific models settings, such as nested models. In addition, the task of manually tuning between-model transition variables via repeated pilot runs of the chain can become laborious, and quickly prohibitive. There is therefore a strong argument for continued research into the development of assisted or automated proposal generation, for both standard Metropolis-Hastings methods (Roberts et al. 1997, for example), and for trans-dimensional sampling schemes in particular.

Recently, Brooks, Guidici, and Roberts (2003) (see also Ehlers and Brooks 2003) introduced a number of methods to achieve the automatic scaling of the proposal density. For a proposed move from  $(\theta_k, k)$  in  $M_k$  to model  $M_{k'}$  one technique is based upon identifying the random vector “centering point”,  $c(\theta_k) = g_{k \rightarrow k'}(\theta_k, u_k(\theta_k))$ , such that for some particular choice of proposal vector,  $u_k(\theta_k)$ , the current and proposed states are identical in terms of likelihood contribution. That is,  $L_k(\mathbf{x}|\theta_k) = L_{k'}(\mathbf{x}|c(\theta_k))$ . Given the constraint on  $u_k(\theta_k)$ , the relevant scaling parameters are then obtained to ensure that  $A[(\theta_k, k) \rightarrow (c(\theta_k), k')] = 1$ . For example, in the case of autoregressive regression if  $M_{k'}$  is a higher order model than  $M_k$  then we may have  $u_k(\theta_k) = 0$  as  $L_k(\mathbf{x}|\theta_k) = L_{k'}(\mathbf{x}|(\theta_k, 0))$ . See also Ntzoufras et al. (2003) for a centering method in the context of linear models.

A similarly motivated method is based on a series of  $n^{th}$ -order conditions (for  $n \geq 1$ ), which require that for the proposed move, the  $n^{th}$  derivative of the acceptance probability

equals the zero vector at  $c(\theta_k)$ :

$$\nabla^n A[(\theta_k, k) \rightarrow (c(\theta_k), k')] = 0. \quad (5)$$

Similarly to the above, appropriate values for the proposal parameter vector,  $\psi_k$ , are determined via (5), such that are likely to generate acceptance ratios close to one within a region. In this manner, proposal parameters are adapted to the current state of the chain at each stage rather than relying on a constant proposal parameter vector for all state transitions. It can be shown that for a simple two model case, the above conditions are optimal in terms of the capacitance of the algorithm (Lawler and Sokal 1988).

Brooks, Guidici, and Roberts (2003) also propose a second class of models based on augmenting the state space with an auxiliary set of state-dependent variables,  $v_k$ , so that the state space of  $\pi(\theta_k, v_k | \mathbf{x}) = \pi_k(\theta_k | \mathbf{x})\tau_k(v_k)$  is of constant dimension for all models  $M_k \in \mathcal{M}$ . Although this fixed dimensionality is later relaxed, there is obvious analogue with the product space frameworks of Carlin and Chib (1995) and Godsill (2001). By considering updates of  $v_k$  via a slowly mixing Markov chain with a Gaussian stationary distribution a temporal memory is induced that persists in the  $v_k$  from state to state. In this way, the idea behind the auxiliary variables is to assist in between-model proposal transitions in that some memory of previous model states is retained. The authors demonstrate that this approach can significantly enhance mixing compared to an unassisted reversible jump algorithm.

There is an obvious utility in the above and other approaches for generic proposal design. However, one caveat with the above schemes is that they assume prior specification of the between model mapping  $g_{k \rightarrow k'}$ , and consequently that interest lies primarily in maximising the between-model acceptance probability. When the candidate models have a strong degree of mutual consistency — so that there exist well defined functionals of parameters with consistent meaning across models, and when prior specifications regarding these functionals are also consistent across models — the decomposition of the proposal model  $M_{k'}$  into prior and likelihood terms as proposed by Brooks et al. (2003) is likely to generate natural local mappings between models (Green 2003b). This can be observed in moment-matching methods such as the well known “split-and-merge” move-types. When mutual consistency between models is lacking however, or where the mappings,  $g_{k \rightarrow k'}$ , are sub-optimal or even wholly undetermined, there is a clear limitation in the degree of chain efficiency that may be realised and it is difficult to envisage how such schemes may assist.

In this sense, while maximising within-model acceptance probabilities through local structural proposals will be useful in a broad range of modelling scenarios, more moderate acceptance rates, which aim to balance comprehensive within- and between-model transitions, may offer enhanced efficiencies in more general settings. (A useful, and familiar analogy can be found in the fixed-dimensional Metropolis-Hastings setting, whereby it is trivial to ensure a near 100% acceptance rate, but at the expense of poor exploration of the state space Gelman et al. 1996, for example). From this perspective, a combination of both structurally local and more global between-model move-types which do not rely on structural knowledge of the models in order to specify the between-model mapping, may provide the optimum specification. The case of automatically generating global between-model proposals is examined further in Section 3.3.

## 3.2 Perfect Sampling

Publication of the seminal paper by Propp and Wilson (1996) (see also Fill 1998) on the subject of stochastic sampling schemes which aim to draw realisations *exactly* from the stationary distribution of a Markov chain (see for example Dimakos 2001; Casella et al. 2000 for a review<sup>1</sup>) generated considerable interest in the statistical community, based on the hope that such frameworks would become panacea to issues of convergence. Despite considerable research in this area, perfect sampling schemes have proved difficult to implement in all but the simplest of modelling situations. This coupled with the non-modularity of the proposed algorithms, has resulted in a general reticence to embrace them as a mainstream technology. The increasing resemblance of exact simulation methods to standard algorithms involving Metropolis-Hastings steps is even leading some researchers to the opinion that little further may be gained by such methods in terms of ease and speed of implementation (Robert 2003). Given the enhanced problems associated with the convergence of trans-dimensional Markov chains (Section 2.5), an exact multi-model algorithm would appear even more desirable than in the fixed dimensional setting.

To date, the only methodological approach proposed in which an extension to a multi-modal setting is explicitly discussed (Brooks, Fan, and Rosenthal 2002; Møller and Nicholls 1999) is based on an extension of the simulated tempering algorithm of Geyer and Thompson (1995). Intuitively, the fixed-model case proceeds by moving through the augmented state space of heated models, each indexed by a temperature  $\tau \in \mathcal{T} = \{1, 2, \dots, \tau_{\max}\}$  (for example), and at each stage proposing a move to the distribution, indexed by  $\tau = \tau^*$ , from which it is possible to sample directly. The smallest possible probability of this occurring at any stage, assuming an identical state space  $\Theta^T$  for all tempering models, is given by

$$\epsilon = \inf_{\theta \in \Theta^T, \tau \in \mathcal{T}} q(\tau \rightarrow \tau^*) \min \{1, A[(\theta, \tau) \rightarrow (\theta, \tau^*)]\}.$$

Assuming that all possible chains commencing at time  $t = -\infty$  coalesce into a single chain in model  $\tau^*$  with probability  $\epsilon$ , focus would then naturally be upon the first instance this occurred in the reverse-time chain from  $t = 0, -1, -2, \dots, -T$ . By construction  $T \sim \text{Geometric}(\epsilon)$ . Consequently, commencing a forward-time Markov chain sampler starting in model  $\tau^*$  at time  $t = -T$  will have the effect of generating a draw exactly from the stationary distribution at time  $t = 0$ . Extension to the multi-model setting is immediate, requiring only a substitution of the temperature-augmented models by the set of candidate models  $\mathcal{M}$  indexed by  $k \in \mathcal{K}$ , from which it is assumed that there is one model,  $M_{k^*}$ , from which it is possible to sample from directly. The algorithm then proceeds as before, following a trans-dimensional sampling scheme, where  $\epsilon$  is now given by

$$\epsilon = \inf_{(\theta_k, k) \in \Theta, u_k \in \mathcal{Q}_k} q(k \rightarrow k^*) \min \{1, A[(\theta_k, k) \rightarrow (g(\theta_k, u_k), k^*)]\}, \quad (6)$$

the smallest possible probability moving to model  $M_{k^*}$  over all possible models and states  $(\theta_k, k) \in \Theta$ , and random vectors  $u_k \sim q(u; \psi_k)$  with  $u_k \in \mathcal{Q}_k$ .

While theoretically intuitive and appealing, there are a number of issues that restrict implementation of this algorithm in the general setting, most notably surrounding determination of  $\epsilon$  in (6). Problems arise in the sense that the model  $M_{k^*}$  may not be known *a*

---

<sup>1</sup>A valuable perfect sampling resource is <http://dbwilson.com/exact/> maintained by D. Wilson.

*priori*, nor the model,  $M_k$ , which results in realising the infimum. Additionally  $\epsilon$ , or a more computationally accessible approximation  $\epsilon' \leq \epsilon$  may so small as to make simulation infeasible. While the situation may be partially simplified in certain cases, such as nested models, it remains to be demonstrated that the above approach (or any other) may be efficiently implemented in a trans-dimensional setting.

Given the nature of the problems of chain mixing and efficiency, one approach to good sampler construction would be to identify relevant issues for a given analysis and develop of situation-specific solutions. This may be less than ideal for the MCMC novice wishing to use trans-dimensional sampling methods, as they may not be adequately knowledgeable of the necessary implementational details, nor other specifications beyond that of the actual model to be analysed – naturally this is not (yet) recommended in practice! An alternative approach might be to seek more general solutions, by developing methods that aim to circumvent such implementation issues. We now consider some generic schemes which aim to adopt just such an automatically calibrated system.

### 3.3 Generic Samplers

The problem of efficiently constructing between-model mapping templates,  $g_{k \rightarrow k'}$ , as previously outlined may be approached from an alternative perspective. Rather than adapting the properties of a fixed proposal mechanism to maximise the probability of accepting model jumping transitions (Section 3.1), one possibility would be to remove the reliance on a user-specified method (such as birth/death or split/merge transitions) altogether. While such an attractive ideal currently remains on the research horizon, a number of preliminary approaches to the development of generic automatic sampling frameworks have recently been proposed.

A reversible jump analogy of the random-walk Metropolis sampler of Roberts (2003) was proposed by Green (2003a). For each of a small number of models  $\mathcal{M}$  suppose that estimates of the first and second order moments of  $\theta_k$  are known — denote these by  $\mu_k$  and  $B_k B_k^T$  respectively. In proposing a move from  $(\theta_k, k)$  to model  $M_{k'}$  a new parameter vector is generated by setting

$$\theta'_{k'} = \begin{cases} \mu_{k'} + B_{k'} [RB_k^{-1}(\theta_k - \mu_k)]_1^{n_{k'}} & \text{for } n_{k'} < n_k \\ \mu_{k'} + B_{k'} RB_k^{-1}(\theta_k - \mu_k) & \text{for } n_{k'} = n_k \\ \mu_{k'} + B_{k'} R \begin{pmatrix} B_k^{-1}(\theta_k - \mu_k) \\ u_k \end{pmatrix} & \text{for } n_{k'} > n_k \end{cases}$$

where  $[\cdot]_1^m$  denotes the first  $m$  components of a vector,  $u_k \sim q_k(u; \psi_k)$  is an  $(n_{k'} - n_k)$ -dimensional vector of random numbers, and  $R$  is a orthogonal matrix of order  $\max\{n_k, n_{k'}\}$ . If the marginals  $\pi_k(\theta_k | \mathbf{x}) \sim \text{MVN}(\mu_k, B_k B_k^T)$  and  $u_k$  is a standard normal vector, choosing an appropriate model proposal density of  $q(k \rightarrow k')/q(k' \rightarrow k) = \pi_{k'}(\mathbf{x})/\pi_k(\mathbf{x})$  would ensure that the acceptance probability  $\min\{1, A[(\theta_k, k) \rightarrow (\theta_{k'}, k')]\}$ , with

$$A[(\theta_k, k) \rightarrow (\theta_{k'}, k')] = \frac{\pi_{k'}(\theta'_{k'} | \mathbf{x}) q(k' \rightarrow k) |B_{k'}|}{\pi_k(\theta_k | \mathbf{x}) q(k \rightarrow k') |B_k|} \times \begin{cases} q_k(u_k) & \text{for } n_{k'} < n_k \\ 1 & \text{for } n_{k'} = n_k \\ 1/q_k(u_k) & \text{for } n_{k'} > n_k \end{cases}$$

would equal unity. That is, the sampler would automatically achieve the detailed balance condition, and is the motivation for the author. The implication is that high transition

probabilities may be achieved when the marginal distributions  $\pi_k$  are uni-modal with first and second order moments given by  $\mu_k$  and  $B_k B_k^T$ . Green (2003a) discusses a number of modifications to this general framework and illustrates it via variable selection and change-point problems.

A related sampler is proposed by Godsill (2003) who, in adopting standard Gaussian approximation arguments, suggests the proposal generating mechanism

$$\theta'_{k'} = \mu_{k'} + B_{k'} v_{k'},$$

where  $v_{k'} \sim q_{k'}^*$  is an  $n_{k'}$ -dimensional standard normal vector, which has similar detailed balance properties to the Green (2003a) sampler. Although a detailed comparative study has not yet been implemented, differences between the two samplers emerge when the target densities are non-Gaussian, in which there is a trade-off between both the target distribution ratios, and the difference in variability between the ratio  $q_{k'}^*(v_{k'})/q_k^*(v_k)$  and  $q_k(u_k)$  which is generally of lower dimension (Godsill 2003).

The above samplers have a number of obvious restrictions. Primarily they involve knowledge of the first and second order moments of the parameters under each model – something that may be difficult to attain for anything but a small number of models — Green (2003a) obtains estimates of these via pilot chains on each model. The assumption of uni-modality is also an important factor — the departure of the proposal distribution from the true conditional density strongly affects the acceptance rate of the algorithm. On a (multi-modal) change-point analysis Green (2003a) demonstrates a relative efficiency of 29% of the automatic sampler compared to the results of a standard reversible jump sampler (Green 1995), but with reduced implementation time. Such caveats notwithstanding, the avoidance of the ‘necessity’ of specifying a between-model parameter mapping  $g_{k \rightarrow k'}$  is an important step in the development of future sampling technologies.

### 3.4 Automatic and Objective Prior Specification

One of the obvious problems with the Bayesian approach to inference, in contrast to the sampler-based mechanics considered until now, is prior elicitation. This is particularly a problem in analyses when the number of candidate models is large, as a specification that accurately represents given prior knowledge for all parameters and all models is typically infeasible. The use of non-informative or improper priors is therefore attractive for this task, and in general forms the basis of an analysis. The drawback to model selection with such priors is that the Bayes factor is arbitrarily determined, leading to a number of alternative default methods such as default *proper* priors (Zellner and Siow 1980, for example), fractional Bayes factors (O’Hagan 1995) and intrinsic Bayes factors (Berger and Pericchi 1996). See, for example, Berger and Pericchi (2001) for an extensive discussion.

One recent development of the intrinsic Bayes factor approach, that is particularly attractive from a computational viewpoint, is the expected-posterior prior approach of Pérez and Berger (2002). Assuming standard non-informative priors,  $p_k^N(\theta_k)$  for each model, the expected-posterior prior for  $\theta_k$  under  $m^*$  is defined as

$$p_k^*(\theta_k) = \int p_k^N(\theta_k | y^*) m^*(y^*) dy^*, \quad (7)$$

where  $y^*$  denotes a minimal vector of ‘imaginary’ training data,  $p_k^*(\theta_k|y^*)$  is the posterior distribution of the parameters  $\theta_k$  given the training data, and  $m^*$  is a predictive measure for  $y^*$ .

Assuming that the (possibly improper) default priors  $\pi_k^N$  can be automatically determined, construction of (7) requires only the specification of  $m^*$ , which may itself be potentially determined given the modelling situation. For instance, in the case of nested models, one such choice might be

$$m^*(y^*) = \int L_1(y^*|\theta_1)p_1(\theta_1)d\theta_1,$$

the predictive density of the training data under the base model,  $M_1$ . Most significantly for the expected-posterior prior approach, the generated prior distributions for different models are appropriately compatible, and problems of impropriety are also avoided — see Pérez and Berger (2002) for full details. Perhaps what makes the above prior specification particularly attractive, is that due to the probabilistic nature of the specification (7) their usage is particularly suited to incorporation in hierarchical Markov chain sampling frameworks (Pérez and Berger 2001, for example). The imaginary training sample  $y^*$  is simply considered as one of the unknown parameters in the simulation.

### 3.5 Adaptation

The ability of a Markov chain to satisfactorily traverse model space is further complicated given that models with high posterior probability need not exhibit structural similarities and correspondingly may reside in relatively disparate portions of model space (see also Section 3.3). This creates obvious problems for the construction of between model proposals  $q(k \rightarrow k')$  and parameter mappings,  $g_{k \rightarrow k'}$ , as it is natural to consider model-transitions based on local perturbations, both for conceptual ease and for ensuring a reasonable likelihood of accepting the proposed move. A similar situation occurs in fixed-dimensional settings in the presence of a strongly multi-modal posterior.

Most research to address this problem has been conducted in the fixed dimensional under the generic label of “adaptive” MCMC, which seeks to use the full sample path of the Markov chain to construct an efficient proposal density during chain implementation. See, for example, Frigessi (2003) for a concise statement on the current state of the art. Briefly, the main issues center on the extent to which the stationary distribution adheres to the desired posterior given the manner of adaptation. Care must be taken not to adapt too quickly or inconsistently, or the wrong target distribution may be attained; a result that is all too easily achieved (see Atchade and Rosenthal 2003 for example). Subject to assumptions of uniformly ergodic transition kernels and bounded state spaces, the adaptive algorithm of Haario et al. (2001) which depends on the full history of the chain can be directly shown to yield unbiased Monte Carlo estimates of the expectations of bounded functionals (see Andrieu and Moulines 2002, Atchade and Rosenthal 2003 for results in more general settings). In comparison, if so-called regeneration points exist — such as an independent sample drawn from the ‘hot’ distribution in a simulated tempering algorithm (Tierney 1996; Brooks et al. 2002), or an atom in the state space — the adaptation may be implemented at these times. The dependence on the full chain history is consequently mitigated and the Markovian structure is preserved (Sahu and Zhigljavsky 2003; Gåsemyr 2003; Gilks et al. 1998). Frigessi (2003) suggests there may be scope for development in adopting  $d^{th}$ -order Markov chains whose

stationary distribution may be slightly biased, but which is mathematically more flexible than chains based on the full sample-path history. This could conceivably be extended to variable-length Markov chains (Bühlmann and Wyner 1999).

While currently centered on the fixed dimensional problem it is easily envisaged that adaptive methods will eventually graduate to the trans-dimensional setting, permitting the construction of between-model proposals that increasingly resemble full model conditionals. In situations where the trans-dimensional nature of an analysis may be avoided (Section ??) this is already possible. For example, in a Gaussian variable selection setting using intrinsic priors — a fixed dimensional problem as posterior model probabilities are known — Casella and Moreno (2002) suggest a method of ensuring that the model proposal density approximates the posterior model probability  $q(k \rightarrow k') \approx M_{k'}(\mathbf{x})$  as the chain length  $N \rightarrow \infty$ . In this manner the suggested transitions density  $q(k \rightarrow k')$  seeks to avoid the problems associated with local modes of model space by proposing candidate models approximately in proportion to their posterior probabilities.

## 4 Discussion

One of the fundamental goals of trans-dimensional sampling frameworks is to achieve high degrees of both efficiency and automation. In addition to providing a survey of the past decade of progress towards this goal, in this article we hope to have presented a discussion on the some of problems associated with attaining this objective and illustrated some of the most recent attempts to engage it. Each of the areas highlighted has considerable potential for further development and innovation.

Given the degree of implementational difficulty associated with such methodologies, the aspect of automation of default model and chain specifications is particularly important with regard to the extent to which trans-dimensional Markov chains have been embraced within a broad array of application areas across a wide variety of disciplines. As a fundamental rule, it is important that statistical techniques are accessible to those who wish to adopt them while stopping short of black-box implementations — public domain software packages, such as *WinBUGS*, are increasingly instrumental in propagating new methodologies through disparate disciplines. An all-in-one software package implementing automated and generic sampling schemes would surely prove a formidable resource for Bayesian analyses.

It might be imagined under such a scenario, that specification of the desired models would be made via a graphical user interface or scripting language, with the option of highlighting a checkbox labelled “Adopt Objective Priors.” The software would then determine the most efficient within- and between-model transitions and implement the sampler, possibly even determining before sampling commences whether a perfect or standard sampling scheme will offer the greatest efficiency for a given computational time. While such a software package is currently some way from realisation, a number of recent innovative papers discussed in this article have taken small, but confident strides in this direction. In doing so, they have generated heightened interest and enthusiasm in this goal.

As a closing note, while we have focused on reversible Markov chains defined by satisfying the detailed balance condition, there is a small literature that suggests that non-reversible chains may offer improvements in efficiency not available to their more accessible reversible counterparts. For example, Diaconis et al. (2000) show that non-reversibility can lead to

improvements over the diffusive behavior of simple Markov chain sampling schemes. Neal (2004) (see also Mira and Geyer 2000) demonstrates that any reversible Markov chain on a finite and irreducible state space may be used to construct a non-reversible Markov chain on a related state space with asymptotic variance at least as small as that using the reversible chain, although typically this will be much smaller. The author concludes that “this construction demonstrates that non-reversible chains have a fundamental advantage over reversible chains for MCMC estimation. Research into better MCMC methods may therefore best be focused on non-reversible chains.”

## Acknowledgements

The author wishes to thank Y. Fan and two referees for many useful comments on an earlier draft of this article.

Green (2003a)	<i>AutoRJ</i> : Automated reversible jump MCMC
Richardson and Green (1997)	<i>Nmix</i> : Bayesian analysis of univariate normal mixtures
Green (1995)	<i>Cpt</i> : Bayesian multiple change-point analysis for point processes Available from: <a href="http://www.stats.bris.ac.uk/~peter">http://www.stats.bris.ac.uk/~peter</a>
Cappé et al. (2003)	<i>CT/RJ-Mix</i> : Continuous time and reversible jump samplers for Gaussian mixtures Available from: <a href="http://www.tsi.enst.fr/~cappe/ctrj_mix/">http://www.tsi.enst.fr/~cappe/ctrj_mix/</a>
Ntzoufras (2002)	Variable/model selection using BuGS
Katsis and Ntzoufras (2003)	Available from: <a href="http://www.jstatsoft.org/index.php?vol=7">http://www.jstatsoft.org/index.php?vol=7</a> <a href="http://www.ba.aegean.gr/ntzoufras/papers/paper13.htm">http://www.ba.aegean.gr/ntzoufras/papers/paper13.htm</a>
Petris and Tardella (2003)	<i>HI</i> : Simulation from distributions supported by complex hyperplanes Available from: <a href="http://cran.r-project.org/src/contrib/Descriptions/HI.htm">http://cran.r-project.org/src/contrib/Descriptions/HI.htm</a>
Sisson and Fan (2004a)	<i>RjDiag</i> : Convergence diagnostic for trans-dimensional point processes (see <i>Nmix</i> above) Available from: <a href="http://www.maths.unsw.edu.au/~scott/">http://www.maths.unsw.edu.au/~scott/</a>
Hoeting et al. (1999)	Links to various Bayesian Model Averaging software Available from: <a href="http://www.research.att.com/~volinsky/bma.html">http://www.research.att.com/~volinsky/bma.html</a>
Spiegelhalter et al. (1996b)	<i>WinBuGS</i> : Software package for the Bayesian analysis of complex models using MCMC
Gilks et al. (1992)	Available from: <a href="http://www.mrc-bsu.cam.ac.uk/bugs/welcome.shtml">http://www.mrc-bsu.cam.ac.uk/bugs/welcome.shtml</a> and the ‘Jump’ component to become available from: <a href="http://homepages.tesco.net/~creeping_death/">http://homepages.tesco.net/~creeping_death/</a>

Table 1: Freely available software for the implementation of trans-dimensional samplers, with URL and primary citation.

Jump Diffusion	
Model Comparison	Phillips and Smith (1996)
Image Segmentation	Han et al. (2004)
Target Recognition	Miller et al. (1995)
Reversible Jump	
Change-point models	Green (2001), Fan and Brooks (2000)
Signal processing	Larocque and Reilly (2002), Andrieu et al. (2001)
Mixture models	Richardson and Green (1997)
QTL mapping	Waagepetersen and Sorensen (2001), Stephens and Fisch (1998)
DNA segmentation	Boys and Henderson (2001)
Product Space	
Switching models	Kim and Nelson (2001)
Variable/Model selection	Spiegelhalter et al. (1996a), Katsis and Ntzoufras (2003)
Birth and Death	
Mixture models	Stephens (2000), Marin et al. (2003), Hurn et al. (2003)
Hidden Markov models	Shi et al. (2002)
Switching models	Soegner (2000)

Table 2: Articles with detailed implementations of the respective samplers.

## References

- Al-Awadhi, F., M. A. Hurn, and C. Jennison (2004). Improving the acceptance rate of reversible jump MCMC proposals. *Statistics and Probability Letters* 69, 189 – 198.
- Andrieu, C., J. de Freitas, and A. Doucet (2000). Reversible jump MCMC simulated annealing for neural networks. In *Uncertainty in Artificial Intelligence*, pp. 11 – 18. Morgan Kaufmann.
- Andrieu, C., P. M. Djurić, and M. Doucet (2001). Model selection by MCMC computation. *Signal Processing* 81, 19 – 37.
- Andrieu, C., A. Doucet, and C. P. Robert (2004). Computational advances for and from Bayesian analysis. *Statistical Science* 19, 118 – 127.
- Andrieu, C. and E. Moulines (2002). On the ergodicity properties of some adaptive MCMC algorithms. Technical report, University of Bristol.
- Atchade, Y. F. and J. S. Rosenthal (2003). On adaptive Markov chain Monte Carlo algorithms. Technical report, University of Montreal.
- Barbieri, M. M. and J. O. Berger (2004). Optimal predictive model selection. *The Annals of Statistics* 32, 870 – 897.
- Bartolucci, F. and A. Mira (2003). Efficient estimate of Bayes Factors from reversible jump output. Technical Report 33, Dipartimento di Economia, Università dell’Insubria.
- Bartolucci, F. and L. Scaccia (2003). A new approach for estimating the Bayes Factor. Technical report, University of Perugia.

- Berger, J. and L. Pericchi (1996). The Intrinsic Bayes Factor for model selection and prediction. *Journal of the American Statistical Association* 91, 109 – 122.
- Berger, J. O. and L. R. Pericchi (2001). In P. Lahiri (Ed.), *Model Selection*, Volume 38 of *IMS Lecture Notes - Monograph Series*, Chapter Objective Bayesian methods for model selection: Introduction and comparison (with discussion), pp. 135 – 207.
- Berger, J. O. and L. R. Pericchi (2004). Training samples in objective Bayesian model selection. *The Annals of Statistics* 32, 841 – 869.
- Besag, J. (1994). Contribution to the discussion of a paper by Grenander and Miller. *Journal of the Royal Statistical Society, B* 56, 591 – 592.
- Bottolo, L., G. Consonni, P. Dellaportas, and A. Lijoi (2003). Bayesian analysis of extreme values by mixture modelling. *Extremes* 6, 25 – 47.
- Boys, R. J. and D. A. Henderson (2001). A comparison of reversible jump MCMC algorithms for DNA sequence segmentation using hidden Markov models. *Computer Science and Statistics* 33, 35 – 49.
- Brooks, S. P. (1997). Contribution to the discussion of a paper by Richardson and Green. *Journal of the Royal Statistical Society, Series B* 59, 774 – 775.
- Brooks, S. P., Y. Fan, and J. S. Rosenthal (2002). Perfect forward simulation via simulated tempering. Technical report, University of Cambridge.
- Brooks, S. P., N. Friel, and R. King (2003). Classical model selection via simulated annealing. *Journal of the Royal Statistical Society, B* 65, 503 – 520.
- Brooks, S. P. and A. Gelman (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics* 7, 434 – 455.
- Brooks, S. P. and P. Giudici (2000). MCMC convergence assessment via two-way ANOVA. *Journal of Computational and Graphical Statistics* 9, 266 – 285.
- Brooks, S. P., P. Giudici, and A. Philippe (2003). On non-parametric convergence assessment for MCMC model selection. *Journal of Computational and Graphical Statistics* 12, 1 – 22.
- Brooks, S. P., P. Giudici, and G. O. Roberts (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. *Journal of the Royal Statistical Society, B* 65, 3 – 39.
- Bühlmann, P. and A. J. Wyner (1999). Variable length Markov chains. *The Annals of Statistics* 27, 480 – 513.
- Cappé, O. and C. P. Robert (2000). Markov chain Monte Carlo: 10 years and still running! *Journal of the American Statistical Association* 95, 1282 – 1286.
- Cappé, O., C. P. Robert, and T. Rydén (2003). Reversible jump MCMC converging to birth-and-death MCMC and more general continuous time samplers. *Journal of the Royal Statistical Society, B* 65, 679 – 700.
- Carlin, B. P. and S. Chib (1995). Bayesian model choice via Markov chain Monte Carlo. *Journal of the Royal Statistical Society, B* 57, 473 – 484.
- Casella, G., M. Lavine, and C. Robert (2000). Explaining the perfect sampler. *Preprint Université Paris 9 - Dauphine..*

- Casella, G. and E. Moreno (2002). Objective Bayesian variable selection. *Submitted*.
- Castelloe, J. M. and D. L. Zimmerman (2002). Convergence assessment for reversible jump MCMC samplers. Technical Report 313, Department of Statistics and Actuarial Science, University of Iowa.
- Chen, M. H. and Q. M. Shao (1997). Estimating ratios of normalizing constants for densities with different dimensions. *Statistica Sinica* 7, 607 – 630.
- Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* 90, 1313 – 1321.
- Chib, S. and I. Jeliazkov (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association* 96, 270 – 281.
- Chipman, H., E. George, and R. E. McCulloch (2001). In P. Lahiri (Ed.), *Model Selection*, Number 38 in IMS Lecture Notes-Monograph Series, Chapter The practical implementation of Bayesian model selection (with discussion), pp. 67 – 134.
- Clyde, M. A. (1999a). Bayesian model averaging and model search strategies. In J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith (Eds.), *Bayesian Statistics 6*, pp. 157 – 185. Oxford University Press, Oxford.
- Clyde, M. A. (1999b). Discussion of “Bayesian model averaging: A tutorial” by Hoeting et al. (1999). *Statistical Science* 14, 401 – 404. (Corrected version available online at <http://www.stat.washington.edu/www/research/online/hoeting1999.pdf>).
- Clyde, M. A. and E. I. George (2004). Model uncertainty. *Statistical Science* 19, 81 – 94.
- Cowles, M. K. and B. P. Carlin (1996). Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association* 91, 883 – 904.
- Dellaportas, P., J. J. Forster, and I. Ntzoufras (2002). On Bayesian model and variable selection using MCMC. *Statistics and Computing* 12, 27 – 36.
- Denison, D. G. T., C. C. Holmes, B. K. Mallick, and A. F. M. Smith (2002). *Bayesian Methods for Nonlinear Classification and Regression*. Wiley, New York.
- Denison, D. G. T., B. K. Mallick, and A. F. M. Smith (1998). Automatic Bayesian curve fitting. *Journal of the Royal Statistical Society, Series B* 60, 333 – 350.
- Diaconis, P., S. Holmes, and R. M. Neal (2000). Analysis of a non-reversible Markov chain sampler. *The Annals of Applied Probability* 10, 726 – 752.
- Dimakos, X. K. (2001). A guide to exact simulation. *International Statistical Review* 69, 27 – 48.
- DiMatteo, I., C. R. Genovese, and R. E. Kass (2001). Bayesian curve-fitting with free-knot splines. *Biometrika* 88, 1055 – 1071.
- Ehlers, R. S. and S. P. Brooks (2003). Constructing general efficient proposals for reversible jump MCMC. Technical report, Department of Statistics, Federal University of Paraná.
- Fan, Y. and S. P. Brooks (2000). Bayesian modelling of prehistoric corbelled domes. *Journal of the Royal Statistical Society, Series D* 49, 339 – 354.
- Fill, J. A. (1998). An interruptible algorithm for exact sampling via Markov chains. *Annals of Applied Probability* 8, 131 – 162.

- Frigessi, A. (2003). In P. J. Green, N. L. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, Chapter On some current research in MCMC, pp. 172 – 175. Oxford University Press.
- Gelman, A., G. O. Roberts, and W. G. Gilks (1996). Efficient metropolis jumping rules. In A. P. D. J. M. Bernardo, J. O. Berger and A. F. M. Smith (Eds.), *Bayesian Statistics 5*, pp. 599 – 607. Valencia University Press, Valencia.
- Gelman, A. and D. B. Ruben (1992). Inference from iterative simulations using multiple sequences. *Statistical Science 7*, 457 – 511.
- George, A. W., K. L. Mengersen, and G. P. Davis (1999). A Bayesian approach to ordering gene markers. *Biometrics 55*, 419 – 429.
- George, E. I. and R. E. McCulloch (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association 88*, 881 – 889.
- George, E. I. and R. E. McCulloch (1996). *Markov chain Monte Carlo in Practice*, Chapter Stochastic search variable selection, pp. 203 – 214. Chapman and Hall, London.
- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In J. M. Bernardo, A. F. M. Smith, A. P. Dawid, and J. O. Berger (Eds.), *Bayesian Statistics 4*, pp. 169 – 193. New York: Oxford University Press.
- Geyer, C. J. and J. Møller (1994). Simulation procedures and likelihood inference for spatial point processes. *Scandinavian Journal of Statistics 21*, 359 – 373.
- Geyer, C. J. and E. A. Thompson (1995). Annealing Markov chain Monte Carlo with applications to ancestral inference. *Journal of the American Statistical Association 90*, 909 – 920.
- Ghosh, J. K. and T. Samanta (2001). Model selection – An overview. *Current Science 80*, 1135 – 1144.
- Gilks, W. R., S. Richardson, and D. J. (Eds.) (1996). *Markov chain Monte Carlo in Practice*. Chapman & Hall/CRC.
- Gilks, W. R., G. O. Roberts, and S. K. Sahu (1998). Adaptive Markov chain Monte Carlo through regeneration. *Journal of the American Statistical Association 93*, 1045 – 1054.
- Gilks, W. R., D. Thomas, and D. J. Spiegelhalter (1992). Software for the Gibbs sampler. *Computing Science and Statistics 24*, 439 – 448.
- Gåsemyr, J. (2003). On an adaptive Metropolis-Hastings algorithm with independent proposal. *Scandinavian Journal of Statistics 30*, 159 – 173.
- Godsill, S. (2003). In P. J. Green, N. L. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, Chapter Discussion of Trans-dimensional Markov chain Monte Carlo by P. J. Green, pp. 199 – 203. Oxford University Press.
- Godsill, S. J. (2001). On the relationship between Markov chain Monte Carlo methods for model uncertainty. *Journal of Computational and Graphical Statistics 10*, 1 – 19.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika 82*, 711 – 732.

- Green, P. J. (2001). In O. E. Barndorff-Nielsen, D. R. Cox, and C. Klüppelberg (Eds.), *Complex Stochastic Systems*, Number 87 in Monographs on Statistics and Probability, Chapter A primer on Markov chain Monte Carlo, pp. 1 – 62. Chapman and Hall/CRC.
- Green, P. J. (2003a). In P. J. Green, N. L. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, Chapter Trans-dimensional Markov chain Monte Carlo, pp. 179 – 198. Oxford University Press.
- Green, P. J. (2003b). Contribution to the discussion of a paper by Brooks, Giudici and Roberts. *Journal of the Royal Statistical Society, B* 65, 48 – 49.
- Green, P. J. and A. Mira (2001). Delayed rejection in reversible jump Metropolis-Hastings. *Biometrika* 88, 1035 – 1053.
- Green, P. J. and A. O’Hagan (1998). Model choice with MCMC on product spaces without using pseudo-priors. Technical Report 98-01, Department of Statistics, Nottingham University.
- Grenander, U. and M. I. Miller (1994). Representations of knowledge in complex systems. *Journal of the Royal Statistical Society, B* 56, 549 – 603.
- Haario, H., E. Saksman, and J. Tamminen (2001). An adaptive Metropolis algorithm. *Bernoulli* 7, 223 – 242.
- Han, C. and B. P. Carlin (2001). MCMC methods for computing Bayes Factors: A comparative review. *Journal of the American Statistical Association* 96, 1122 – 1132.
- Han, F., Z. Tu, and S.-C. Zhu (2004). Range image segmentation by an effective jump-diffusion method. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 1138 – 1153.
- Heikkinen, J. (2003). In P. J. Green, N. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, Chapter Trans-dimensional Bayesian non-parametrics with spatial point processes, pp. 203 – 206. Oxford University Press.
- Hoeting, J. A., D. Madigan, A. E. Raftery, and C. T. Volinsky (1999). Bayesian model averaging: A tutorial (with discussion). *Statistical Science* 14, 382 – 417. (Corrected version available online at <http://www.stat.washington.edu/www/research/online/hoeting1999.pdf>).
- Hurn, M., A. Justel, and C. P. Robert (2003). Estimating mixtures of regressions. *Journal of Computational and Graphical Statistics* 12, 55 – 79.
- Jackson, C. H. and L. D. Sharples (2004). Models for longitudinal data with censored changepoints. *Journal of the Royal Statistical Society Series C* 53, 149 – 162.
- Kass, R. E. and A. E. Raftery (1995). Bayes Factors. *Journal of the American Statistical Association* 90, 773 – 796.
- Kasuya, M. and I. Takagawa (2001). Model uncertainty of real exchange rate forecast over mid-term horizons. Technical Report 01-23, Research and Statistics Department, Bank of Japan.
- Katsis, A. and I. Ntzoufras (2003). Bayesian hypothesis testing for the distribution of insurance claim counts using the Gibbs sampler. Technical report, Department of Statistics and Actuarial Science, University of the Aegean.

- Kim, C.-J. and C. R. Nelson (2001). A Bayesian approach to testing for Markov-switching in univariate and dynamic factor models. *International Economic Review* 42, 989 – 1013.
- King, R. and S. P. Brooks (2004). A classical study of catch-effort models for Hector’s dolphins. *Journal of the American Statistical Association*. 99, 325 – 333.
- Larocque, J. R. and J. P. Reilly (2002). Reversible jump MCMC for joint detection and estimation of directions of arrival in coloured noise. *IEEE Transactions on Signal Processing* 50, 231 – 240.
- Lawler, G. and A. Sokal (1988). Bounds on the  $L^2$  spectrum for Markov chains and Markov processes. *Transactions of the American Mathematical Society* 309, 557 – 580.
- Lunn, D. J., N. Best, and J. Whittaker (2004). “Reversible Jump” and its implementation in WinBUGS. Technical report, Department of Epidemiology and Public Health, Imperial College School of Medicine, London.
- Madigan, D. and A. E. Raftery (1994). Model selection and accounting for uncertainty in graphical models using Occam’s window. *Journal of the American Statistical Association* 89, 1535 – 1546.
- Marin, J. M., M. R. R. Bernal, and M. P. Wiper (2003). Using weibull mixture distributions to model heterogeneous survival data. Technical Report 03-32, Department of Statistics and Economics, Carlos III University of Madrid.
- Müller, P. and F. A. Quintana (2004). Nonparametric Bayesian data analysis. *Statistical Science* 19, 95 – 110.
- Meng, X. L. and S. Schilling (2002). Wrap bridge sampling. *Journal of Computational and Graphical Statistics* 11, 552 – 586.
- Meng, X. L. and W. H. Wong (1996). Simulating ratios of normalising constants via a simple identity: A theoretical exploration. *Statistica Sinica* 6, 831 – 860.
- Mengersen, K. L. and C. P. Robert (2003). IID sampling using self-avoiding population Monte Carlo: The pinball sampler. In J. M. Bernardo, M. J. Baryarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West (Eds.), *Bayesian Statistics* 7, pp. 277 – 292. Oxford University Press.
- Mengersen, K. L., C. P. Robert, and C. Guihenneuc-Joyaux (1999). MCMC convergence diagnostics: A review. In J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith (Eds.), *Bayesian Statistics* 6, pp. 415 – 440. Oxford University Press, Oxford.
- Miller, M. I., A. Srivastava, and U. Grenander (1995). Conditional-mean estimation via jump-diffusion processes in multiple target tracking/recognition. *IEEE Transactions on Signal Processing* 43, 2678 – 2690.
- Mira, A. and C. J. Geyer (2000). *Fields Institute Communications: Monte Carlo Methods*, Chapter On non-reversible Markov chains, pp. 93 – 108.
- Mira, A. and G. Nicholls (2004). Bridge estimation of the probability density at a point. *Statistica Sinica* 14, 603 – 612.
- Møller, J. and G. K. Nicholls (1999). Perfect simulation for sample-based inference. Technical report, Aalborg University.

- Neal, R. M. (2004). Improving asymptotic variance of MCMC estimators: Non-reversible chains are better. Technical Report 0406, Department of Statistics, University of Toronto.
- Ntzoufras, I. (2002). Gibbs variable selection using BuGS. *Journal of Statistical Software* 7(7).
- Ntzoufras, I., P. Dellaportas, and J. J. Forster (2003). Bayesian variable and link determination for generalised linear models. *Journal of Statistical Planning and Inference* 111, 165 – 180.
- O’Hagan, A. (1995). Fractional Bayes Factors for model comparisons. *Journal of the Royal Statistical Society, Series B* 57, 99 – 138.
- Pérez, J. M. and J. O. Berger (2001). In E. George and P. Nanopolus (Eds.), *Bayesian Methods with Applications to Science, Policy and Official Statistics*, Chapter Analysis of mixture models using expected posterior priors, with application to classification of gamma ray bursts, pp. 401 – 410. Official Publications of the European Communities, Luxembourg.
- Pérez, J. M. and J. O. Berger (2002). Expected posterior prior distributions for model selection. *Biometrika* 89, 491 – 511.
- Petris, G. and L. Tardella (2003). A geometric approach to transdimensional Markov chain Monte Carlo. *The Canadian Journal of Statistics* 31.
- Phillips, D. B. and A. F. M. Smith (1996). *Markov chain Monte Carlo in Practice*, Chapter Bayesian model comparison via jump diffusions, pp. 215 – 239. Chapman and Hall, London.
- Preston, C. J. (1977). Spatial birth-and-death processes. *Bulletin of the International Statistical Institute* 46, 371 – 391.
- Propp, J. G. and D. B. Wilson (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random structures and Algorithms* 9, 223 – 252.
- Richardson, S. and P. J. Green (1997). On Bayesian analysis of mixtures with an unknown number of components (with discussion). *Journal of the Royal Statistical Society, B* 59, 731 – 792.
- Ripley, B. D. (1977). Modelling spatial patterns (with discussion). *Journal of the Royal Statistical Society, B* 39, 172 – 212.
- Robert, C. (2001). *The Bayesian Choice* (2nd ed.). Springer, New York.
- Robert, C. and G. Casella (2004). *Monte Carlo Statistical Methods* (2nd ed.). Springer, New York. To appear.
- Robert, C. P. (2003). In P. J. Green, N. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, Chapter Advances in MCMC: A discussion, pp. 167 – 171. Oxford University Press.
- Roberts, G. O. (2003). In P. J. Green, N. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, Chapter Linking theory and practice of MCMC, pp. 145 – 166. Oxford University Press.

- Roberts, G. O., A. Gelman, and W. R. Gilks (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *Annals of Applied Probability* 7, 110 – 120.
- Roberts, G. O. and R. L. Tweedie (1996). Exponential convergence of Langevin diffusions and their discrete approximations. *Bernoulli* 2, 341 – 364.
- Royston, P. and D. G. Altman (1994). Regression using fractional polynomials of continuous covariates: Parsimonious parametric modelling. *Applied Statistics* 43, 429 – 453.
- Sahu, S. K. and A. A. Zhigljavsky (2003). Self regenerative Markov chain Monte Carlo with adaption. *Bernoulli* 9, 395 – 422.
- Salmenkivi, M., J. Kere, and H. Mannila (2002). Genome segmentation using piecewise constant intensity models and reversible jump MCMC. *Bioinformatics* 18, S211 – S218.
- Shi, J. Q., R. Murray-Smith, and D. M. Titterton (2002). Birth-death MCMC methods for mixtures with an unknown number of components. Technical Report TR-2002-117, Department of Computing Science, University of Glasgow.
- Sisson, S. A. and Y. Fan (2004a). A distance-based diagnostic for trans-dimensional Markov chains. Technical report, School of Mathematics, University of New South Wales.
- Sisson, S. A. and Y. Fan (2004b). Towards automating model selection for a mark-recapture-recovery analysis. Technical report, School of Mathematics, University of New South Wales.
- Sisson, S. A. and M. A. Hurn (2004). Bayesian point estimation of quantitative trait loci. *Biometrics* 60, 60 – 68.
- Smith, B. J. (2001). *Bayesian Output Analysis Program (BOA), User's Manual*. Department of Biostatistics, University of Iowa.
- Smith, M. and R. Kohn (1996). Nonparametric regression using Bayesian variable selection. *Journal of Econometrics* 75, 317 – 344.
- Soegner, L. (2000). Selection of the number of states by birth-death processes. Technical report, Department of Economics, Vienna University of Economics and Business Administration.
- Spiegelhalter, D., A. Thomas, N. Best, and W. Gilks (1996a). Bugs 0.5: Examples Volume 2, MRC Biostatistics Unit, Institute of Public Health, Cambridge, U.K. Available on line from <http://www.mrc-bsu.cam.ac.uk/bugs/documentation/exampvol2/node20.html>.
- Spiegelhalter, D. J., A. Thomas, N. G. Best, and W. R. Gilks (1996b). BUGS: Bayesian inference Using Gibbs Sampling version 0.50. *Cambridge: Medical Research Council Biostatistical Unit*.
- Stephens, D. A. and R. D. Fisch (1998). Bayesian analysis of quantitative trait locus data using reversible jump Markov chain Monte Carlo. *Biometrics* 54, 1334 – 1347.
- Stephens, M. (2000). Bayesian analysis of mixture models with an unknown number of components - an alternative to reversible jump methods. *Annals of Statistics* 28, 40 – 74.

- Suchard, M. A., R. E. Weiss, and J. S. Sinsheimer (2001). Bayesian selection of continuous-time Markov chain evolutionary models. *Molecular Biology and Evolution* 18, 1001 – 1013.
- Tierney, L. (1996). *Markov chain Monte Carlo in Practice*, Chapter Introduction to general state-space Markov chain theory, pp. 59 – 74. Chapman and Hall, London.
- Tierney, L. (1998). A note on Metropolis-Hastings kernels for general state spaces. *Annals of Applied Probability* 8, 1 – 9.
- Waagepetersen, R. and D. Sorensen (2001). A tutorial on reversible jump MCMC with a view toward applications on QTL mapping. *International Statistical Review* 69, 49 – 62.
- Zellner, A. and Siow (1980). Posterior odds for selected regression hypotheses. In D. V. L. J. M. Bernardo, M. H. DeGroot and A. F. M. Smith (Eds.), *Bayesian Statistics*, pp. 585 – 603. Valencia University Press, Valencia.