Efficient sampling of Gaussian graphical models using conditional Bayes factors

Max Hinne[*1], Alex Lenkoski[2], Tom Heskes[1] and Marcel van Gerven[1]

[1]Radboud University Nijmegen, the Netherlands
[2]Norwegian Computing Center, Norway

September 10, 2014

ABSTRACT. Bayesian estimation of Gaussian graphical models has proven to be challenging because the conjugate prior distribution on the Gaussian precision matrix, the G-Wishart distribution, has a doubly intractable partition function. Recent developments provide a direct way to sample from the G-Wishart distribution, which allows for more efficient algorithms for model selection than previously possible. Still, estimating Gaussian graphical models with more than a handful of variables remains a nearly infeasible task. Here, we propose two novel algorithms that use the direct sampler to more efficiently approximate the posterior distribution of the Gaussian graphical model. The first algorithm uses conditional Bayes factors to compare models in a Metropolis-Hastings framework. The second algorithm is based on a continuous time Markov process. We show that both algorithms are substantially faster than state-of-the-art alternatives. Finally, we show how the algorithms may be used to simultaneously estimate both structural and functional connectivity between subcortical brain regions using resting-state fMRI.

1. INTRODUCTION

A key objective in many areas of science is to uncover the interactions amongst a large number of variables based on a limited amount of data. Examples include gene regulatory networks, where one wants to identify the interactions amongst DNA segments, market basket analysis where the relations are studied between customers based on their purchase behavior, or neuroscience where the connections between segregated neuronal populations are linked to cognitive ability and impairment. One way to estimate these relations is to employ Gaussian graphical models, where the non-zero entries in the off-diagonal of a precision matrix correspond to the edges in a conditional independence graph [1]. However, fully Bayesian estimation of the posterior of a Gaussian graphical model has proven to be notoriously hard.

To allow Bayesian inference of the Gaussian graphical model, a conjugate prior [2] on a precision matrix restricted by the conditional independence graph G was constructed for decomposable graphs [3], and later generalized to arbitrary graphs [4]. Subsequent work coined this distribution the G-Wishart distribution [5]. A number of Monte Carlo algorithms for model estimation using the G-Wishart distribution have been developed [6–9], but each of these algorithms required substantial computational resources due to difficulty with sampling from the G-Wishart distribution. To address this bottleneck, a recent study proposed an efficient way to directly sample from the G-Wishart distribution [10] by scaling samples from a regular Wishart distribution to fit the required dependency structure [11]. Even with the direct sampler, approximating the Gaussian graphical model remained difficult because of the doubly intractable partition function of the G-Wishart distribution. However, by combining features of the exchange algorithm [12] with reversible jump sampling [13], calculating the partition function may be circumvented [10]. The algorithm that implements this idea, named the double reversible jump algorithm, provides substantial computational gains compared to earlier approaches [10].

Although the double reversible jump algorithm enables model selection in a more efficient manner than previously possible, computational costs remain a limiting factor in practical applications with a large number of variables. In this paper, we propose two novel, faster, algorithms for Bayesian estimation of the Gaussian graphical model. In the first algorithm, we combine the direct sampler [10] with an efficient representation of the conditional Bayes factor [14], which results in an elegant Metropolis-Hastings algorithm to which we will refer as the double conditional Bayes

---

[*]mhinne@cs.ru.nl

factor sampler. In the second algorithm, we cast the double conditional Bayes factors algorithm in a birth-death MCMC setting [15]. Here, rather than accepting or rejecting a new state with an edge added or removed, we associate with these changes birth and death events, respectively. These events occur with such rates that their equilibrium coincides with the posterior of interest [16]. Both algorithms provide substantial speed improvement over the status quo, as we show in simulations.

We also provide an application of our algorithms by estimating structural and functional connectivity between subcortical structures using resting-state fMRI. It is a major goal in cognitive neuroscience to understand how spatially segregated neural populations are coupled, using indirect measures of neural activity such as functional magnetic resonance imaging [17, 18]. In this context, the anatomical pathways between neural populations are referred to as structural connectivity whereas correlated activity patterns between these populations are referred to as functional connectivity [19]. Both forms of connectivity may be estimated simultaneously using Gaussian graphical models. Here, the precision matrix captures the functional interactions between variables and the associated conditional independence graph represents the direct connections between variables. Bayesian estimation of Gaussian graphical models is particularly relevant since the posterior over precision matrices provides complete information about the strength of functional interactions and the posterior over conditional independence graphs allows one to associate a probability with a putative direct connection between variables of interest.

2. Gaussian graphical models

2.1. **Preliminaries.** Let observed data $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)^T$ consist of n independent draws from a p-dimensional multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{K}^{-1})$, with zero mean and precision (inverse covariance) matrix $\mathbf{K}$. Here, $\mathbf{K} \in \mathbb{P}_p$, with $\mathbb{P}_p$ the space of positive definite $p \times p$ matrices. The likelihood of $\mathbf{K}$ is given by

$$(1) \qquad P(\mathbf{X} \mid \mathbf{K}) = \prod_{i=1}^{n} \mathcal{N}(\mathbf{x}_i \mid \mathbf{0}, \mathbf{K}^{-1}) \propto |\mathbf{K}|^{n/2} \exp\left[-\frac{1}{2}\langle \mathbf{K}, \mathbf{S} \rangle\right] \ ,$$

where $\mathbf{S} = \mathbf{X}^T\mathbf{X}$ is the empirical covariance and $\langle \cdot, \cdot \rangle$ the trace inner product operator. The precision matrix has the important property that zero elements correspond to conditional independencies. In other words, (1) specifies a Gaussian Markov random field with respect to a graph $G = (V, E)$, with $V = \{1, \ldots, p\}$ and $E \subset V \times V$, in which the absence of a connection indicates independence, i.e. $(i, j) \notin E \to k_{ij} = 0$. For convenience, throughout this paper we slightly abuse notation and use $(i, j) \in G$ to indicate that the edge (i, j) is present in E.

The dependency graph may be used to specify a prior distribution on the precision matrix, which is known as the G-Wishart distribution [4]:

$$(2) \qquad P(\mathbf{K} \mid G, \delta, \mathbf{D}) = \mathcal{W}_G(\delta, \mathbf{D}) = \frac{|\mathbf{K}|^{(\delta-2)/2}}{Z_G(\delta, \mathbf{D})} \exp\left[-\frac{1}{2}\langle \mathbf{K}, \mathbf{D} \rangle\right] \mathbf{1}_{\mathbf{K} \in \mathbb{P}_G} \ ,$$

in which $\mathbb{P}_G$ is the space of positive definite $p \times p$ matrices that have zero elements wherever $(i, j) \notin G$, δ is the prior degrees of freedom, $\mathbf{D}$ is the prior scaling matrix and $\mathbf{1}_x$ evaluates to 1 if and only if x holds and to 0 otherwise. The G-Wishart distribution is conjugate to the multivariate Gaussian likelihood in (1), so that

$$(3) \qquad P(\mathbf{K} \mid G, \delta, \mathbf{D}, \mathbf{X}) = \mathcal{W}_G(\delta + n, \mathbf{D} + \mathbf{S}) = \frac{|\mathbf{K}|^{(n+\delta-2)/2}}{Z_G(\delta + n, \mathbf{D} + \mathbf{S})} \exp\left[-\frac{1}{2}\langle \mathbf{K}, \mathbf{D} + \mathbf{S} \rangle\right] \ .$$

Note that the Wishart distribution is a special case of the G-Wishart distribution, with which it coincides if G is a fully connected graph. Importantly, the partition function $Z_G(\delta, \mathbf{D})$ depends on G, which makes the G-Wishart a doubly intractable distribution. We return to the implications of this fact later on.

Central to this work is that we wish to perform model selection in Gaussian graphical models, which revolves around the joint posterior

$$(4) \qquad P(G, \mathbf{K} \mid \mathbf{X}) \propto P(\mathbf{X} \mid \mathbf{K})P(\mathbf{K} \mid G)P(G) \ .$$

In the remainder, we outline several algorithms to approximate this distribution.

2.2. **Direct samples from the G-Wishart distribution.** Since the prior $P(\mathbf{K} \mid G)$ is $\mathcal{W}_G(\delta, \mathbf{D})$, we need a way to draw samples from the G-Wishart distribution. Up until recently, this was achieved using a block Gibbs sampler that updates $\mathbf{K}$ according to either the edges of G [9] or its clique decomposition [6]. Although this enables model inference of $P(G, \mathbf{K} \mid \mathbf{X})$, as desired, both approaches require substantial computational effort, making them prohibitive for use in contexts with a large number of variables. An alternative method was proposed that is more efficient [10], which is an adaptation of an algorithm for estimating the mode $\hat{\mathbf{K}}$ of the G-Wishart distribution [11, 20]. The algorithm is as follows:

(1) Sample $\mathbf{\Sigma} \sim \mathcal{W}(\delta, \mathbf{D})$.
(2) Let $\mathbf{W} = \mathbf{\Sigma}$.
(3) Repeat for $j = 1, 2, \ldots, p$ until convergence:
  (a) Let $N_j \subset V$ be the set of variables that are connected to j in G.
      Form $\mathbf{W}_{N_j}$ and $\mathbf{\Sigma}_{N_j, j}$ and solve
      $$\hat{\beta}_j^* = \mathbf{W}_{N_j}^{-1} \mathbf{\Sigma}_{N_j, j}$$
  (b) Form $\hat{\beta}_j \in \mathbb{R}^{p-1}$ by copying the elements of $\hat{\beta}_j^*$ to the appropriate
      locations and imputing zeros in those locations not connected to j in G.
  (c) Replace $\mathbf{W}_{j, -j}$ and $\mathbf{W}_{-j, j}$ with $\mathbf{W}_{-j, -j} \hat{\beta}_j$.
(4) Return $\mathbf{K} = \mathbf{W}^{-1}$.

Conceptually, the algorithm draws a sample from a Wishart distribution, which is then iteratively scaled according to the dependence structure in G. In practice, we observe that convergence (see step 3) is typically reached within a handful of iterations, even for moderate to large p.

3. Sampling algorithms

The direct sampler paves the way for novel inference algorithms. Here, we introduce two novel algorithms for approximation of the joint posterior in (4).

3.1. **Double reversible jump sampler.** As a baseline for comparison, we use the double reversible jump (DRJ) sampler [10]. This algorithm was shown to be more efficient as previously used approaches and may be considered state of the art. It builds upon the reversible jump sampler discussed in [21]. The key idea offered by this approach is that it introduces an auxiliary variable $\mathbf{K}^0 \sim \mathcal{W}_G(\delta, \mathbf{D})$, as in the exchange algorithm [12], that is efficiently sampled using the direct sampler discussed above. Because of the way this auxiliary variable is constructed, the doubly intractable partition functions of the G-Wishart distribution are canceled out in the calculation of the acceptance ratios of newly proposed graphs.

3.2. **Direct double conditional Bayes factor sampler.** The double reversible jump algorithm provides a substantial improvement over previous algorithms, as it avoids the need to approximate the ratio of partition functions or invoke the Gibbs sampling algorithm for drawing samples from the G-Wishart distribution. Nonetheless, the algorithm can be simplified. In [14] it is shown that if G and $\tilde{G}$ differ only in the edge $e = (p-1, p)$ and $G \subset \tilde{G}$, the odds ratio of these two models may be expressed as

$$(5) \qquad \frac{P(\mathbf{X} \mid \tilde{G}, \mathbf{K}, \mathbf{D})}{P(\mathbf{X} \mid G, \mathbf{K}, \mathbf{D})} = N(\mathbf{K}, \mathbf{D} + \mathbf{S}) \frac{Z_G(\delta, \mathbf{D})}{Z_{\tilde{G}}(\delta, \mathbf{D})}$$

with

$$(6) \qquad N(\mathbf{K}, \mathbf{U}) \equiv \phi_{p-1,p-1} \left(\frac{2\pi}{u_{pp}} \right)^{1/2} \exp \left[\frac{1}{2} u_{pp} \left(\frac{\phi_{p-1,p-1} u_{p-1,p}}{u_{pp}} - \frac{\sum_{l=1}^{p-2} \phi_{lp-1} \phi_{lp}}{\phi_{p-1,p-1}} \right)^2 \right],$$

where $\mathbf{K} = \mathbf{\Phi}^T \mathbf{\Phi}$, with $\mathbf{\Phi}$ the Cholesky decomposition of $\mathbf{K}$. The term in (5) can be considered the conditional Bayes factor of the comparison between G and $\tilde{G}$. Similar to the double reversible jump

approach, [14] propose to augment the sampling process with an auxiliary variable $\mathbf{K}^0 \sim \mathcal{W}_G(\delta, \mathbf{D})$. This results in a convenient acceptance ratio for the addition of an edge to G

$$(7) \qquad \alpha = \frac{N(\mathbf{K}, \mathbf{D} + \mathbf{S})}{N(\tilde{\mathbf{K}}^0, \mathbf{D})} \frac{P(\tilde{G})}{P(G)} \ ,$$

where the ratio is inverted if the edge is removed from G instead. Note that the variables $G, \mathbf{K}, \mathbf{U}$ and $\mathbf{D}$ must be permuted for each edge flip to place the particular edge under consideration in the position $(p-1, p)$.

The algorithm described in [14] employs the block Gibbs sampler to sample from the G-Wishart distribution. Instead, here we propose to make use of the direct sampler explained in Section 2.2 to arrive at the following procedure for estimation of the Gaussian graphical model:

(1) Let $G = G^{[s]}$ be the current graph and let $\mathbf{K} = \mathbf{K}^{[s]} \sim \mathcal{W}_G(\delta + n, \mathbf{D} + \mathbf{S})$.
(2) For each edge $(i, j) \in G$, do:
  (a) Create a permutation of the variables so that $(i, j) \to (p-1, p)$. Permute $G, \mathbf{K}, \mathbf{D}$ and $\mathbf{S}$ accordingly.
  (b) Let $\tilde{G} = G \cup (p-1, p)$ if $(p-1, p) \notin G$ or $\tilde{G} = G \setminus (p-1, p)$ if $(p-1, p) \in G$.
  (c) Draw $\tilde{\mathbf{K}}^0 \sim \mathcal{W}_{\tilde{G}}(\delta, \mathbf{D})$.
  (d) Accept the move from G to $\tilde{G}$ with probability α as in (7).
  (e) Restore the original ordering of $G, \mathbf{K}, \mathbf{D}$ and $\mathbf{S}$ and draw $\tilde{\mathbf{K}} \sim \mathcal{W}_{\tilde{G}}(\delta + n, \mathbf{D} + \mathbf{S})$
(3) Set $G^{[s+1]} = \tilde{G}$ and $\mathbf{K}^{[s+1]} \sim \mathcal{W}_{\tilde{G}}(\delta + n, \mathbf{D} + \mathbf{S})$.

The usage of the direct sampler instead of the block Gibbs updates makes this direct double conditional Bayes factors (DCBF) algorithm computationally much more efficient [22].

3.3. **Double continuous time sampler.** A downside of the usage of an auxiliary variable scheme is that it decreases the acceptance probability of proposals, as essentially two moves have to be accepted at once. This hampers mixing of the Markov chain, so that multimodal distributions are approximated poorly. To improve acceptance, [15] introduce a birth-death continuous-time Markov process [23] for Gaussian graphical models. Rather than accepting the addition or removal of an edge, [15] associates birth and death events with these changes, respectively. Each edge dies independently of all others as a Poisson process with death rate $d_e(G, \mathbf{K})$. Because the edges are independent, the overall death rate at a particular pair of graph G and precision $\mathbf{K}$ is $d(\mathbf{K}) = \sum_e d_e(G, \mathbf{K})$. Birth rates $b(\mathbf{K})$ are defined similarly, but for non-edges instead.

Because the birth and death processes are independent Poisson processes, the expected time between two events is $1/(d(\mathbf{K}) + b(\mathbf{K}))$. This time can be considered the process spends at any particular instance of $(G, \mathbf{K})$. The probability of the death event of edge $e \in G$ is

$$(8) \qquad P(\text{death of edge } e) = \frac{d(G, \mathbf{K})}{b(G, \mathbf{K}) + d(G, \mathbf{K})} \ ,$$

with again an analogous definition for the birth event for a non-edge.

Mohammadi and Wit [15] show that the birth-death process has the posterior $P(G, \mathbf{K} \mid \mathbf{X})$ as stationary distribution, if for all edges and non-edges e

$$(9) \qquad d_e(\tilde{G}, \tilde{\mathbf{K}}) P(\tilde{G}, \tilde{\mathbf{K}} \mid \mathbf{X}) = b_e(G, \mathbf{K}) P(G, \mathbf{K} \mid \mathbf{X}) \ ,$$

for $\tilde{G} = G \cup e$. The birth and death rates may be chosen accordingly as

$$(10) \qquad b_e(G, \mathbf{K}) = \frac{P(\tilde{G}, \tilde{\mathbf{K}} \mid \mathbf{X})}{P(G, \mathbf{K} \mid \mathbf{X})} \quad \text{for } e \notin G \quad \text{and} \quad d_e(G, \mathbf{K}) = \frac{P(G, \mathbf{K} \mid \mathbf{X})}{P(\tilde{G}, \tilde{\mathbf{K}} \mid \mathbf{X})} \quad \text{for } e \in G.$$

with again $\tilde{G} = G \cup e$.

The key observation is now that these birth-death rates can be computed using the double conditional Bayes factors as in (7). Here again we make use of the exchange framework by introducing the auxiliary variable $\mathbf{K}^0$, such that explicit evaluation of the partition functions is circumvented. This leads to a novel approach that we will refer to as the double continuous time (DCT) sampler, given by:

(1) Let $G = G^{[s]}$ be the current graph and let $\mathbf{K} = \mathbf{K}^{[s]} \sim \mathcal{W}_G(\delta + n, \mathbf{D} + \mathbf{S})$.

(2) For each non-edge $e \notin G$:
  (a) Create a random permutation of the variables so that $(i, j) \rightarrow (p - 1, p)$. Permute $G, \mathbf{K}, \mathbf{D}$ and $\mathbf{S}$ accordingly.
  (b) Let $\tilde{G} = G \cup e$. Draw $\mathbf{K}^0 \sim \mathcal{W}_{\tilde{G}}(\delta, \mathbf{D})$
  (c) Compute the birth rate $b_e(G, \mathbf{K})$ using (10).
(3) Compute the total birth rate of the current state $b(G, \mathbf{K})$.
(4) For each edge $e \in G$:
  (a) Create a random permutation of the variables so that $(i, j) \rightarrow (p - 1, p)$. Permute $G, \mathbf{K}, \mathbf{D}$ and $\mathbf{S}$ accordingly.
  (b) Let $\tilde{G} = G \setminus e$. Draw $\mathbf{K}^0 \sim \mathcal{W}_{\tilde{G}}(\delta, \mathbf{D})$
  (c) Compute the death rate $d_e(G, \mathbf{K})$ using (10).
(5) Compute the total death rate of the current state $d(G, \mathbf{K})$ and the waiting time between events $w(G, \mathbf{K}) = 1/(d(\mathbf{K}) + b(\mathbf{K}))$.
(6) Create a birth or death event according to the probabilities of death events (8) and birth events, and set $G^{[s+1]} = \tilde{G}$ and $\mathbf{K}^{[s+1]} \sim \mathcal{W}_{\tilde{G}}(\delta + n, \mathbf{D} + \mathbf{S})$.

4. Experiments

In this section we first analyze the validity of the two proposed methods using an example with a known precision matrix. Subsequently we apply the algorithms in an explorative study to identify structural and functional connectivity between subcortical brain structures.

4.1. **Simulation.** We compared the performance of the double reversible jump algorithm and the two novel algorithms using a simulation proposed in [9]. In this example, we have $p = 6$ and $n = 18$. Furthermore, the precision matrix $\mathbf{K}$ is given by $k_{ii} = 1$ for $i = 1, \ldots, p$, $k_{i,i+1} = k_{i+1,i} = 0.5$ for $i = 1, \ldots, p - 1$ and finally $k_{1p} = k_{p1} = 0.4$. The associated conditional independence graph G follows as $(i, j) \in G \leftrightarrow k_{ij} \neq 0$. The scatter matrix is then constructed as $\mathbf{S} = \mathbf{X}\mathbf{X}^T = n\mathbf{K}^{-1}$, which corresponds to n independent observations of $\mathcal{N}(\mathbf{0}, \mathbf{K}^{-1})$. Through exhaustive enumeration of all 32,768 possible graphs of size p, [9] shows that the posterior edge probabilities are

$$(11) \qquad P((i,j) \in G \mid \mathbf{X}) = \begin{pmatrix} 1 & 0.969 & 0.106 & 0.085 & 0.113 & 0.850 \\ 0.969 & 1 & 0.980 & 0.098 & 0.081 & 0.115 \\ 0.106 & 0.980 & 1 & 0.982 & 0.098 & 0.086 \\ 0.085 & 0.098 & 0.982 & 1 & 0.980 & 0.106 \\ 0.113 & 0.081 & 0.98 & 0.980 & 1 & 0.970 \\ 0.850 & 0.115 & 0.086 & 0.106 & 0.970 & 1 \end{pmatrix}$$

and the expectation of $\mathbf{K}$ is

$$(12) \qquad \mathbb{E}(\mathbf{K} \mid \mathbf{X}) = \begin{pmatrix} 1.139 & 0.569 & -0.011 & 0.006 & -0.013 & 0.403 \\ 0.569 & 1.175 & 0.574 & -0.008 & 0.005 & -0.014 \\ -0.011 & 0.574 & 1.176 & 0.574 & -0.008 & 0.006 \\ 0.006 & -0.008 & 0.574 & 1.175 & 0.573 & -0.011 \\ -0.013 & 0.005 & -0.008 & 0.573 & 1.175 & 0.569 \\ 0.403 & -0.014 & 0.006 & -0.011 & 0.569 & 1.138 \end{pmatrix}.$$

We approximate this ground truth using the three different algorithms, each implemented in Matlab. Throughout, we use vague priors in the form of $P(G) \propto 1$ for G and $P(\mathbf{K} \mid G) = \mathcal{W}_G(3, \mathbf{I}_p)$. The algorithms are each executed for 100,000 iterations, of which the first 50,000 are discarded as burn-in. Conditional expectations for edges (i.e. edge probabilities) and precision matrices are then calculated as

$$(13) \qquad \mathbb{E}((i,j) \in G \mid \mathbf{X}) = \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}_{(i,j) \in G_t} \quad \text{and} \quad \mathbb{E}(\mathbf{K} \mid \mathbf{X}) = \frac{1}{T} \sum_{t=1}^{T} \mathbf{K}_t$$

for the double reversible jump and the double conditional Bayes factor algorithms, with T the number of samples. For the double continuous time algorithm, these expectations are calculated

TABLE 1. Results for the comparison between the three described samplers on a simulated example, averaged over 10 simulations. Standard errors are indicated in parentheses. Shown are the mean squared error (MSE) of edge probabilities relative to (11), the Kullback-Leibler divergence (KL) between the expected precision matrix and (12), the number of unique models visited, the marginal posterior probability of the true graph $P(G \mid \mathbf{S})$ and the relative speed of the algorithms compared to the double reversible jump baseline.

Algorithm	MSE	KL	#models	$P(G \mid \mathbf{S})$	Rel. speed
DRJ	5e-04 (4e-05)	1e-04 (2e-05)	1299 (31)	0.3674 (0.0008)	1 (0)
DCBF	5e-04 (2e-05)	1e-04 (1e-05)	1472 (23)	0.3826 (0.0040)	3.57 (1e-01)
DCT	1e-03 (1e-05)	7e-04 (3e-04)	1187 (35)	0.4277 (0.0008)	3.80 (1e-02)

as

$$(14) \qquad \mathbb{E}((i,j) \in G \mid \mathbf{X}) = \frac{1}{W} \sum_{t=1}^{T} w_t \mathbf{1}_{(i,j) \in G_t} \quad \text{and} \quad \mathbb{E}(\mathbf{K} \mid \mathbf{X}) = \frac{1}{W} \sum_{t=1}^{T} w_t \mathbf{K}_t \ ,$$

with $W = \sum_{t=1}^{T} w_t$. It is easy to see that this idea generalizes the discrete time MCMC approach by assuming $w_t = 1$ for all t.

We quantify the approximation accuracy of the three algorithms in a number of ways. First, the accuracy of the edge probabilities is expressed using the mean squared error with respect to the true probabilities in (11). Second, we compute the Kullback-Leibler divergence [24] between the precision matrix obtained in [9] as defined in (12) and $\hat{\mathbf{K}} \equiv \mathbb{E}(\mathbf{K} \mid \mathbf{X})$ using either of the algorithms. We also count the number of unique models that each algorithm considers to express mixing behavior. Next, we compute the marginal posterior probability of the true graph. Finally, we compute the relative computational speeds of the algorithms. The results of the comparison are shown in Table 1. The algorithms have similar performance in approximating the desired posterior distribution and each obtains the true graph as the mode of the approximated distribution. Contrary to [15], we do not find the continuous time algorithm to have the best mixing. In fact, of the three considered models, the continuous time MCMC approach finds the smallest number of unique models. Note that the continuous time approach may converge faster [25], but this is not apparent in this simulation. Finally, the efficiency of our way of computing the conditional Bayes factor (see (5)) is demonstrated by a substantial speed increase, as the DCBF algorithm is 3.57 times faster than the DRJ sampler, and the DCT algorithm is 3.80 times faster than the DRJ algorithm, whereas the algorithm in [15] is 1.79 times slower than the DRJ sampler.

4.2. **Subcortical brain connectivity.** As an explorative example, we estimate structural and functional connectivity in a fully Bayesian setting. In previous work, functional connectivity has been estimated under the assumption that the underlying structural connectivity was known [26]. Here, we address the more challenging problem of simultaneously estimating the posterior distribution of both structural and functional connectivity.

4.2.1. *Empirical data.* The data consist of resting-state functional MRI data collected for one subject. We refer the reader to [27] for details of the acquisition protocol. Preprocessing was performed using FSL 5.0 [28] and consisted of the following steps. T1 images were linearly registered to MNI-152 space. Multi-echo volumes at each TR were combined [29]. Motion correction was performed using MCFLIRT and estimated motion parameters were regressed out together with their temporal derivatives and mean time courses for both WM and CSF. Finally, data were high-pass filtered at 0.001 Hz. Subcortical structures were segmented using FSL FIRST [30], resulting in data for a total of 14 regions, consisting of bilateral accumbens, amygdala, caudate, hippocampus, pallidum, putamen and thalamus (see Fig. 1A). For each of these regions the signal was averaged over all voxels in that region and subsequently standardized to have zero mean and unit variance.
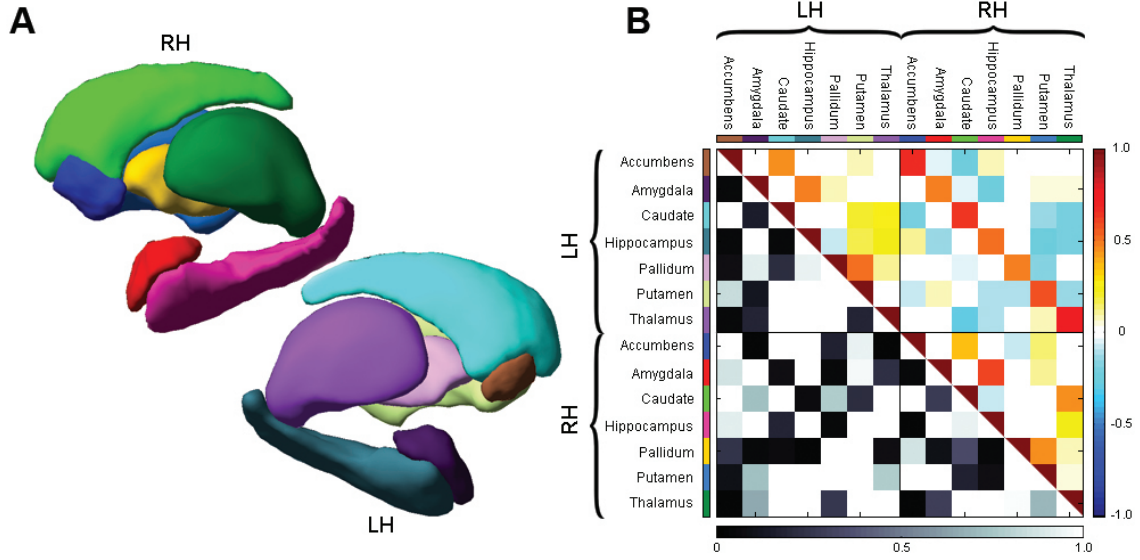
FIGURE 1. Subcortical connectivity. A. Subcortical structures, consisting of bilateral accumbens, amygdala, caudate, hippocampus, pallidum, putamen and thalamus. B. Posterior probabilities of structural connectivity (lower triangle) and expected partial correlations between these structures (upper triangle). LH and RH indicate left hemisphere, right hemisphere, respectively.

4.2.2. *Bayesian structural and functional connectivity estimation.* The human brain can be viewed as a complex dynamical system where ongoing changes in neuronal dynamics produce adaptive behavior [31]. These dynamics can be expressed in terms of interactions between brain regions, which is commonly referred to as functional connectivity. At the same time, direct functional interactions presuppose anatomical links between brain regions, known as structural connectivity. For this reason, structural and functional connectivity must be intimately related [32].

Functional connectivity is most easily expressed using a covariance matrix that, in the case of standardized data, provides the correlation structure between different brain regions. However, this approach suffers from the drawback that it cannot distinguish between direct and indirect connections. Alternatively, one may use partial correlations that capture only direct effects, in the absence of confounding factors. The matrix of partial correlations $\mathbf{R}$ may be obtained from a precision matrix using $r_{ij} = 1$ if $i = j$ and $r_{ij} = -k_{ij}/\sqrt{k_{ii}k_{jj}}$ otherwise. Because functional coupling must be accompanied by an anatomical connection, partial correlations between brain regions not only reveal the strength of these couplings, but also indicate which regions are physically connected. In other words, the joint posterior in (4) becomes a distribution over functional connectivity $\mathbf{K}$ (or, equivalently, $\mathbf{R}$) and structural connectivity G.

We proceed by approximating the joint posterior using both the DCBF algorithm as well as the DCT sampler. Both algorithms were executed for 100,000 iterations, of which the first 50,000 were discarded as burn in. Once again, we set $P(G) \propto 1$ and $P(\mathbf{K} \mid G) = \mathcal{W}_G(3, \mathbf{I}_p)$. The algorithms yield almost identical results, as shown by an MSE of edge probabilities of 0.0006 and a symmetrized Kullback-Leibler divergence of 0.0002.

Figure 1B shows the posterior edge probabilities and partial correlations produced by the DCT algorithm. The structural connectivity estimate shows that the majority of edges is associated with either very high or very low edge probabilities. The functional connectivity estimate shows that functional homologues in left and right hemispheres are associated with high partial correlations (expected partial correlations $\langle r \rangle$ in the range $[0.48, 0.73]$), indicating that these functional homologues have similar functional roles. Within a cortical hemisphere, the most salient functional interactions (highest expected partial correlations with $\langle r \rangle$ in the range $[0.23, 0.61]$) are given bilaterally by amygdala−hippocampus, pallidum−putamen, accumbens−caudate, caudate−thalamus,

and hippocampus–thalamus. These functional interactions can be explained by direct pathways as well as unobserved common inputs that induce a high partial correlation. Interestingly, edges with high posterior probability (edge probability higher than 0.999) can be associated with weak absolute partial correlations (with $\langle r \rangle$ as low as 0.1). This indicates that there exist weakly coupled regions (from the linear correlation point of view) that cannot be explained away by other functional interactions.

5. Discussion

We have proposed two novel algorithms for Bayesian model selection in a Gaussian graphical model. The first algorithm combines a direct manner to sample G-Wishart variates [10] with an efficient way of computing conditional Bayes factors when comparing two different models [14], resulting in an improved Metropolis-Hastings approach. The second approach integrates the direct sampler within a birth-death continuous time Markov process [15]. Both algorithms provide accurate estimates of the posterior graphs and precision matrices and are substantially faster (up to a factor of 3.80) than previously available alternatives. We demonstrate the use of the algorithms by estimating, for the first time, both structural and functional connectivity simultaneously using fMRI data.

In future work we aim to improve mixing of the samplers by introducing moves between graphs that differ by more than a single edge. Similarly, one may conceive events other than births and deaths of edges. In either case, the corresponding conditional Bayes factors must be derived, and these should be more efficient to compute than a series of consecutive edge additions and removals. We expect that this will further contribute to efficient estimation of Gaussian graphical models.

References

[1] A P Dempster. Covariance selection. *Biometrics*, pages 157–175, 1972.

[2] P Diaconis and D Ylvisaker. Conjugate priors for exponential families. *Ann Stat*, 7(2):269–281, 1979.

[3] A P Dawid and S L Lauritzen. Hyper Markov laws in the statistical analysis of decomposable graphical models. *Ann Stat*, 21:1272–1317, 1993.

[4] A Roverato. Hyper inverse Wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models. *Scand J Stat*, 29(3):391–411, 2002.

[5] A Atay-Kayis and H Massam. A Monte Carlo method for computing the marginal likelihood in nondecomposable Gaussian graphical models. *Biometrika*, 92(2):317–335, 2005.

[6] M Piccioni. Independence structure of natural conjugate densities to exponential families and the Gibbs sampler. *Scand J Stat*, 27:111–127, 2000.

[7] N Mitsakakis, H Massam, and M D Escobar. A Metropolis-Hastings based method for sampling from the G-Wishart distribution in Gaussian graphical models. *Electron J Stat*, 5:18–30, 2011.

[8] A Dobra, A Lenkoski, and A Rodriguez. Bayesian inference for general Gaussian graphical models with application to multivariate lattice data. *J Am Statist Assoc*, 106:1418–1433, 2011.

[9] H Wang and S Z Li. Efficient Gaussian graphical model determination under G-Wishart distributions. *Electron J Stat*, 6:168–198, 2012.

[10] A Lenkoski. A direct sampler for G-Wishart variates. *Stat*, 2(1):119–128, 2013.

[11] T Hastie, R Tibshirani, and J Friedman. *The elements of statistical learning: data mining, inference and prediction.* Springer, 2 edition, 2009.

[12] I Murray, Z Ghahramani, and D J C MacKay. MCMC for doubly-intractable distributions. In *Proc 22nd Ann Conf Uncertainty in Artificial Intelligence (UAI-06)*, pages 359–366. AUAI Press, 2006.

[13] P J Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82:711–732, 1995.

[14] Y Cheng and A Lenkoski. Hierarchical Gaussian graphical models: Beyond reversible jump. *Electron J Stat*, 6:2309–2331, 2012.

[15] A Mohammadi and E C Wit. Bayesian model selection in sparse Gaussian graphical models. *arXiv:1210.5371v6*, 2014.

[16] M Stephens. Bayesian analysis of mixture models with an unknown number of components — an alternative to reversible jump methods. *Ann Stat*, 28(1):40–74, 2000.

[17] S M Smith, D Vidaurre, C F Beckmann, M F Glasser, M Jenkinson, and et al. Miller, K L. Functional connectomics from resting-state fMRI. *Trends Cog Sci*, 17(12):666–682, 2013.

[18] E Salinas and T J Sejnowski. Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(1):539–550, 2001.

[19] K Friston. Functional and effective connectivity: a review. *Brain Conn*, 1(1):13–35, 2011.

[20] B Moghaddam, B Marlin, E Khan, and K Murphy. Accelerating Bayesian structural inference for non-decomposable Gaussian graphical models. In Y Bengio, D Schuurmans, J Lafferty, C K I Williams, and A Culotta, editors, *Adv Neural Inf Process*, pages 1285–1293. Curran Associates Inc., 2009.

[21] A Dobra and A Lenkoski. Copula Gaussian graphical models and their application to modeling functional disability data. *Ann Appl Stat*, 5(2A):969–993, 2011.

[22] F Liang. A double Metropolis-Hastings sampler for spatial models with intractable normalizing constants. *J Stat Comput Simul*, 80:1007–1022, 2010.

[23] O Cappé, C P Robert, and T Rydén. Reversible jump, birth-and-death and more general continuous time Markov chain Monte Carlo samplers. *J R Stat Soc: B*, 65(3):679–700, 2003.

[24] S Kullback and R A Leibler. On information and sufficiency. *Ann Math Stat*, 22(1):79–86, 1951.

[25] V Rao and Y W Teh. MCMC for continuous-time discrete-state systems. In F Pereira, C J C Burges, L Bottou, and K Q Weinberger, editors, *Adv Neural Inf Process 25*, pages 701–709. Curran Associates, Inc., 2012.

[26] M Hinne, L Ambrogioni, R J Janssen, T Heskes, and M van Gerven. Structurally-informed Bayesian functional connectivity analysis. *Neuroimage*, 68:294–305, 2014.

[27] E S B van Oort, A M van Cappellen van Walsum, and D G Norris. An investigation into the functional and structural connectivity of the default mode network. *NeuroImage*, 90:381–389, 2014.

[28] M Jenkinson, C F Beckmann, T E Behrens, M W Woolrich, and S M Smith. FSL. *NeuroImage*, 62:782–790, 2012.

[29] B A Poser, M J Versluis, J M Hoogduin, and D G Norris. BOLD contrast sensitivity enhancement and artifact reduction with multiecho EPI: parallel-acquired inhomogeneity-desensitized fMRI. *Magn Reson Med*, 55(6):1227–1235, 2006.

[30] B Patenaude, S M Smith, D Kennedy, and Jenkinson M. A Bayesian model of shape and appearance for subcortical brain. *NeuroImage*, 56(3):907–922, 2011.

[31] E Bullmore and O Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, 10(3):186–198, 2009.

[32] H Akil, M E Martone, and D C van Essen. Challenges and opportunities in mining neuroscience data. *Science*, 331(6018):708–712, 2011.