



HAL
open science

Shortened Bridge Sampler: Using deterministic approximations to accelerate SMC for posterior sampling

Sophie Donnet, Stéphane Robin

► **To cite this version:**

Sophie Donnet, Stéphane Robin. Shortened Bridge Sampler: Using deterministic approximations to accelerate SMC for posterior sampling. 2017. hal-01566898

HAL Id: hal-01566898

<https://hal.science/hal-01566898>

Preprint submitted on 21 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Shortened Bridge Sampler:

Using deterministic approximations to accelerate SMC for posterior sampling

Sophie Donnet · Stéphane Robin

Received: date / Accepted: date

Abstract Sequential Monte Carlo has become a standard tool for Bayesian Inference of complex models. This approach can be computationally demanding, especially when initialized from the prior distribution. On the other hand, deterministic approximations of the posterior distribution are often available with no theoretical guaranties. We propose a bridge sampling scheme starting from such a deterministic approximation of the posterior distribution and targeting the true one. The resulting Shortened Bridge Sampler (SBS) relies on a sequence of distributions that is determined in an adaptive way.

We illustrate the robustness and the efficiency of the methodology on a large simulation study. When applied to network datasets, SBS inference leads to different statistical conclusions from the one supplied by the standard variational Bayes approximation.

Keywords Bayesian statistics · Sequential Monte Carlo · Approximate posterior distribution · Bridge sampling

1 Introduction

In Bayesian statistics, except in a restricted number of conjugate models, the posterior distribution does not have a close form and requires the use of approximation methods. The 1990's have witnessed the developments of powerful stochastic methods and simulation-based algorithms able to perform Bayesian statistical inference on complex statistical models.

Among them, Monte Carlo Markov Chains (MCMC) –whose principle is to generate a Markov Chain such that its ergodic distribution is the posterior distribution (see for instance Robert and Casella (2005, 2009) for an introduction)– have been successfully applied to many problems, as attested by the countless publications in that domain. However, MCMCs suffer from several drawbacks. First of all, it is difficult to assess whether the chain has reached its ergodic distribution or not. Secondly, if the distribution of interest is highly multi-modal, MCMC algorithms

can be trapped in local modes. More generally, when the space of parameters to be explored is of high dimension, MCMC algorithms will have difficulties in reaching their equilibrium distribution within a reasonable computational time.

Recently, population based Monte Carlo methods have proved their efficiency and robustness in front of high dimensional and multimodal spaces. In a few words, population based Monte Carlo algorithms generate a large sample of parameters with a tractable distribution and update the importance sampling weights at each iteration, in order to finally match the distribution of interest. Among population based Monte Carlo, Sequential Monte Carlo (SMC) is a method combining parameters sampling and resampling. More precisely, a sequence of distributions of interest is designed, such that the first one is simple (i.e. easy to sample from) and the last one is the posterior distribution. This sequence of distributions defines the iterations of the algorithm. Then, at the first iteration, a sample of parameters is simulated with the first distribution. In the following iterations, the parameters are stochastically moved, weighted and resampled to follow the current distribution. The true posterior distribution is reached at the last iteration. The sequence of distributions can be dynamically designed. Primarily developed in the context of filtering problems (see Doucet et al (2001)), they have been extended to the general problem of posterior sampling by Del Moral et al (2006). In comparison with MCMC methods, SMC does not require any burn-in period or convergence diagnostic. In addition, whereas computing marginal likelihood (for model comparison) has always been a challenging issue when using MCMC, SMC supplies an unbiased estimator of this quantity as a by-product of the algorithm. For all these reasons, SMC has proved its superiority over MCMC for complex models.

In the recent years, particular fields (such as genomics or network analysis to name but a few) brought news statistical problems involving an increasing amount of data or statistical models with a large number of parameters. In such cases, not only MCMC but also population Monte Carlo algorithms have reached their limitations, requiring unreasonable computational time to explore the posterior distribution. To deal with such difficulties, deterministic approximations of the posterior distribution through optimization mathematical tools – such as variational approximation (Wainwright and Jordan (2008); Blei et al (2016)), Expectation-Propagation (Minka (2001)) or Integrated nested Laplace approximation (Rue et al (2009)) for instance – have been proposed. These methods have the great advantage to be computationally light and can handle large data. However, their theoretical properties and accuracy is still under study. In particular, we do know that variational approximations can supply underestimated posterior variances (see for instance Consonni and Marin (2007) for a large illustration of this phenomena on the Probit model).

On the one hand, SMC supplies a sample from the exact posterior distribution but can require unacceptable computational time. On the other hand, deterministic approximations (optimal in a sense to be determined) of the posterior distribution are fast methods but non-exact. One may therefore be tempted to take advantage of the two approaches in a combined strategy. The idea of combining variational Bayes inference with SMC is actually not new. Rabinovich et al (2015) split the data into block and compute the posterior distribution of θ given each block. They use a variational argument to propose the product of this partial posterior as a proxy for the true posterior. Focusing on Gaussian mixtures, McGrory et al (2016) consider online-inference and propose an sequential sampling scheme

where, for each new batch of data, the variational approximation is iteratively updated and used as a proposal. Naesseth et al (2017) use a SMC approach to get an improved, but still biased, variational approximation. Our approach is different from all these. Our main idea is to design a bridge sampling from the approximated posterior distribution to the true posterior distribution, the transfer from the approximate to the exact distribution being performed with an SMC algorithm (Del Moral et al (2006)). The sampling method we propose can be considered from two points of view : either SMC is seen as a tool to correct the approximate distribution, or the approximate posterior distribution is seen as a mean to drastically accelerate the SMC procedure.

Adopting the last perspective plunges the problematic into in a larger topic. Indeed, in general, for any challenging statistical model at stake, there exists a frequentist solution, suppling a point estimation of the parameter of interest. In the Bayesian practice, this point estimation is standardly set as an initial value in the MCMC algorithm, thus hoping a decrease of the computational time. In a SMC strategy starting with a sample from the prior distribution, such an initial point value is meaningless. We claim that the posterior distribution can be reached in a reduced computational time if the bridge sampling scheme starts from an approximate posterior distribution based on that point estimate. We therefore refer to the proposed sampling method as the Shortened Bridge Sampler (SBS).

The paper is organized as follows. Section 2 is dedicated to the description of the methodology. We first remind the principle of importance sampling, then introduce the sampling path and expose the algorithm in Subsection 2.3. Its robustness and efficiency are illustrated on several simulated experiments in Section 3. In Subsection 3.1, the logit regression serves as a toy example to illustrate the computational time reduction and to test the robustness of the method with respect to the quality of the deterministic approximation of the posterior distribution. The Latent Class Analysis model (Subsection 3.2) is exploited to illustrate the relevance of our methodology on a mixture model, in particular, we propose a new strategy to tackle the label switching issue. On the Stochastic Block Model (SBM) with covariates (Subsection 3.3), we compare our strategy with the Variational Bayesian one in terms of model selection and model averaging. Finally, real datasets of social networks with covariates are presented in Section 4 : we stress the new insights brought by a “correction” of the Variational posterior approximation by the SMC strategy in term of both model averaging and significance of the parameters. Perspectives are discussed in Section 5.

2 From the approximate posterior distribution to the true posterior distribution

Let us first introduce some notations. \mathbf{Y} denotes the observations, $\ell(\mathbf{Y}|\theta)$ is the likelihood function with $\theta \in \Theta$ the unknown parameters and $\pi(\theta)$ is the prior distribution on θ . The Bayesian inference is based on the posterior distribution:

$$p(\theta|\mathbf{Y}) = \frac{\ell(\mathbf{Y}|\theta)\pi(\theta)}{p(\mathbf{Y})}. \quad (1)$$

where $p(\mathbf{Y})$ is the marginal likelihood defined as:

$$p(\mathbf{Y}) = \int \ell(\mathbf{Y}|\theta)\pi(\theta)d\theta \quad (2)$$

and is required in the Bayesian model selection procedure.

In what follows, $\tilde{p}_{\mathbf{Y}}$ is an approximate posterior distribution on θ . We assume that $\tilde{p}_{\mathbf{Y}}$ can be easily intensively simulated and that the density function of $\tilde{p}_{\mathbf{Y}}$ has an explicit expression. The aim of this paper is to propose a way to use such an approximate posterior to actually sample from the true posterior.

Remark 1 *Note that, in general, complex statistical models are written as hierarchical models and involve latent variables \mathbf{Z} (see Sections 3.2 and 3.3). In such cases, the distributions of interest are the joint distribution $p(\mathbf{Z}, \theta|\mathbf{Y})$ or the marginal one $p(\theta|\mathbf{Y})$. For the sake of simplicity, we chose to present the method without latent variables but obviously, all the following results and algorithms can be extended to this situation, replacing θ by (\mathbf{Z}, θ) . A substantial part of the simulations presented in Section 3 is devoted to such models.*

2.1 A first approach: Importance Sampling

A first naive approach to use $\tilde{p}_{\mathbf{Y}}$ consists in resorting to a simple importance sampling (IS) strategy, that is to say sampling $(\theta^m)_{m=1\dots M}$ from $\tilde{p}_{\mathbf{Y}}$ and weighting the sample by $W^m \propto \ell(\mathbf{Y}|\theta^m)\pi(\theta^m)/\tilde{p}_{\mathbf{Y}}(\theta^m)$. However, this strategy is obviously naive for several reasons. First of all, there is no guarantee that the support of the approximate distribution includes the support of the true distribution, the contrary is even observed in practice for the variational approximation for instance (see Consonni and Marin, 2007; Wang and Titterton, 2005). As a consequence, the posterior sample obtained through such an importance sampling strategy would be restricted to the support of $\tilde{p}_{\mathbf{Y}}$ which can be strictly included in the support of $p(\cdot|\mathbf{Y})$. Secondly, if $\tilde{p}_{\mathbf{Y}}$ and $p(\cdot|\mathbf{Y})$ are very different, the sample will be degenerated, meaning that very few particles will have a non-negligible weight. This results in a small Effective Sample Size (*ESS*), which the algorithm we propose aims at keeping along iterations. In such situations, there is no hope to efficiently sample using 'one-step' IS but the principle can be used iteratively to progressively shift from the initial proposal to the true posterior distribution.

2.2 A path sampling between the approximate and the true posterior distributions

The main idea of this paper is to take advantage of the deterministic approximation $\tilde{p}_{\mathbf{Y}}$ of the posterior distribution to accelerate SMC procedure, or, inversely to transform sequentially a sample from the deterministic approximated posterior distribution into a sample from the true posterior distribution.

Sequential Monte Carlo samplers generate samples from a sequence of intermediate distributions $(p_h)_{h=0\dots H}$ where the intermediate distributions $(p_h)_{h=0\dots H}$ are

smooth transitions from a simple distribution p_0 to the distribution of interest $p_H = p(\cdot|\mathbf{Y})$. A classical choice for $(p_h)_{h=0\dots H}$ (Neal, 2001) is to consider:

$$p_h(\theta) \propto \pi(\theta)\ell(\mathbf{Y}|\theta)^{\rho_h} \quad (3)$$

where $\rho_0 = 0$, $\rho_H = 1$, thus slowly moving from the prior distribution to the posterior by progressively integrating the data \mathbf{Y} through the likelihood function. In this paper, we propose an alternative scheme moving smoothly from the approximate posterior distribution $\tilde{p}_{\mathbf{Y}}$ to the true $p(\cdot|\mathbf{Y})$. The path is thus defined by:

$$\begin{aligned} p_h(\theta) &\propto \tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h} (p(\theta|\mathbf{Y}))^{\rho_h} \\ &\propto \tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h} (\pi(\theta)\ell(\mathbf{Y}|\theta))^{\rho_h}. \end{aligned} \quad (4)$$

where, $\rho_0 = 0$, $\rho_H = 1$. In a few words, we start from the easy-to-sample distribution $\tilde{p}_{\mathbf{Y}}(\theta)$ and progressively replace it with the true posterior distribution, this strategy being known as annealed importance sampling procedure (Neal, 2001). We claim that this scheme significantly reduces the computational time and is robust with respect to $\tilde{p}_{\mathbf{Y}}$.

Remark 2 Note that if $\tilde{p}_{\mathbf{Y}}$ is chosen to be the prior distribution $\pi(\cdot)$, then schemes (3) and (4) are identical.

Remark 3 An alternative strategy would consist in using $\tilde{p}_{\mathbf{Y}}$ has an importance sampler in the first iteration of the standard annealing scheme defined in (3). However, in such a strategy, the approximate distribution $\tilde{p}_{\mathbf{Y}}$ is under-exploited, since at the first iteration, the particles are reweighted with $W_0^m \propto \pi(\theta^m)/\tilde{p}_{\mathbf{Y}}(\theta^m|\mathbf{Y})$, thus going back, in practice, to a (possibly truncated) version of the prior distribution. This phenomena will be illustrated in Section 3.1.

To sample from the sequence of distributions $(p_h)_{h=1,\dots,H}$, we adopt the sequential sampler proposed by Del Moral et al (2006) where the annealing coefficients ρ_h will be adjusted dynamically. We describe the algorithm in the following subsection.

2.3 Shortened Bridge Sampling Algorithm

We now need to design an algorithm sequentially sampling from $p_h(\theta) \propto \tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h} (\pi(\theta)\ell(\mathbf{Y}|\theta))^{\rho_h}$. The last years have witnessed a proliferation of scientific papers dealing with the problem of SMC methods and their applications (see Doucet et al, 2001, for an overview). In our work, we resort to the algorithm proposed by Del Moral et al (2006).

Let us introduce the following notations:

$$\gamma_h(\theta) = \tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h} [\ell(\mathbf{Y}|\theta)\pi(\theta)]^{\rho_h}, \quad (5)$$

where $Z_h = \int \gamma_h(\theta) d\theta$, so that $p_h(\theta) = \gamma_h(\theta)/Z_h$ is a probability density. The main idea of Del Moral et al (2006) is to plunge the problem of sampling a sequence of distributions defined on a single set Θ into the standard SMC filtering

framework. To that purpose, the sequence $(p_h)_{h=0\dots H}$ is replaced by a sequence of extended distributions:

$$\bar{p}_h(\theta_{0:h}) = \frac{\bar{\gamma}_h(\theta_{0:h})}{Z_h} \quad (6)$$

with

$$\bar{\gamma}_h(\theta_{0:h}) = \gamma_h(\theta_h) \prod_{k=1}^h L_k(\theta_{k-1}|\theta_k) \quad (7)$$

where $\theta_{0:h} = (\theta_0, \dots, \theta_h) \in \Theta \times \dots \times \Theta = \Theta^{h+1}$ and $(L_k)_{k=0, \dots, H-1}$ is a sequence of backward kernels satisfying:

$$\int L_k(\theta_{k-1}|\theta_k) d\theta_{k-1} = 1, \quad \forall k = 0 \dots H-1. \quad (8)$$

Due to Property (8), the marginal version of \bar{p}_h (i.e. when integrating out $\theta_0, \dots, \theta_{h-1}$) is the distribution of interest p_h . Once defined the sequence $(\bar{p}_h)_{h=0\dots H}$, one may use the original SMC algorithm designed by Doucet et al (2001) for filtering. At iteration h , the SMC sampler involves three steps:

- *Moving the particles* from θ_{h-1} to θ_h using a transition kernel $K_h(\theta_h|\theta_{h-1})$. As a consequence, let $\eta_{h-1}(\theta_{0:h-1})$ denote the sampling kernel for $\theta_{0:h-1}$ until iteration $h-1$, η_h 's expression is:

$$\eta_h(\theta_{0:h}) = \eta_{h-1}(\theta_{0:h-1}) K_h(\theta_h|\theta_{h-1}) \quad (9)$$

- *Reweighting the particles* in order to correct the discrepancy between the sampling distribution η_h and the distribution of interest at iteration h , \bar{p}_h .
- *Selecting the particles* in order to reduce the variability of the importance sampling weights and avoid degeneracy. In practice the particles will be resampled when the *ESS* decreases below a pre-specified rate.

About the importance weights. At iteration h , the importance sampling weights for $(\theta_{0:h}^m)_{m=1\dots M}$ are : $\forall m = 1 \dots M$,

$$w_h^m = w_h(\theta_{0:h}^m) = \frac{\bar{\gamma}_h(\theta_{0:h}^m)}{\eta_h(\theta_{0:h}^m)} \quad (10)$$

in their unnormalized version. $(W_h^m)_{m=1\dots M}$ denotes the normalized weights, i.e.

$$W_h^m = \frac{w_h^m}{\sum_{m'=1}^M w_h^{m'}}, \quad \forall m = 1 \dots M \quad (11)$$

Equations (6-7-9-10) imply a recurrence formula for the weight of any particle $\theta_{0:h}$:

$$w_h(\theta_{0:h}) = w_{h-1}(\theta_{0:h-1}) \tilde{w}_{h-1:h}(\theta_{h-1}, \theta_h) \quad (12)$$

where the incremental weight $\tilde{w}_{h-1:h}(\theta_{h-1}, \theta_h)$ is equal to:

$$\tilde{w}_{h-1:h}(\theta_{h-1}, \theta_h) = \frac{L_h(\theta_{h-1}|\theta_h)}{K_h(\theta_h|\theta_{h-1})} \frac{\gamma_h(\theta_h)}{\gamma_{h-1}(\theta_{h-1})} \quad (13)$$

About the transition kernel K_h . As, at this step, the target distribution is p_h , it seems natural to choose $K_h(\theta_h|\theta_{h-1})$ as a Monte Carlo Markov Chain (MCMC) kernel with $p_h(\theta) \propto \tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h}(\ell(\mathbf{Y}|\theta)\pi(\theta))^{\rho_h}$ as stationary distribution. Following Del Moral et al (2006), we choose the backward kernel:

$$L_h(\theta_{h-1}|\theta_h) = \frac{K_h(\theta_h|\theta_{h-1})p_h(\theta_{h-1})}{p_h(\theta_h)} \quad (14)$$

which satisfies Property (8) and enables us to rewrite the weight increment $\tilde{w}_{h-1:h}(\theta_{h-1}, \theta_h)$ appearing in (12) and defined in (13) as

$$\tilde{w}_{h-1:h}(\theta_{h-1}, \theta_h) = \frac{\gamma_h(\theta_{h-1})}{\gamma_{h-1}(\theta_{h-1})} = [\alpha(\theta_{h-1})]^{\rho_h - \rho_{h-1}} \quad (15)$$

where

$$\alpha(\theta) = \frac{\ell(\mathbf{Y}|\theta)\pi(\theta)}{\tilde{p}_{\mathbf{Y}}(\theta|\mathbf{Y})}. \quad (16)$$

In what follows, we denote $\alpha_h = \alpha(\theta_h)$.

Remark 4 Using this particular backward kernel (14) has two major consequences. First it is not required having an explicit expression for the transition kernel $K_h(\theta_h|\theta_{h-1})$, which is quite welcome for MCMC kernels. Secondly, examining equations (12) and (15), one may notice that the weight for a particle $\theta_{0:h}$ does not depend on θ_h but only on $\theta_{0:h-1}$. As a consequence, the weights of the particles $\theta_{0:h}$ can be computed before they are simulated and for any new p_h

Adaptive design of $(\rho_h)_{h=0\dots H}$. As a consequence of this last remark, we are able to design an adaptive strategy for $(\rho_h)_{h=0\dots H}$ (as in Schäfer and Chopin, 2013; Jasra et al, 2011). Indeed, being able to compute the weights of the up-coming particles for any new ρ_h , we can increase ρ_h until the quality of the sample (measured through an indicator computed from the weights) decreases for the next distribution. In practice, following Zhou et al (2016), we use the conditional Effective Sampling Size (*cESS*) to measure the quality of p_{h-1} as an importance sampler when estimating an expectation against p_h . It is defined as:

$$\begin{aligned} cESS &= \left[\sum_{m=1}^M MW_{h-1}^m \left(\frac{\tilde{w}_{h-1:h}^m}{\sum_{m=1}^M MW_{h-1}^m \tilde{w}_{h-1:h}^m} \right)^2 \right]^{-1} \\ &= \frac{M \left(\sum_{m=1}^M W_{h-1}^m \tilde{w}_{h-1:h}^m \right)^2}{\sum_{m=1}^M W_{h-1}^m (\tilde{w}_{h-1:h}^m)^2}, \end{aligned}$$

becoming

$$\begin{aligned} cESS(\rho_h; \rho_{h-1}, (W_{h-1}^m, \alpha_{h-1}^m)_{m \leq M}) &= cESS_{h-1}(\rho) \\ &= \frac{M \left(\sum_{m=1}^M W_{h-1}^m (\alpha_{h-1}^m)^{\rho_h - \rho_{h-1}} \right)^2}{\sum_{m=1}^M W_{h-1}^m (\alpha_{h-1}^m)^{2(\rho_h - \rho_{h-1})}}. \end{aligned} \quad (17)$$

where $\alpha_{h-1}^m = \alpha(\theta_{h-1}^m)$ has been defined in equation (16). If $\rho_h = \rho_{h-1}$, *cESS* is maximal (equal to M , the number of particles). As ρ_h increases, the discrepancy between p_{h-1} and p_h increases and so the quality of p_{h-1} as an importance

sampling distribution when estimating an expectation against p_h decreases and so does $cESS$. As a consequence, our strategy to find the next ρ_h is to set:

$$\rho_h = 1 \wedge \sup_{\rho} \{\rho > \rho_{h-1}, cESS_{h-1}(\rho) \geq \tau_1 M\}$$

Selection of the particles. In order to prevent a degeneration of the particle approximation, we use a standard resampling of the particles. In other words, if the variance of weights $(W_h^m)_{m=1\dots M}$ is too high (or in other words, if the ESS is too small), we resample the particles using a multinomial distribution, thus discarding the particles with low weights and duplicating the particles with high weights.

Sampling algorithm. Finally, we propose the following Shorten Bridge Sampling SBS algorithm adapted to the sequence (4).

SBS algorithm

Set $(\tau_1, \tau_2) \in [0, 1]^2$, $\rho_0 = 0$.

0. At iteration 0, sample $(\theta_0^m)_{m=1\dots M}$ from the approximate distribution $\tilde{p}_{\mathbf{Y}}$.
 $\forall m = 1 \dots M$, set:

$$w_0^m = 1, \quad W_0^m = \frac{1}{M}, \quad \alpha_0^m = \frac{\ell(\mathbf{Y}|\theta_0^m)\pi(\theta_0^m)}{\tilde{p}_{\mathbf{Y}}(\theta_0^m)}$$

1. At iteration h : starting from $(\theta_{h-1}^m, W_{h-1}^m, \alpha_{h-1}^m)_{m=1\dots M}$
 (a) Find ρ_h such that:

$$\rho_h = 1 \wedge \sup_{\rho} \{\rho > \rho_{h-1}, cESS_{h-1}(\rho) \geq \tau_1 M\},$$

- (b) $\forall m = 1 \dots M$, compute $w_h^m = w_{h-1}^m (\alpha_{h-1}^m)^{\rho_h - \rho_{h-1}}$ and $W_h^m = w_h^m / \sum_{m'=1}^M w_h^{m'}$
 (c) Compute

$$ESS_h = \frac{\left(\sum_{m=1}^M W_h^m\right)^2}{\sum_{m=1}^M (W_h^m)^2} \in [1, M]$$

If $ESS_h < \tau_2 M$, resample the particles

$$\begin{aligned} (\theta_{h-1}^m)' &\sim_{i.i.d} \sum_{m=1}^M W_h^m \delta_{\{\theta_{h-1}^m\}}(\theta_h^m) \\ \theta_{h-1}^m &\leftarrow (\theta_{h-1}^m)' \\ w_h^m &\leftarrow 1 \\ W_h^m &\leftarrow 1/M \end{aligned} \quad \forall m = 1 \dots M$$

- (d) $\forall m = 1 \dots M$, : propagate the particle $\theta_h^m \sim K_h(\cdot|\theta_{h-1}^m)$ where K_h is a MCMC kernel with $p_h(\theta) \propto \tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h} (\ell(\mathbf{Y}|\theta)\pi(\theta))^{\rho_h}$ as an invariant distribution and compute:

$$\alpha_h^m = \frac{\ell(\mathbf{Y}|\theta_h^m)\pi(\theta_h^m)}{\tilde{p}_{\mathbf{Y}}(\theta_h^m|\mathbf{Y})}$$

2. If $\rho_h = 1$, stop. If $\rho_h < 1$ return to 1.

Remark 5 Let ϕ be a function defined in Θ . The study of the statistical properties of $\sum_{m=1}^M W_m^H \phi(\theta_m^H)$ as an estimator of $\mathbf{E}[\phi(\theta)|\mathbf{Y}]$ is a difficult task due to the sampling and resampling steps in the algorithm. However, many results can be found in the literature (see Doucet and Johansen, 2009, and references there in). First of all, $\sum_{m=1}^M W_m^H \phi(\theta_m^H)$ is known to be strongly convergent. Moreover, following Del Moral et al (2006), a Central Limit Theorem can be obtained. Besides, in addition to these asymptotic properties, it is possible to control the mean-square error of the estimator for a given number of particles M , provided additional assumptions on ϕ . Results of convergence were also provided by Del Moral et al (2012) for adaptive sequential Monte Carlo algorithms.

2.4 Estimation of the marginal likelihood

With respect to MCMC strategies, Annealing Importance Sampling and SMC have the great advantage to supply good estimators of the marginal likelihood. Indeed, as proved by Del Moral et al (2006), a non-biased estimator of the marginal likelihood derives as a by-product of SMC. Moreover, the path sampling identity also provides an estimate of the marginal likelihood, as detailed hereafter.

Let us recall that $Z_h = \int_{\Theta} \gamma_h(\theta) d\theta$. Following Del Moral et al (2006) and using the notations introduced below, the ratio of the marginal likelihoods Z_h/Z_{h-1} is estimated by:

$$\frac{\widehat{Z}_h}{\widehat{Z}_{h-1}} = \sum_{m=1}^M W_h^m \widetilde{w}_{h-1:h}^m.$$

and

$$\frac{\widehat{Z}_H}{\widehat{Z}_0} = \prod_{h=1}^H \frac{\widehat{Z}_h}{\widehat{Z}_{h-1}} = \prod_{h=1}^H \sum_{m=1}^M W_h^m \widetilde{w}_{h-1:h}^m \quad (18)$$

is an unbiased estimator of Z_H/Z_0 . Now, $Z_H = p(\mathbf{Y})$ and $Z_0 = 1$.

Another estimate is given by the path sampling identity. Indeed, under non-restrictive regularity assumptions, the following equality holds:

$$\log p(\mathbf{Y}) - \log Z_0 = \int_0^1 \mathbb{E}_{p_\rho} \left[\frac{d \log \gamma_\rho(\cdot)}{d\rho} \right] d\rho \quad (19)$$

where $\gamma_\rho(\theta) = \widetilde{p}_{\mathbf{Y}}(\theta)^{1-\rho} (\ell(\mathbf{Y}|\theta)\pi(\theta))^\rho$, $p_\rho(\cdot)$ is the associated probability density distribution and, in our geometric path sampling:

$$\frac{d \log \gamma_\rho(\cdot)}{d\rho} = \log \frac{\ell(\mathbf{Y}|\theta)\pi(\theta)}{\widetilde{p}_{\mathbf{Y}}(\theta)} = \log \alpha(\theta)$$

An elementary trapezoidal scheme and Monte Carlo approximations of the expectations involved in (19) lead to the following approximation of the marginal likelihood:

$$\widehat{\widehat{\log p(\mathbf{Y})}} = \sum_{h=1}^H \frac{\rho_h - \rho_{h-1}}{2} (U_h^M + U_{h-1}^M) \quad (20)$$

where $U_h^M = \widehat{\mathbb{E}}_{p_{\rho_h}} [\log \alpha(\theta)] = \sum_{m=1}^M W_h^m \log \alpha_h^m$.

Remark 6 *Note that as suggested in Zhou et al (2016), we noticed on simulation studies that the two estimators behave similarly in our examples. A precise comparison of the two estimators is out of the scope of this paper.*

3 Simulation study

We now present a large simulation study. The goal of the study is to assess the fact that the proposed SBS algorithm –that combines an optimization-based approximation of the posterior distribution with a SMC sampler– drastically decreases the computational time with respect to a classical annealing-scheme or, equivalently, that the approximated posterior distribution can be corrected into the true posterior distribution at a low computational cost.

3.1 Logistic regression

This first model is used as a toy example to illustrate the efficiency and the robustness of our methodology. Let (Y_1, \dots, Y_n) be a set of n independent observations with values in $\{0, 1\}$. Any individual observation i is described by a vector $\mathbf{x}_i \in \mathbb{R}^p$ of p covariates and we consider the logistic regression model:

$$P(Y_i = 1) = 1 - P(Y_i = 0) = \frac{e^{\mathbf{x}_i^t \theta}}{1 + e^{\mathbf{x}_i^t \theta}}$$

We generate a simulated dataset with a randomly chosen regression matrix X and the following parameters:

$$n = 200, \quad p = 4, \quad \theta = (0.5, -0.6, 0, -1).$$

Setting a Gaussian prior distribution on $\theta \in \mathbb{R}^p$, $\theta \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_p)$, with $\sigma^2 = 100$, it is natural to propose a Gaussian approximation of the posterior distribution $\tilde{p}_{\mathbf{Y}} := \mathcal{N}(\hat{\mu}, \hat{\Sigma})$. We first computed the Gaussian variational Bayes estimator and we obtained:

$$\hat{\mu}^{VB} = (0.398, -0.643, -0.280, -0.847)$$

and

$$\hat{\Sigma}^{VB} = 10^{-3} \begin{pmatrix} 23.3 & -1.2 & 0.1 & 1.1 \\ -1.2 & 20.6 & -0.1 & 2.0 \\ 0.1 & -0.1 & 24.0 & 1.4 \\ 1.1 & 2.0 & 1.4 & 25.5 \end{pmatrix}.$$

We also considered an approximate Gaussian posterior distribution based on the maximum likelihood estimator (ML) $\hat{\theta}^{ML}$. The ML and its asymptotic variance $\hat{\Sigma}^{ML}$ are obtained with the R-function `glm`. To test the robustness of our bridge sampling with respect to $\tilde{p}_{\mathbf{Y}}$, we also consider an artificially increased (respectively

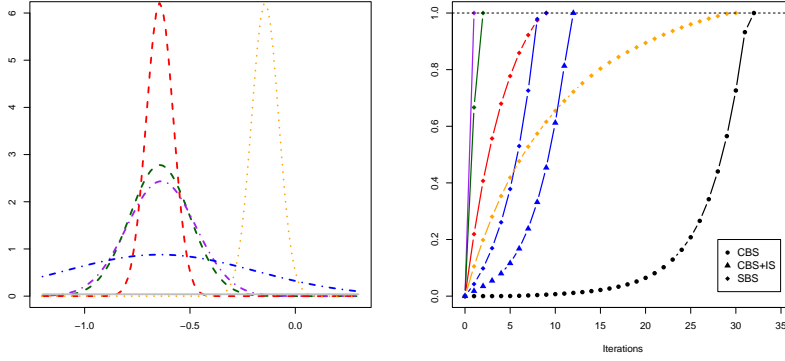


Fig. 1 Logistic regression. *Left frame* : approximate posterior distributions for θ_2 . $\tilde{p}_{\mathbf{Y},1}$ (VB) is in the darkgreen curve, $\tilde{p}_{\mathbf{Y},2}$ (ML) is in purple, $\tilde{p}_{\mathbf{Y},3}$ (VB with decreased variance) is in blue and $\tilde{p}_{\mathbf{Y},4}$ (VB with increased variance) is in orange. *Right frame*: sequences $(\rho_h)_{h \leq 0}$ with the three algorithms and the 5 $\tilde{p}_{\mathbf{Y}}$ along the iterations of algorithm. The result for CBS is the dotted black line. The results for our algorithm SBS are labeled by diamonds \diamond whereas the ones for algorithm CBS+IS are marked by triangles \triangle .

decreased) variance. In addition, we consider a distribution centered on an aberrant value with a small variance, thus leading to five different $\tilde{p}_{\mathbf{Y}}$:

$$\begin{aligned}\tilde{p}_{\mathbf{Y},1} &= \mathcal{N}(\hat{\mu}^{VB}, \hat{\Sigma}^{VB}) \\ \tilde{p}_{\mathbf{Y},2} &= \mathcal{N}(\hat{\theta}^{ML}, \hat{\Sigma}^{ML}) \\ \tilde{p}_{\mathbf{Y},3} &= \mathcal{N}(\hat{\mu}^{VB}, \text{diag}(\hat{\Sigma}^{VB})/5) \\ \tilde{p}_{\mathbf{Y},4} &= \mathcal{N}(\hat{\mu}^{VB}, \text{diag}(\hat{\Sigma}^{VB}) \times 10) \\ \tilde{p}_{\mathbf{Y},5} &= \mathcal{N}(\hat{\mu}^{VB} + 0.5, \text{diag}(\hat{\Sigma}^{VB})/5)\end{aligned}$$

$\tilde{p}_{\mathbf{Y},1}, \tilde{p}_{\mathbf{Y},2}, \tilde{p}_{\mathbf{Y},3}, \tilde{p}_{\mathbf{Y},4}$ and $\tilde{p}_{\mathbf{Y},5}$ are plotted in Figure 1 (for θ_2).

For each $(\tilde{p}_{\mathbf{Y},k})_{k=1\dots 5}$, we sample the posterior distribution using three methods.

- CBS refers to a Classical Bridge Sampling $\pi(\theta)\ell(\mathbf{Y}|\theta)^{\rho_h}$, sequentially sampled with a SMC algorithm. The strategy serves as a reference to be compared to the other ones.
- CBS+IS refers to the same annealing scheme as CBS but the first sample $(\theta_0^m)_{m=1,\dots,M}$ is generated with $\tilde{p}_{\mathbf{Y}}$ and the adequate weights are computed.
- Finally, we use SBS, described in section 2.3 corresponding to the annealing scheme: $\tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h}(\pi(\theta)\ell(\mathbf{Y}|\theta))^{\rho_h}$.

Tunings. In each case, the SMC is performed using $M = 10000$ particles, $\tau_1 = 0.9$ and $\tau_2 = 0.8$. The kernel K_h is a composed of $B = 5$ iteration of a standard Metropolis-Hastings kernel, proposing a new parameter as: $\theta^c \sim \frac{1}{3} \sum_{i=1}^3 \mathcal{N}(\theta^{\ell-1}, \rho_i \times \hat{\Sigma}^{ML})$ where $\hat{\Sigma}^{ML}$ is the asymptotic variance of the maximum likelihood estimator and $(\rho_1, \rho_2, \rho_3) = (1, 0.1, 10)$.




Fig. 2 Logistic regression. Posterior distributions for θ_2 obtained with algorithm CBS+IS (left) and our strategy SBS (right) with the 5 $(\tilde{p}_{\mathbf{Y},k})_{k=1,\dots,5}$ with the following colors: $\tilde{p}_{\mathbf{Y},1}$ (VB) in darkgreen. $\tilde{p}_{\mathbf{Y},2}$ (ML) in purple. $\tilde{p}_{\mathbf{Y},3}$ (VB with decreased variance) in red. $\tilde{p}_{\mathbf{Y},4}$ (VB with increased variance) in blue. $\tilde{p}_{\mathbf{Y},5}$ (VB translated with small variance) in orange. The results for CBS are plotted in black.

Results. The results are plotted in Figures 2 and 1 (right). For the five $(\tilde{p}_{\mathbf{Y},k})_{k=1\dots 5}$, the posterior distributions given by our algorithm SBS (Figure 2, right) are confounded and stick to the one obtained from the reference algorithm CBS, thus illustrating the practical robustness of our new bridge sampler even. More precisely, even if the approximated posterior distribution has an underestimated variance (which is known to be the case for the Variational Bayes estimator in this case), our methodology will supply a sample from the true posterior distribution. Moreover, the algorithm is also robust when $\tilde{p}_{\mathbf{Y}}$ is piked around on an absurd value (see results for $\tilde{p}_{\mathbf{Y},5}$).

On the contrary, using algorithm CBS+IS can be a bad idea. Indeed, if the approximation $\tilde{p}_{\mathbf{Y}}$ has a wider support than the true posterior, the algorithm will perform well and provide a sample from the right posterior distribution: as can be seen in Figure 2, left frame, the black and the blue curves are indistinguishable. However, if $\tilde{p}_{\mathbf{Y}}$ has an underestimated variance ($\tilde{p}_{\mathbf{Y},1}$, $\tilde{p}_{\mathbf{Y},2}$, $\tilde{p}_{\mathbf{Y},4}$ and $\tilde{p}_{\mathbf{Y},5}$), the standard sampling strategy will lead to a false posterior distribution (see the orange, red, purple and green curves in the left frame of Figure 2). Note that, there is no algorithmic indicator detecting such a bad behavior.

To compare the computational times of the various strategies, we can have a look at the number of iterations required for the sequences $(\rho_h)_{h \geq 0}$ to reach 1. These sequences $(\rho_h)_{h \geq 0}$ are plotted in Figure 1 (right). We only plot the curves for the combinations “algorithm/approximated posterior $\tilde{p}_{\mathbf{Y}}$ ” leading to the true posterior distribution. As expected, CBS is the most time consuming (see black dotted curve). As an indicative basis, on this example, the 30 iterations of CBS requires roughly 5 minutes on an Intel (R)Xeon(R) CPU E5-1650 v3 @ 3.50GHz \times 12 using six cores. With $\tilde{p}_{\mathbf{Y},3}$ (increased variance), SBS and CBS+IS finish in quite comparable computational time with slightly better results for our methodology. SBS requires the same computational time with $\tilde{p}_{\mathbf{Y},4}$ (small variance). SBS clearly outperforms CBS. Finally, considering an extreme case where the approximation

distribution $\tilde{p}_{\mathbf{Y},5}$ is concentrated around an aberrant value (which is unlikely to be the case in practice), the SBS and CBS have comparable computation times.

As a conclusion, CBS+IS can not be used in general cases since it can supply a wrong approximation of the posterior distribution. This is due to the fact that, at the first iteration of CBS+IS, the sample must be of the prior distribution. Using IS, i.e. simulating with $\tilde{p}_{\mathbf{Y},i}$ and assessing weight can lead to a sample from the truncated prior distribution $\pi(\theta)1_{\tilde{p}_{\mathbf{Y}}(\theta)>0}$. This simulation from the wrong distribution at the first step of the sequential importance sampler is not corrected in the following iterations. Besides, we are not able to detect such a phenomena. On the contrary, our new bridge sampler behaves well, whatever $\tilde{p}_{\mathbf{Y}}$. An under-evaluated variance in $\tilde{p}_{\mathbf{Y}}$ is not a limit to the use of our scheme. The gain in computational time depends obviously on the distance between $\tilde{p}_{\mathbf{Y}}$ and $p(\cdot|\mathbf{Y})$ but is expected to be drastic when $\tilde{p}_{\mathbf{Y}}$ comes from deterministic approximations of the posterior (Variational Bayes, Expectation Propagation, etc).

3.2 Latent Class Analysis model

In this section we consider the latent class analysis (LCA) model. On this model, we focus on the label switching issue and show that the strategy we propose can tackle this difficulty.

3.2.1 Model and prior distribution

LCA is a mixture model for multivariate binary observations such as the correct or incorrect answers submitted during an exam (Bartholomew et al, 2011), the symptoms presented by persons with major depressive disorder (Garrett and Zeger, 2000) or a disability index recorded by long-term survey (Erosheva et al, 2007). Let $(\mathbf{Y}_i)_{i=1,\dots,n} = (Y_{i1}, \dots, Y_{iq})_{i=1,\dots,n}$ be n i.i.d observations where $\forall(i, j), Y_{ij} \in \{0, 1\}$. i and j are respectively the individual and the response indices. Each \mathbf{Y}_i is assumed to arise from the following mixture model:

$$\mathbb{P}(\mathbf{Y}_i = (y_{i1}, \dots, y_{iq})) = \sum_{k=1}^g \pi_k \prod_{j=1}^q \gamma_{kj}^{y_{ij}} (1 - \gamma_{kj})^{1-y_{ij}}$$

where π_k represents the proportion of the k -th component ($\sum_{k=1}^g \pi_k = 1$) and γ_{kj} is the success probability for the j -th response in the k -th group. The model is equivalently written as

$$\begin{aligned} Y_{ij}|Z_i &\sim \mathcal{B}(\gamma_{Z_i j}), & \forall i = 1, \dots, n, j = 1, \dots, q \\ \mathbb{P}(Z_i = k) &= \pi_k, & \forall i = 1, \dots, n, k = 1, \dots, g \end{aligned} \quad (21)$$

where $\mathbf{Z} = (Z_1, \dots, Z_n)$ is a latent random vector.

We set the following standard exchangeable prior distributions on $\boldsymbol{\pi} = (\pi_1, \dots, \pi_g)$ and $\boldsymbol{\gamma} = (\gamma_{kj})_{k=1,\dots,g,j=1,\dots,q}$:

$$\begin{aligned} (\pi_1, \dots, \pi_g) &\sim \text{Dir}(d, \dots, d) \\ (\gamma_{kj})_{k=1,\dots,g,j=1,\dots,q} &\sim \text{i.i.d Beta}(a, b) \end{aligned} \quad (22)$$

3.2.2 Posterior distribution and label switching

As for any mixture model, the posterior distribution should reproduce the likelihood invariance under permutation of the mixture indices. In other words, the posterior distribution is multi-modal, each mode corresponding to a permutation of the index of the mixture components. In such cases, it is well documented that MCMC algorithms often fail into exploring the various modes of the posterior distribution. Note that the label switching issue arises not only when it comes out to sample the posterior distribution but also for evidence approximation in model selection (see for instance the introduction of Lee and Robert, 2016, and references inside). In this section, we illustrate the fact that a simple solution –based on the Variational Bayes posterior approximation and our sampling algorithm SBS– can be proposed to handle the label switching problem.

On this model, it is easy to derive a mean-field variational approximation of the posterior distribution, resulting into a posterior approximation of the form:

$$\begin{aligned} \tilde{p}_{\mathbf{Y}}^{VB}(\mathbf{Z}, \boldsymbol{\gamma}, \boldsymbol{\pi}) &= f_{\mathcal{D}ir(\tilde{\delta}_1, \dots, \tilde{\delta}_g)}(\boldsymbol{\pi}) \\ &\times \prod_{k=1, j=1}^{g, q} f_{\mathcal{B}eta(\tilde{\alpha}_{kj}, \tilde{\beta}_{kj})}(\gamma_{kj}) \prod_{i=1}^n \prod_{k=1}^q (\tau_{ik})^{1_{Z_i=k}} \end{aligned} \quad (23)$$

Details can be found in White and Murphy (2014) and the algorithm is implemented in the corresponding R-package `BayesLCA`. However, contrary to the true posterior distribution ($p(\cdot|\mathbf{Y})$), $\tilde{p}_{\mathbf{Y}}^{VB}$ is not exchangeable. Moreover, this posterior approximation is known to be excessively concentrated around one mode. As a consequence, we can presume (and we illustrate it on the following numerical experiments) that our sampling algorithm SBS starting from $\tilde{p}_{\mathbf{Y}}$ and using a standard Gibbs transition kernel will not be able to propagate particles on the other modes of the posterior distribution.

Remark 7 *When talking about a “standard Gibbs transition kernel”, we refer to the most naive Gibbs algorithm, sequentially sampling $[\mathbf{Z}|\mathbf{Y}, \boldsymbol{\gamma}, \boldsymbol{\pi}]$, $[\boldsymbol{\gamma}|\mathbf{Y}, \mathbf{Z}, \boldsymbol{\pi}]$ and $[\boldsymbol{\pi}|\mathbf{Z}]$. Using the expressions for $\tilde{p}_{\mathbf{Y}}^{VB}$ and $p_{\rho_n}(\mathbf{Z}, \boldsymbol{\gamma}, \boldsymbol{\pi})$, these three distributions are conjugate. We stick to this MCMC kernel, and do not introduce any modification to force or prevent the label switching phenomena at the propagation step.*

$\tilde{p}_{\mathbf{Y}}^{VB}$ being unsatisfying from the label switching perspective, we introduce its so called *symmetrized version*, forcing the invariance by permutation. More precisely, let \mathcal{S}_g be the set of all the permutations of $\{1, \dots, g\}$, we define

$$\tilde{p}_{\mathbf{Y}}^{VB.Sym}(\mathbf{Z}, \boldsymbol{\gamma}, \boldsymbol{\pi}) = \sum_{\sigma \in \mathcal{S}_g} \tilde{p}_{\mathbf{Y}}^{VB.Sym}(\mathbf{Z}, \boldsymbol{\gamma}, \boldsymbol{\pi}, \sigma) \quad (24)$$

where

$$\begin{aligned} \tilde{p}_{\mathbf{Y}}^{VB.Sym}(\mathbf{Z}, \boldsymbol{\gamma}, \boldsymbol{\pi}, \sigma) &= \frac{1}{g!} f_{\mathcal{D}ir(\tilde{\delta}_{\sigma(1)}, \dots, \tilde{\delta}_{\sigma(g)})}(\boldsymbol{\pi}) \\ &\times \prod_{k=1}^g \prod_{j=1}^q f_{\mathcal{B}eta(\tilde{\alpha}_{\sigma(k)j}, \tilde{\beta}_{\sigma(k)j})}(\gamma_{kj}) \prod_{i=1}^n \prod_{k=1}^q (\tau_{i\sigma(k)})^{1_{Z_i=k}}. \end{aligned} \quad (25)$$

In order to keep the conjugacy properties in conditional distributions we use our algorithm SBS to sequentially sample:

$$p_{\rho_h}(\mathbf{Z}, \gamma, \boldsymbol{\pi}, \sigma) \propto [\tilde{p}_{\mathbf{Y}}^{VB.Sym}(\mathbf{Z}, \gamma, \boldsymbol{\pi}, \sigma)]^{1-\rho_h} \times [\ell(\mathbf{Y}|\mathbf{Z}, \gamma, \boldsymbol{\pi})p(\mathbf{Z}|\boldsymbol{\pi})\pi(\boldsymbol{\pi}, \gamma)]^{\rho_h} \quad (26)$$

Once σ has been integrated out, our sequential sampling scheme starts at $(\tilde{p}_{\mathbf{Y}}^{VB.Sym})$ given in equation (24) and terminates at the true posterior distribution.

3.2.3 Simulation design and test

This example aims at proving that our strategy supplies a sample from the true posterior distribution, even in the difficult framework of mixture models. To assess the validity of our method, we propose using the following testing procedure.

Testing procedure. We introduce the following property from which we derive a negative criterion, in the sense that if the obtained distribution is the true posterior distribution, it must satisfy the criterion.

Property 1 Let $\Phi : \Theta \mapsto \mathbb{R}$ be such that $\exists \Psi$ verifying $H : \theta \mapsto (\Phi(\theta), \Psi(\theta))$ is injective and continuously differentiable. Assume that

$$(\theta^*, \mathbf{Y}^*) \sim \pi(\theta^*)\ell(\mathbf{Y}^*|\theta^*).$$

Given \mathbf{Y}^* , let $p_{\mathbf{Y}^*}(\Phi(\theta))$ be any probability distribution on $\Phi(\Theta)$. Let $U(\theta^*, \mathbf{Y}^*, \Phi, p_{\mathbf{Y}^*})$ be the following statistic:

$$U(\theta^*, \mathbf{Y}^*, \Phi, p_{\mathbf{Y}^*}) = \mathbb{E}_{p_{\mathbf{Y}^*}} [1_{\Phi(\theta) < \Phi(\theta^*)}] . \quad (27)$$

Then, if $p_{\mathbf{Y}^*}(\Phi(\theta)) = p(\Phi(\theta)|\mathbf{Y}^*)$, $\forall \theta \in \Theta$ then $U(\theta^*, \mathbf{Y}^*, \Phi, p_{\mathbf{Y}^*}) \sim \mathcal{U}_{[0,1]}$.

Remark 8 Note that when $U(\theta^*, \mathbf{Y}^*, \Phi, p_{\mathbf{Y}^*})$ has no explicit expression and when we have access to a sample from $q_{\mathbf{Y}^*}$, we can replace $U(\theta^*, \mathbf{Y}^*, \Phi, p_{\mathbf{Y}^*})$ by its non-biased and convergent estimator:

$$U_M(\theta^*, \mathbf{Y}^*, \Phi, (\theta_m)_{m=1\dots M}) = \frac{1}{M} \sum_{m=1}^M 1_{\Phi(\theta_m) < \Phi(\theta^*)} \quad (28)$$

where $\theta_m \sim_{i.i.d.} p_{\mathbf{Y}^*}$ and moreover if $p_{\mathbf{Y}^*}(\Phi(\theta)) = p(\Phi(\theta)|\mathbf{Y}^*)$ for all $\theta \in \Theta$, then

$$U_M(\theta^*, \mathbf{Y}^*, \Phi, (\theta_m)_{m=1\dots M}) \sim \mathcal{U}_{\{0, \frac{1}{M}, \frac{2}{M}, \dots, 1\}}.$$

This property enables the elaboration of our checking procedure in 4 steps.

Checking procedure for posterior approximation

For a given approximation method \mathcal{M} of the posterior distribution. Case (a) corresponds to deterministic approximations and case (b) to stochastic approximations.

1. Generate S parameters and datasets $(\theta^{*s}, \mathbf{Y}^{*s})_{s=1, \dots, S}$ according to the Bayesian model $\pi(\theta^{*s})\ell(\mathbf{Y}^{*s}|\theta^{*s})$.
2. For each $s = 1 \dots S$, from dataset \mathbf{Y}^{*s} ,
 - (a) derive the deterministic approximation of the posterior $p_{\mathbf{Y}^{*s}}$ using \mathcal{M} ,
 - (b) get a sample $(\theta_m^s)_{m=1 \dots M}$ from \mathcal{M} .
3. Choose a real-valued function of the parameter $\Phi(\theta)$ and compute
 - (a) $U(\theta^{*s}, \mathbf{Y}^{*s}, \Phi, p_{\mathbf{Y}^{*s}})$ using (27),
 - (b) $U_M(\theta^{*s}, \mathbf{Y}^{*s}, \Phi, (\theta_m^s)_{m=1 \dots M})$ using (28).
4. Compare the empirical distribution
 - (a) of $U(\theta^{*s}, \mathbf{Y}^{*s}, \Phi, p_{\mathbf{Y}^{*s}})_{s=1 \dots S}$ to the uniform distribution on $[0, 1]$,
 - (b) of $(U_M(\theta^{*s}, \mathbf{Y}^{*s}, \Phi, (\theta_m^s)_{m=1 \dots M}))_{s=1 \dots S}$ to the uniform distribution on $\{1, \dots, \frac{1}{M}\}$.

The comparison can be performed through graphical tools – for instance the empirical Cumulative Distribution Function (cdf) – or through a statistical test such as the discrete goodness of fit test, discrete version of the Kolmogorov-Smirnov (KS) test. If we reject the hypothesis

$$H_0 = \left\{ U_M(\theta^*, \mathbf{Y}^*, \Phi, (\theta_m)_{m=1 \dots M}) \sim \mathcal{U}_{\{0, \frac{1}{M}, \frac{2}{M}, \dots, 1\}} \right\},$$

we conclude that the obtained sample is not distributed from the true posterior distribution.

In our case, several methods \mathcal{M} will be considered: the variational Bayes approximation (VB), the SBS starting from VB and SBS starting from a symmetrized version of VB.

Simulation design. We set $n = 100$ individuals, $p = 10$ observations by individual and $g = 2$ groups in the mixture. Referring to equation (22), the hyper-parameters are set to $\delta = a = b = 2$. $S = 500$ datasets \mathbf{Y}^{*s} are simulated. For each simulated dataset, we run the VB algorithm implemented in the BayesLCA package (White and Murphy, 2014). The sampling algorithms SBS starting from $\tilde{p}_{\mathbf{Y}^{*s}}^{VB}$ and $\tilde{p}_{\mathbf{Y}^{*s}}^{VB.Sym}$ respectively are implemented with $M = 5000$, $\tau_1 = 0.9$ and $\tau_2 = 0.9$. The K_h kernel is standard Gibbs algorithm of length $B = 5$.

3.2.4 Results

On Figure 3 (left panel), we plot the estimated posterior density (for a arbitrarily chosen dataset) obtained from the Variational approximation, its symmetrized version, the SBS starting from with VB and the SBS stating from with VB symmetrized. As expected, the variational approximation underestimates the posterior variance. SBS succeeds in inflating this posterior variance. This phenomenon can be observed on $\Phi_1(\theta) = |\pi_1 - \pi_2|$ (left) and $\Phi_2(\theta) = \pi_1$. About π_1 , we observe that using only the classical variational Bayes estimator –highly concentrated on the MAP– we are unable to explore the other modes. However, if forcing the exploration of the different modes, we are able to observe the full posterior distribution, charging all the modes.

On Figure 3 (right panel), we plot the empirical cdf's (ecdf) for U_M . The results for Φ_1 (Φ_2 respectively) are on the left (right respectively).

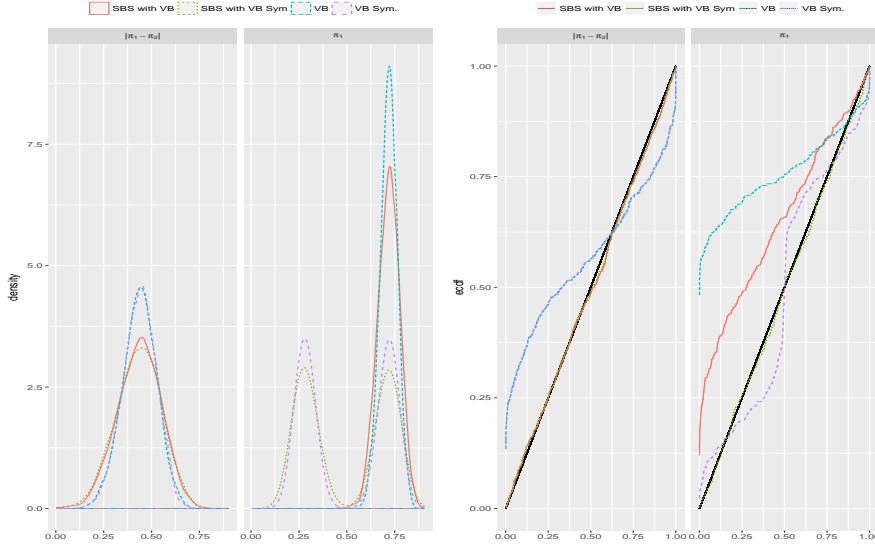


Fig. 3 Simulation results for the LCA model. Left panel: posterior distributions for $|\pi_1 - \pi_2|$ (left) and π_1 (right). Right panel: cumulative distribution function for $U_M(\theta^{*s}, \mathbf{Y}^{*s}, \Phi_r, q_{\mathbf{Y}^{*s}})$. The results for $\tilde{p}_{\mathbf{Y}^{*s}}$ and $\tilde{p}_{\mathbf{Y}^{*s}}^{Sym}$ are in blue and purple respectively. The results obtained by the SBS algorithm starting from $\tilde{p}_{\mathbf{Y}^{*s}}$ and $\tilde{p}_{\mathbf{Y}^{*s}}^{Sym}$ are in red and green respectively.

The ecdf for $(U(\theta^{*s}, \mathbf{Y}^{*s}, \Phi_l, \tilde{p}_{\mathbf{Y}^{*s}}^{VB}))_{s=1 \dots S}$ is in blue and the ecdf for $(U(\theta^{*s}, \mathbf{Y}^{*s}, \Phi_l, \tilde{p}_{\mathbf{Y}^{*s}}^{VBSym}))_{s=1 \dots S}$ is in purple. The ecdf obtained from $(\theta_m^s)_{m=1 \dots M}$ and $(\theta_m^{s, Sym})_{m=1 \dots M}$ sampled by algorithm SBS starting from $\tilde{p}_{\mathbf{Y}^{*s}}^{VB}$ (respectively $\tilde{p}_{\mathbf{Y}^{*s}}^{VB.Sym}$) are plotted in red (green respectively).

The phenomena observed before on a single dataset are confirmed here on the 500 datasets. On $|\pi_1 - \pi_2|$, which is insensitive to label switching, the non-symmetrized and symmetrized versions of the algorithms give equivalent results (curves red/green, and blue/purple can not be distinguished), which is not the case on π_1 (which is actually non-identifiable because of label switching). On $|\pi_1 - \pi_2|$, $\tilde{p}_{\mathbf{Y}^{*s}}^{VB}$ and $\tilde{p}_{\mathbf{Y}^{*s}}^{VB.Sym}$ are different from the true posterior distribution but our two algorithms starting respectively from $\tilde{p}_{\mathbf{Y}^{*s}}^{VB}$ and $\tilde{p}_{\mathbf{Y}^{*s}}^{VB.Sym}$ supply a sample from the true posterior distribution.

On π_1 , we observe that starting SBS from $\tilde{p}_{\mathbf{Y}^{*s}}^{VB}$ clearly leads to the wrong posterior. The equality of the true posterior distribution with the one obtained via SBS starting from VB.Sym is not rejected.

3.3 Stochastic block models with covariates

Model. As a last example, we consider the combination of stochastic block-model (Nowicki and Snijders, 2001, SBM) and logistic regression (shortened as ‘SBM-reg’ in the sequel) considered in Latouche et al (2015). This model aims at deciphering

	$ \pi_1 - \pi_2 $	π_1
VB	$4.497e^{-06}$	$< 2.2e^{-16}$
Symmetrized VB	$1.563e^{-05}$	$1.431e^{-09}$
SBS with VB	0.596	$< 2.2e^{-16}$
SBS with Symmetrized VB	0.567	0.903

Table 1 Simulation results for the LCA model: p-values for the KS test of Property 1.

some residual structure in an observed network once accounted for the effect of some edge covariates. The model is as follows. Consider a set of n nodes; for each pair ($1 \leq i < j \leq n$) of nodes, we observe a p -dimensional covariates vector \mathbf{x}_{ij} . Likewise in SBM, we further assume that each node belongs to one among g groups and we denote Z_i the (unobserved) group where node i is affected; $\boldsymbol{\pi} = (\pi_k)_k$ denotes the vector of group proportions. The model states that the edges of the observed binary undirected network $\mathbf{Y} = (Y_{ij})$ are drawn independently conditionally on the set of latent variables $\mathbf{Z} = (Z_i)$ as Bernoulli variables:

$$(Y_{ij}|Z_i, Z_j, \boldsymbol{\alpha}, \boldsymbol{\beta}) \sim \mathcal{B}(p_{ij}), \quad \text{logit}(p_{ij}) = \mathbf{x}_{ij}^\top \boldsymbol{\beta} + \alpha_{Z_i, Z_j}$$

where $\boldsymbol{\alpha} = (\alpha_{kl})$ stands for the matrix of between-group effects (analogous to the between-group connection probabilities from SBM, in logit scale) and $\boldsymbol{\beta} = (\beta_\ell)_{\ell=1, \dots, p}$ for the vector of regression coefficients. As for the priors, $\boldsymbol{\pi}$ has a Dirichlet distribution, both $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are Gaussian. When considering model selection or averaging, the number of groups g is supposed to be uniformly distributed among $\{1, \dots, g_{\max}\}$.

Bayesian model averaging (BMA). BMA (Hoeting et al, 1999) is a general principle, which consists in combining the results obtained with several models, rather than to choose the 'best' one. Among other interests, it allows to account for model uncertainty. We apply this principle to a regression parameter β_ℓ . While model selection consists in choosing g as $\hat{g} = \arg \max_g p(g|\mathbf{Y})$ and considering the posterior $p(\beta_\ell|\mathbf{Y}, g = \hat{g})$, BMA directly considers the unconditional posterior

$$p(\beta_\ell|\mathbf{Y}) = \sum_g p(g|\mathbf{Y}) p(\beta_\ell|\mathbf{Y}, g).$$

In terms of moments, it results in $\mathbb{E}(\beta_\ell|\mathbf{Y}) = \sum_g p(g|\mathbf{Y}) \mathbb{E}(\beta_\ell|\mathbf{Y}, g)$ and $\mathbb{V}(\beta_\ell|\mathbf{Y}) = \mathbb{V}_{\text{within}}(\beta_\ell|\mathbf{Y}) + \mathbb{V}_{\text{between}}(\beta_\ell|\mathbf{Y})$ where $\mathbb{V}_{\text{within}}$ measures the mean variance of the parameter conditionally on g and $\mathbb{V}_{\text{between}}$ is the variance of the parameter due to model uncertainty:

$$\begin{aligned} \mathbb{V}_{\text{within}}(\beta_\ell|\mathbf{Y}) &= \sum_g p(g|\mathbf{Y}) \mathbb{V}(\beta_\ell|\mathbf{Y}, g), \\ \mathbb{V}_{\text{between}}(\beta_\ell|\mathbf{Y}) &= \sum_g p(g|\mathbf{Y}) (\mathbb{E}(\beta_\ell|\mathbf{Y}, g) - \mathbb{E}(\beta_\ell|\mathbf{Y}))^2. \end{aligned}$$

VB approximation. As illustrated in Subsection 3.3, the VB approximate posterior is quite accurate for logistic regression and Gazal et al (2012) also proved its empirical accuracy for SBM. A first goal of this simulation study is to check if this accuracy still holds when the two models are combined into the SBM-reg model. To this aim, we focus on the posterior distribution of the regression parameters. Secondly, we want to check the accuracy of the VB posterior distribution of the number of groups, that can be used either to assess goodness-of-fit or for model averaging (Latouche et al (2015))

Simulation design. We simulate networks with $n \in \{20, 50\}$ nodes according to an SBM-reg model with $g_* \in \{1, 2\}$ groups and $p = 3$ covariates. To apply Property 1, the parameters are sampled from the prior distribution. $S = 100$ replicates are simulated for each configuration and for each of them, the SBM-reg models with $g \in \{1, \dots, g_{\max} = 5\}$ were fitted with the VB algorithm described in Latouche et al (2015). The SBS algorithm is then run on each dataset.

Results for parameter estimation. We first consider the posterior distribution of β when the number of groups g is known. On Figure 4, we plot on the left the boxplots for the posterior means $(\widehat{\mathbb{E}}^{VB}(\beta_\ell | \mathbf{Y}^s))_{s=1 \dots 100}$ and $(\widehat{\mathbb{E}}^{SBS}(\beta_\ell | \mathbf{Y}^s))_{s=1 \dots 100}$. The boxplot (over the 100 simulated datasets) of the posterior standard deviations $(\widehat{\sigma}^{VB}(\beta_\ell | \mathbf{Y}^s))_{s=1 \dots 100}$ and $(\widehat{\sigma}^{SBS}(\beta_\ell | \mathbf{Y}^s))_{s=1 \dots 100}$ are on the top-right. We clearly observe that the posterior means provided by VB and SBS are both accurate and similar, but the VB's posterior standard deviations (sd) are smaller than SBS's posterior standard deviations.

To further assess the quality of the posterior distribution provided by VB and SBS, we checked Property 1 for the regression coefficients, which are not subject to label switching. We observe that the ecdf of the SBS sample is close to uniform, whereas there is a departure for VB. The p-values resulting from KS test of Property 1 (see Table 2, upper table) lead to the same conclusions. All these observations concur to show that, although the VB approximate posterior distribution is accurate for logistic regression and SBM separately, it is biased for the SBM-reg model, and that the proposed SBS is a way to correct it. As a consequence of this phenomenon, the empirical level of VB's credibility intervals is equal to 84.75%, which is below the nominal level 95%, whereas SBS's credibility intervals almost reach the targeted level (93.75%).

$\widehat{g} = g_*$				
	$n = 20$		$n = 50$	
	$g = 1$	$g = 2$	$g = 1$	$g = 2$
VB	0.027	0.004	0.002	0.077
SBS	0.785	0.121	0.839	0.238

BMA				
	$n = 20$		$n = 50$	
	$g = 1$	$g = 2$	$g = 1$	$g = 2$
VB	0.017	0.003	0.002	0.079
SBS	0.740	0.277	0.778	0.312

Table 2 Simulation results for the SBM-reg model: p-values for the KS test of Property 1.

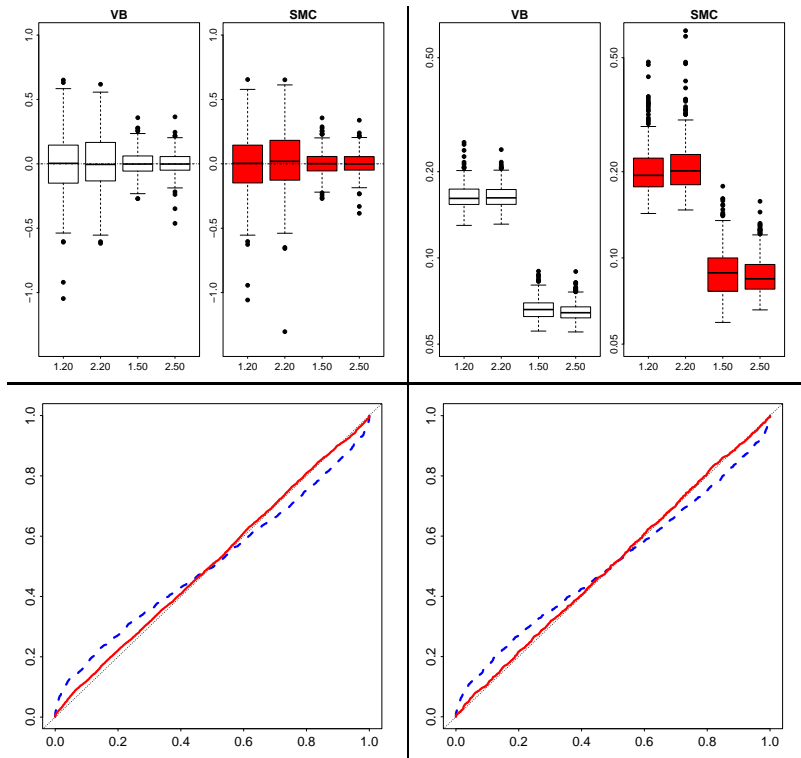


Fig. 4 Simulation results for the SBM-reg model: VB (white) and SBS (red) posterior of the regression coefficients $\beta = (\beta_\ell)$. Top: posterior mean (left), posterior standard deviation (right); x-axis label: $g_* \cdot n$ (e.g. '1.20' means $g_* = 1$, $n = 20$). Bottom: graphical check of Property 1 for VB (dashed blue) and for SBS (solid red). Left: $g = g_*$, right: with model averaging.

Results for model selection. We now consider the posterior distribution of the number of groups $p(g|\mathbf{Y})$ and its use for model selection. Figure 5 provides a comparison of the posterior provided by VB and SBS. We observe that the VB approximation always results in a more concentrated distribution than SBS. This behavior can be compared to the under-estimation of the posterior variance of the parameters that we already discussed. To compare the results in terms of model selection we computed the frequency at which the right model is selected (i.e. when $\hat{g} = g_*$) and the mean posterior probability of the g_* (see Table 3). We observe that VB performs better than SBS for both criteria. This parallels Minka (2005), who shows that the minimization of the Kullback-Leibler (KL) divergence leads to an accurate estimate of the mode, which is convenient for model selection. Although it does not seem to hamper model selection, the biased estimation of the posterior $p(g|\mathbf{Y})$ may have undesired consequences when used for model averaging. To illustrate this point, we simply computed the empirical coverage of credibility intervals for each β_ℓ after model averaging. The mean coverage across simulation condition and covariate index ℓ for VB (85.8%) is still below the nominal level,

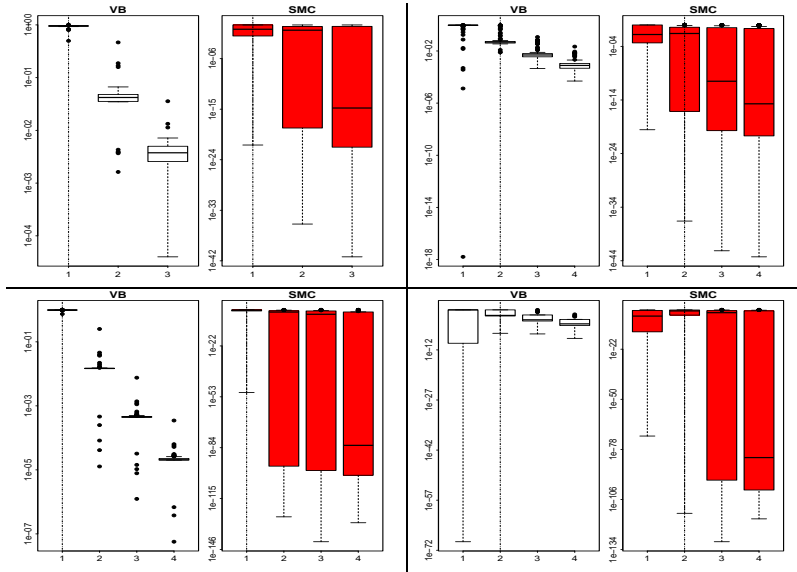


Fig. 5 Simulation results for the SBM-reg model: box-plots for the posterior probability $p(g|\mathbf{Y})$ as a function of g . Top $n = 20$, bottom: $n = 50$. Left: $g_* = 1$, right: $g_* = 2$.

	frequency of $\hat{g} = g_*$ (%)			
	$n = 20$		$n = 50$	
	$g_* = 1$	$g_* = 2$	$g_* = 1$	$g_* = 2$
VB	100	10	100	42
SBS	46	23	60	36

	mean value of $p(g = g_* \mathbf{Y})$			
	$n = 20$		$n = 50$	
	$g_* = 1$	$g_* = 2$	$g_* = 1$	$g_* = 2$
VB	0.947	0.138	0.982	0.410
SBS	0.435	0.257	0.562	0.387

Table 3 Simulation results for the SBM-reg model: model selection

whereas this of SBS (93.25%) is close to 95%. Figure 4 (bottom right) also shows that the distribution of the ecdf after model averaging is almost confounded with the uniform for SBS, whereas it still displays a significant bias for VB. The p-values for the KS test of Property 1 (see Table 2, bottom) lead to the same conclusion.

4 Illustrations on network datasets

Network analysis with SBM-reg. To illustrate the use of the proposed sampling algorithm, studied a series of examples analyzed by Latouche et al (2015) with an SBM-reg model. The main two purposes of such an analysis is (i) to estimate the effect β of the covariates and (ii) to assess the goodness-of-fit of the model based on the covariates. Task (ii) is achieved by computing the posterior probability for

the SBM part of the model to involve only $g = 1$ class, that is $p(g = 1|\mathbf{Y})$. A low value of this probability is an indication for a residual structure in the network.

We refer to Latouche et al (2015) for the presentation of the data. We considered the datasets (networks) referred to as Florentine (business), Florentine (marriage), Karate, Tree and Blog. Their respective sizes range from few tens to few hundreds nodes and their densities from 1% to 50%. Note that the numerical results presented here for the VB inference slightly differ, as we kept all nodes from each graph whereas Latouche et al (2015) removed all isolated nodes.

For each of these datasets, we fitted an SBM-reg model with $g = 1 \dots g_{\max}$ groups with a VBEM algorithm to obtain the Gaussian VB approximate distribution for $\tilde{p}_{\mathbf{Y}}^{VB}(\boldsymbol{\beta})$. We also run the proposed SBS with $M = 1000$ particles and obtained a weighted sample from $\hat{p}_{\mathbf{Y}}^{SBS}(\boldsymbol{\beta})$. To compare the posterior distributions of $\boldsymbol{\beta}$, we adopted the Bayesian model averaging principle described in Section 3.3. For each dataset s and each covariate ℓ , we first computed the ratio between the VB and SBS posterior standard deviation (sd) $\sqrt{\mathbb{V}^{VB}(\beta_j)/\mathbb{V}^{SBS}(\beta_j)}$. For each configuration, we also computed the ratios $\tilde{\beta}_j^{VB}/\sqrt{\tilde{\Sigma}_{jj}^{VB}}$ and $\tilde{\beta}_j^{SBS}/\sqrt{\tilde{\Sigma}_{jj}^{SBS}}$, which are typically used to evaluate the effect of the covariates. Figure 6, left, shows that the VB approximation tends to under estimate the posterior variance of the parameter. As expected, Figure 6, right, shows that it yields in over-estimating the significance of the effect of the covariates.




Fig. 6 Posterior moments of the regression coefficients for the datasets from Latouche et al (2015). Left: ratio between the VB and SBS sd as a function of the network size (color = network). Right: ratio between the posterior mean and the posterior sd (x : VB, y : SBS, signed-log scale, color: network, horizontal and vertical lines: 2.5% and 97.5% $\mathcal{N}(0, 1)$ quantiles)

To further illustrate these differences, we studied the posterior distribution of the three regression coefficients used in the 'Tree' example. The regression coefficient are associated with the genetic, geographic and taxonomic distance between the tree species. The results given in Table 4 indicate again that the posterior sd provided by SBS are all larger than these resulting from VB. We observe that, because of the high concentration of the posterior distribution of g , VB strongly under-estimates the variance due to model uncertainty. A result is that the SBS posterior standard deviations are almost twice larger than the VB posterior stan-

standard deviations. As a consequence, the significance ratio is reduced, and the influence of both the genetic and the geographic distance (β_1 and β_2) turn out to be more questionable according to SBS than to VB.

	$\tilde{p}_{\mathbf{Y}}^{VB}$		
	β_1	β_2	β_3
post. mean	4.62e-05	0.23	-0.9
post. within var.	2.24e-10	0.0432	0.00175
post. between var.	5.55e-17	1.18e-06	2.42e-07
posterior sd	1.5e-05	0.208	0.0418
ratio	3.09	1.11	-21.5

	$\hat{p}_{\mathbf{Y}}^{SBS}$		
	β_1	β_2	β_3
post. mean	4.13e-05	0.355	-0.906
post. within var.	1.09e-09	0.219	0.00889
post. between var.	3.99e-12	0.0019	0.00281
posterior sd	3.31e-05	0.47	0.108
ratio	1.25	0.755	-8.38

Table 4 Posterior moments of the regression coefficients.

For the goodness-of-fit study, we compared the values of $\tilde{p}_{\mathbf{Y}}^{VB}(1)$ and $\hat{p}_{\mathbf{Y}}^{SBS}(1)$ (Table 5). Except in the most uncertain case (Karate), the posterior probabilities are similar and lead to the same conclusion about the existence of a residual structure in the network.

	$\tilde{p}_{\mathbf{Y}}^{VB}(1)$	$\hat{p}_{\mathbf{Y}}^{SBS}(1)$
Marriage	$9.54 \cdot 10^{-1}$	1.00
Business	$7.04 \cdot 10^{-1}$	1.00
Karate	$2.56 \cdot 10^{-1}$	$7.07 \cdot 10^{-3}$
Tree	$4.83 \cdot 10^{-153}$	$1.06 \cdot 10^{-161}$
Blog	$8.63 \cdot 10^{-174}$	$4.04 \cdot 10^{-290}$

Table 5 Posterior probability for the SBM-reg model with only one class.

5 Discussion and perspectives

In this paper, we presented a simple strategy to combine the strength of deterministic approximations of the posterior distribution with sequential Monte Carlo samplers. We illustrated the efficiency of our approach and its robustness with respect to the deterministic approximation on a large simulation study. Its application on network datasets stresses the fact that the well-known underestimation of the posterior variance by the variational approximation can be easily corrected, sometimes leading to different statistical conclusions. Besides, if dependencies between parameters have been neglected in the deterministic posterior approximation, they will be recovered by the sequential sampling.

Our approach is not restricted to the case where a standard deterministic posterior approximation can be derived (such as Variational Bayes, Laplace or Expectation Propagation estimate). Any point estimate can be used to design a rough posterior (using a Gaussian or a log-Gaussian seems to be the simplest solution) and serves as an accelerator of the sampling sequence. This strategy is different from an empirical Bayes strategy, the point estimate being only used to explore more efficiently the posterior distribution and not to elicit a prior distribution. The method is not as sensible as standard Importance Sampling to an eventual under-evaluation of the approximate posterior variance : even with a too narrow approximation of the posterior distribution, the algorithm is able to get back to the true posterior variance.

SMC directly supplies a final population of particles arising from the true posterior distribution, as opposed to MCMC strategies, whose convergence is difficult to assess. The proposed SBS algorithm is adaptive in the sense that the sequence $\tilde{p}_{\mathbf{Y}}(\theta)^{1-\rho_h}(p(\theta|\mathbf{Y}))^{\rho_h}$ is determined on the fly in an automatic way. Furthermore, the algorithm path (summarized by the sequence ρ_h) is an indicator of the quality of the deterministic posterior distribution used to initiate the bridge sampling.

A natural extension of the present work is its adaptation to Approximate Bayesian Computation (ABC) context for models with no explicit likelihood, following Del Moral et al (2012). The difficulty will arise from the specification of the distributions sequence.

Acknowledgements

The authors are thankful to Nicolas Chopin for fruitful discussions.

References

- Bartholomew D, Knott M, Moustaki I (2011) Basic Ideas and Examples, John Wiley & Sons, Ltd, pp 1–18. DOI 10.1002/9781119970583.ch1
- Blei DM, Kucukelbir A, McAuliffe JD (2016) Variational Inference: A Review for Statisticians. arXiv pp 1–33, DOI 10.1080/01621459.2017.1285773, 1601.00670
- Consonni G, Marin JM (2007) Mean-field variational approximate bayesian inference for latent variable models. *Comput Stat Data Anal* 52(2):790–798, DOI 10.1016/j.csda.2006.10.028
- Del Moral P, Doucet A, Jasra A (2006) Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 68(3):411–436, DOI 10.1111/j.1467-9868.2006.00553.x, 0212648v1
- Del Moral P, Doucet A, Jasra A (2012) An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statistics and Computing* 22(5):1009–1020, DOI 10.1007/s11222-011-9271-y
- Del Moral P, Doucet A, Jasra A (2012) On adaptive resampling strategies for sequential monte carlo methods. *Bernoulli* 18(1):252–278, DOI 10.3150/10-BEJ335
- Doucet A, Johansen AM (2009) A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later, Oxford University Press

- Doucet A, Smith A, de Freitas N, Gordon N (2001) *Sequential Monte Carlo Methods in Practice*. Information Science and Statistics, Springer New York
- Erosheva EA, Fienberg SE, Joutard C (2007) Describing disability through individual-level mixture models for multivariate binary data. *Ann Appl Stat* 1(2):502–537, DOI 10.1214/07-AOAS126
- Garrett ES, Zeger SL (2000) Latent class model diagnosis. *Biometrics* 56(4):1055–1067, DOI 10.1111/j.0006-341X.2000.01055.x
- Gazal S, Daudin JJ, Robin S (2012) Accuracy of variational estimates for random graph mixture models. *Journal of Statistical Computation and Simulation* 82(6):849–862, DOI 10.1080/00949655g.2011.560117
- Hoeting JA, Madigan D, Raftery AE, Volinsky CT (1999) Bayesian model averaging: A tutorial. *Statistical Science* 14(4):382–417
- Jasra A, Stephens DA, Doucet A, Tsagaris T (2011) Inference for lvy-driven stochastic volatility models via adaptive sequential monte carlo. *Scandinavian Journal of Statistics* 38(1):1–22, DOI 10.1111/j.1467-9469.2010.00723.x
- Latouche P, Robin S, Ouadah S (2015) Goodness of fit of logistic models for random graphs. Tech. rep., arXiv:1508.00286
- Lee JE, Robert CP (2016) Importance sampling schemes for evidence approximation in mixture models. *Bayesian Anal* 11(2):573–597, DOI 10.1214/15-BA970
- McGrory CA, Pettitt AN, Titterington DM, Alston CL, Kelly M (2016) Transdimensional sequential Monte Carlo using variational Bayes – SM-CVB. *Computational Statistics & Data Analysis* 93:246–254, DOI <http://dx.doi.org/10.1016/j.csda.2015.03.006>
- Minka T (2005) Divergence measures and message passing. Tech. Rep. MSR-TR-2005-173, Microsoft Research Ltd
- Minka TP (2001) Expectation propagation for approximate bayesian inference. In: *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, UAI '01, pp 362–369
- Naesseth CA, Linderman SW, Ranganath R, Blei DM (2017) Variational Sequential Monte Carlo. Tech. rep., 1705.11140
- Neal RM (2001) Annealed importance sampling. *Statistics and Computing* 11(2):125–139, DOI 10.1023/A:1008923215028
- Nowicki K, Snijders T (2001) Estimation and prediction for stochastic block-structures 96:1077–87
- Rabinovich M, Angelino E, Jordan MI (2015) Variational consensus Monte Carlo. ArXiv e-prints 1506.03074
- Robert C, Casella G (2009) *Introducing Monte Carlo Methods with R*. Springer Science & Business Media
- Robert CP, Casella G (2005) *Monte Carlo Statistical Methods* (Springer Texts in Statistics). Springer-Verlag New York, Inc., Secaucus, NJ, USA
- Rue H, Martino S, Chopin N (2009) Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71(2):319–392, DOI 10.1111/j.1467-9868.2008.00700.x
- Schäfer C, Chopin N (2013) Sequential Monte Carlo on large binary sampling spaces. *Statistics and Computing* 23(2):163–184, DOI 10.1007/s11222-011-9299-z

-
- Wainwright MJ, Jordan MI (2008) Graphical models, exponential families, and variational inference. *Found Trends Mach Learn* 1(1–2):1–305
- Wang B, Titterton D (2005) Inadequacy of interval estimates corresponding to variational bayesian approximations. In: *AISTATS*
- White A, Murphy T (2014) Bayeslca: An r package for bayesian latent class analysis. *Journal of Statistical Software, Articles* 61(13):1–28, DOI 10.18637/jss.v061.i13
- Zhou Y, Johansen AM, Aston JAD (2016) Towards Automatic Model Comparison: An Adaptive Sequential Monte Carlo Approach. *Journal of Computational and Graphical Statistics* 25(3):701–726