

A NEW MCMC ALGORITHM FOR BLIND BERNOULLI-GAUSSIAN DECONVOLUTION

Di Ge, Jérôme Idier, Eric Le Carpentier

IRCCyN (CNRS UMR 6597)

1, rue de la Noë, BP 92101, F44321 Nantes Cedex 3, France

(Di.Ge | Jerome.Idier | Eric.Le-Carpentier)@ircsyn.ec-nantes.fr

ABSTRACT

This paper proposes a new algorithm for Bernoulli-Gaussian (BG) blind deconvolution in the Markov chain Monte Carlo (MCMC) framework. To tackle such a problem, the classical Gibbs sampler is usually adopted, as proposed by Cheng *et al.* [1]. However, as already pointed out by Bourguignon and Carfantan [2], it fails to explore the state space efficiently. In principle, a more efficient exploration technique could be obtained by integrating the Gaussian amplitudes out of the target distribution. Unfortunately, some of the sampling steps then become intractable. Therefore, our solution mixes steps in which the amplitudes are integrated out with others where they are not. The *invariant condition* is shown to hold, and simulations indicate that it behaves much more satisfactorily than the reference Gibbs sampler.

Index Terms—Blind deconvolution, Bernoulli-Gaussian model, Markov chain Monte Carlo methods

1. INTRODUCTION

The problem of the restoration of a sparse spike train distorted by a linear system and noise arises in many fields such as seismic exploration [1, 3]. It is classically dealt with using a discrete-time noisy convolution model for the observed vector $\mathbf{z} = [z_1, \dots, z_N]^t$:

$$z_n = \sum_{k=0}^P h_k x_{n-k} + \varepsilon_n, \forall n = 1, \dots, N. \quad (1)$$

$\mathbf{h} = [h_0, \dots, h_P]^t$ denotes the impulse response (IR) of the system (assumed finite here). For the sake of simplicity, the “zero boundary” condition is imposed on the sparse spike train to be restored, such that $\mathbf{x} = [x_1, \dots, x_M]^t$, $M = N - P$. ε is a stationary white Gaussian noise. The deconvolution problem is said *blind* when \mathbf{h} is unknown.

In the present study, we adopt a BG model for the spike train \mathbf{x} , following [3] and many posterior contributions such as [1, 4]. A BG signal is defined in two stages involving a Bernoulli sequence $\mathbf{q} = [q_1, \dots, q_M]^t$, such that: $\forall m = 1, \dots, M$.

$$q_m \sim Bi(\lambda), \quad (x_m | q_m) \sim \mathcal{N}(0, q_m \sigma_x^2). \quad (2)$$

The sparse nature of the spikes is governed by the Bernoulli law, while amplitudes are assumed iid zero-mean Gaussian. The MCMC approach [5, 6] is a powerful numerical tool, appropriate to solve complex inference problems such as blind deconvolution. In the field of blind BG deconvolution, Cheng *et al.* pioneered the introduction of MCMC methods [1]. They proposed to rely on a Gibbs sampler, for which their

algorithm constitutes a simple and canonical example of the application. However, our simulation results indicate that it lacks reliability: from different initial conditions, significantly different estimations are obtained, even after considerable iterations. This conclusion agrees with that of Bourguignon *et al.* in the context of BG spectral analysis [2].

The recent contribution of [7] already identified a convergence issue linked to time-shift ambiguities, and proposed an efficient way to solve it. There exists another source of inefficiency, unrelated to the time-shift ambiguity: instead of exploring the 2^M configurations at an acceptable speed, the Gibbs sampler tends to get stuck in one particular configuration of \mathbf{q} or another, as shown in Sect. 2.2. In fact, in a multi-modal space a Markov chain equilibrates rapidly within a mode (a configuration of \mathbf{q}), but takes a long time to move from mode to mode. Guan [8] studied a comparable slow convergence phenomenon of a *local* chain in comparison with his *small world* MCMC approach.

The new Gibbs-type sampler proposed here explores the posterior distribution *marginally* to the amplitudes \mathbf{x} , such that it probes the Bernoulli sequence \mathbf{q} more freely in its space than Cheng *et al.*'s Gibbs sampler. However, a plain Gibbs sampler of the marginal posterior distribution involves hardly implementable sampling steps. In particular, it is all but simple to simulate \mathbf{h} conditional on \mathbf{z} and \mathbf{q} and marginally w.r.t. \mathbf{x} . Our scheme solves this problem by incorporating steps where the sampling of \mathbf{x} is still involved, while \mathbf{q} is sampled marginally w.r.t. \mathbf{x} . Therefore, it is a *Gibbs-type sampler*, though fully valid from the mathematical viewpoint, since the invariant condition is shown to hold. A comparable idea is found in statistical signal segmentation [9] where some hyperparameters are partially marginalized, though without mathematical justification. Finally, for the sake of conciseness, it will remain implicit that the time-shift ambiguities are dealt with according to [7] in Cheng *et al.*'s Gibbs sampler, as well as in the proposed sampler.

In Sect. 2, the blind BG deconvolution problem is formulated. The Gibbs sampler of the joint posterior distribution [1] is also presented, and an example illustrates its inefficiency. Our Gibbs-type sampler is introduced in Sect. 3 and a simulated example gives an insight into the way it escapes from suboptimal configurations. The invariant condition is shown to hold, and an adapted implementation inspired from [4] is proposed. Finally, simulation results are presented in Sect. 4

2. PROBLEM FORMULATION

2.1 Statistical model

Akin to [1], the following assumptions are made:

- $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma_\varepsilon^2 \mathbf{I})$ is independent of \boldsymbol{x} and \boldsymbol{h} ;
- \boldsymbol{x} is a BG process defined by (2) with $\sigma_x = 1$;
- $\boldsymbol{h} \sim \mathcal{N}(0, \sigma_h^2 \mathbf{I}_{P+1})$;

σ_x is arbitrarily set to one in order to remove the scale ambiguity inherent to the blind deconvolution problem. According to the Monte Carlo principle, a posterior mean estimator of $\Theta = \{\boldsymbol{q}, \boldsymbol{x}, \boldsymbol{h}, \lambda, \sigma_\varepsilon, \sigma_h\}$ given \boldsymbol{z} can be approximated by:

$$\hat{\Theta} = \frac{1}{I-J} \sum_{k=J+1}^I \Theta^{(k)}, \quad (3)$$

where the sum extends over the last $I - J$ samples. In the MCMC framework, the samples are generated recursively, so that the asymptotic distribution of $\Theta^{(k)}$ is the joint posterior distribution [6]:

$$\rho(\Theta | \boldsymbol{z}) \propto g(\boldsymbol{z} - \boldsymbol{x} \star \boldsymbol{h}; \sigma_\varepsilon^2 \mathbf{I}_N) g(\boldsymbol{x}; \sigma_x^2 \text{diag}\{\boldsymbol{q}\}) g(\boldsymbol{h}; \sigma_h^2 \mathbf{I}_{P+1}) p(\boldsymbol{q}; \lambda) p(\sigma_h) p(\lambda) p(\sigma_\varepsilon) \quad (4)$$

where $\text{diag}\{\boldsymbol{q}\}$ denotes a diagonal matrix whose diagonal is \boldsymbol{q} , $g(\cdot; \mathbf{R})$ the centered Gaussian pdf of covariance \mathbf{R} , and \star the convolution operator as defined in Eq. (1).

2.2 Classical Gibbs Sampling

Cheng *et al.*'s Gibbs sampler proceeds along Tab. 1.

<ol style="list-style-type: none"> ① Let $\boldsymbol{y} = (\boldsymbol{q}, \boldsymbol{x})$. For each $i = 1 \dots, M$, <ol style="list-style-type: none"> (a) draw $q_i^{(k+1)} \boldsymbol{y}_{1:i-1}^{(k+1)}, \boldsymbol{y}_{i+1:M}^{(k)}, \boldsymbol{h}^{(k)}, \sigma_\varepsilon^{(k)}, \lambda^{(k)}, \boldsymbol{z}$ (b) draw $x_i^{(k+1)} \boldsymbol{y}_{1:i-1}^{(k+1)}, \boldsymbol{y}_{i+1:M}^{(k)}, q_i^{(k+1)}, \boldsymbol{h}^{(k)}, \sigma_\varepsilon^{(k)}, \boldsymbol{z}$ ② draw $\boldsymbol{h}^{(k+1)} \boldsymbol{x}^{(k+1)}, \sigma_h^{(k)}, \sigma_\varepsilon^{(k)}, \boldsymbol{z}$ ③ draw $\sigma_\varepsilon^{(k+1)} \boldsymbol{x}^{(k+1)}, \boldsymbol{h}^{(k+1)}, \boldsymbol{z}$ ④ draw $\lambda^{(k+1)} \boldsymbol{q}^{(k+1)}$ ⑤ draw $\sigma_h^{(k+1)} \boldsymbol{h}^{(k+1)}$
--

Table 1. Cheng *et al.*'s Gibbs sampler (see [1] for implementation details).

The five simulation steps are iterated until convergence towards the posterior distribution, and $\hat{\Theta}$ is finally built according to (3). This scheme presents several drawbacks despite its simplicity. Labat and Idier [7] have already pointed out that time-shift ambiguities could lead to unreliable estimates, depending on the initialization of \boldsymbol{h} . Here, we rather focus on the fact that step 1 explores the state space of $(\boldsymbol{q}, \boldsymbol{x})$ with a low efficiency. More precisely, the correspondingly sampled chains hardly escape from local maximizers of the posterior likelihood, due to highly dependent consecutive samples. Fig. 1 illustrates this phenomenon, in which simulation data are generated from a single spike, convolved with the IR defined by

$$h_i = \cos\left((i-10)\frac{\pi}{4}\right) \exp\left(-|0.225i-2|^{1.5}\right), \quad i = 0, \dots, 20,$$

and plotted on Fig. 2(b). A suboptimal configuration is chosen as the initial state. The chain takes several hundreds of iterations to visit the optimal configuration for the first time.




Fig. 1. Sampled BG sequences obtained by Cheng *et al.*'s sampler on a simple example. From a suboptimal configuration chosen as initial state, the Markov chain spends several hundreds of iterations before visiting the solution, *i.e.*, a unique spike in position 10 (marked as a bullet).

As a consequence, Cheng *et al.*'s method tends to produce unreliable estimated values, as shown in Fig. 2. The same IR is adopted here while \boldsymbol{x} is the well-known Mendel's test sequence [3]. The data are corrupted by a Gaussian noise with $\sigma_\varepsilon^2 = 4 \times 10^{-6}$, corresponding to SNR = 12.81 dB. Non-informative conjugate prior laws on the parameters are adopted:

$$\sigma_\varepsilon^2 \sim IG(1, 1), \quad \sigma_h^2 \sim IG(1, 1), \quad \lambda \sim Be(1, 1).$$

Initialization of the Markov chain $\Theta^{(0)}$ is done according to:

$$\sigma_\varepsilon^{(0)} = 10^{-2}, \quad \boldsymbol{q}^{(0)} = \mathbf{0}, \quad \boldsymbol{h}^{(0)} = \frac{\|\boldsymbol{z}\|_1}{M} \delta(n-11), \quad \lambda^{(0)} = 0.1.$$

The three estimation results in Fig. 2 are obtained from the same simulated data \boldsymbol{z} , and the same initialization, the only difference being the random seed value. 1000 samples are produced for each Markov chain and the last 250 are averaged to compute the estimation. Substantial variations exist from one estimated result to the other, especially in the number and the positions of spikes $\hat{\boldsymbol{x}}$. Actually, it can be checked that each sequence $\{\boldsymbol{q}^{(k)}\}$ tends to become constant for very long period of time, without fully exploring the state space. It should be stressed that these results are typical, *i.e.*, not selected on purpose.

The inefficiency of Cheng *et al.*'s Gibbs sampling has already been noticed by Bourguignon *et al.* [2]. Their solution involves proposals to shift several adjacent BG components, very similar to the idea introduced by Chi and Mendel [10] in the context of deterministic posterior likelihood maximization. Our approach, however, aims to tackle the problem indirectly by marginalizing the amplitudes \boldsymbol{x} out of the target posterior distribution, and by so doing releases the Markov chain from local maxima of the likelihood.

3. TOWARDS A MORE EFFICIENT SAMPLER

3.1 Partial marginalization of \boldsymbol{x}

The motivation to exclude \boldsymbol{x} from the Gibbs sampler is twofold. Firstly, the determination of \boldsymbol{x} conditional on $(\tilde{\Theta}, \boldsymbol{z})$ with $\tilde{\Theta} = (\boldsymbol{q}, \boldsymbol{h}, \sigma_\varepsilon, \sigma_h, \lambda)$ is a linear estimation, easy to solve afterwards. Secondly, a Gibbs sampler with $\rho(\tilde{\Theta} | \boldsymbol{z})$ as its

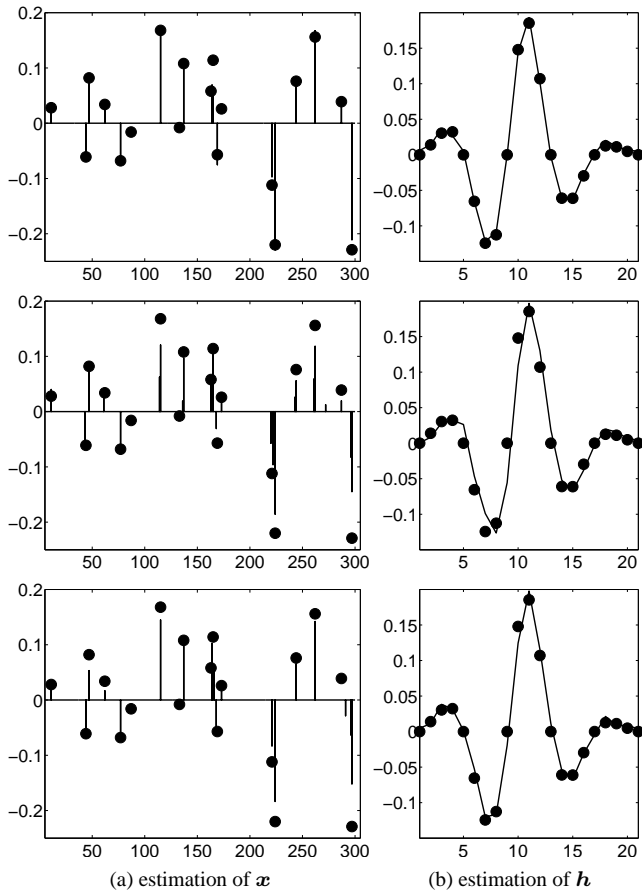


Fig. 2. Three different solutions of Cheng *et al.*'s method (under the improved version of [7]), obtained by changing the random seed value only. The actual values are marked as bullets.

target distribution is likely to be more efficient than Cheng *et al.*'s sampler, particularly w.r.t. to the Bernoulli sequence \mathbf{q} . Theoretical foundations are available in [11] and [5, Chapter 6.7] regarding the convergence rate of the so called *collapsed* Gibbs sampler, that converges to the marginalized distribution $\rho(\tilde{\Theta}|z) = \int \rho(\Theta|z) d\mathbf{x}$ by analytically integrating \mathbf{x} out of the scheme. According to [5, Chapter 6.7], the collapsed Gibbs sampler produces a substantial gain in terms of convergence rate if the integrated variables form highly dependent pairs with others in the Markov chain (see also [12] and references therein). This is exactly the case of the pair $\{\mathbf{x}, \mathbf{q}\}$ in the spike train deconvolution problem.

3.2 Proposed Gibbs-type sampler

Unfortunately, marginalizing \mathbf{x} leads to practical difficulties. In particular, the conditional sampling of \mathbf{h} given $\tilde{\Theta} \setminus \mathbf{h}$ becomes extremely difficult when \mathbf{x} is integrated out, since $\rho(\mathbf{h}|\mathbf{q}, \sigma_h, \sigma_\varepsilon, z)$ is a multivariate, non Gaussian law with a complex structure. Instead of a plain Gibbs sampler on $\tilde{\Theta}$, the sampling scheme of Tab. 2 is proposed to circumvent the direct conditional sampling of $\mathbf{h}, \sigma_\varepsilon \in \tilde{\Theta}$. The four steps are iterated until convergence towards the posterior distribution $\rho(\tilde{\Theta}|z)$. Compared with the classical sampler, the main difference appears in step 1, where \mathbf{x} has been analytically

- | |
|---|
| <ol style="list-style-type: none"> ① for each $i = 1 \dots, M$,
draw $q_i^{(k+1)} q_{1:i-1}^{(k+1)}, q_{i+1:M}^{(k)}, \mathbf{h}^{(k)}, \sigma_\varepsilon^{(k)}, \lambda^{(k)}, z$ ② (a) draw $\mathbf{x} \mathbf{q}^{(k+1)}, \mathbf{h}^{(k)}, \sigma_\varepsilon^{(k)}, z$
(b) draw $\mathbf{h}^{(k+1)} \mathbf{x}, \sigma_h^{(k)}, \sigma_\varepsilon^{(k)}, z$
(c) draw $\sigma_\varepsilon^{(k+1)} \mathbf{x}, \mathbf{h}^{(k+1)}, z$ ③ draw $\lambda^{(k+1)} \mathbf{q}^{(k+1)}$ ④ draw $\sigma_h^{(k+1)} \mathbf{h}^{(k+1)}$ |
|---|

Table 2. The proposed Gibbs-type sampler.

integrated out. Subsect. 3.3 provides an efficient way to implement step 1. On the other hand, in step 2, \mathbf{x} only plays the role of an auxiliary quantity, since it does not belong to the Markov chain anymore. Steps 2(b), (c) are structurally identical to steps 2, 3 in Tab. 1. In order to justify that the sampling scheme of Tab. 2 is mathematically valid, we let $\tilde{\Theta}' = (\mathbf{q}, \mathbf{h}', \sigma_\varepsilon', \sigma_h, \lambda)$ and check the *invariant condition*, i.e., that

$$\int_{\mathbf{h}, \sigma_\varepsilon} \kappa_2(\mathbf{h}', \sigma_\varepsilon' | \tilde{\Theta}, z) \rho(\tilde{\Theta} | z) d\mathbf{h} d\sigma_\varepsilon = \rho(\tilde{\Theta}' | z) \quad (5)$$

holds for step 2, where ρ and κ_2 denote the invariant target distribution and the transition kernel corresponding to the step 2, respectively. We actually have

$$\kappa_2(\mathbf{h}', \sigma_\varepsilon' | \tilde{\Theta}, z) = \int_{\mathbf{x}} \rho(\mathbf{x} | \tilde{\Theta}, z) \rho(\mathbf{h}' | \mathbf{x}, \tilde{\Theta} \setminus \mathbf{h}, z) \rho(\sigma_\varepsilon' | \mathbf{x}, \tilde{\Theta} \setminus \sigma_\varepsilon', z) d\mathbf{x}. \quad (6)$$

We emphasize that Eq. (6) is not obtained by Monte Carlo integration. It is rather the exact expression of the transition kernel of steps 2 to update the pair $\{\mathbf{h}, \sigma_\varepsilon\}$ while \mathbf{x} is no longer retained in the chain. It follows that the left handside of Eq. (5) reads

$$\begin{aligned} & \int \rho(\mathbf{x}, \tilde{\Theta} | z) \rho(\mathbf{h}' | \mathbf{x}, \tilde{\Theta} \setminus \mathbf{h}, z) \rho(\sigma_\varepsilon' | \mathbf{x}, \tilde{\Theta} \setminus \sigma_\varepsilon', z) d\mathbf{h} d\sigma_\varepsilon d\mathbf{x} \\ &= \int \rho(\mathbf{x}, \tilde{\Theta} \setminus \mathbf{h} | z) \rho(\mathbf{h}' | \mathbf{x}, \tilde{\Theta} \setminus \mathbf{h}, z) \rho(\sigma_\varepsilon' | \mathbf{x}, \tilde{\Theta} \setminus \sigma_\varepsilon', z) d\sigma_\varepsilon d\mathbf{x} \\ &= \int \rho(\mathbf{h}', \mathbf{x}, \tilde{\Theta} \setminus \mathbf{h} | z) \rho(\sigma_\varepsilon' | \mathbf{x}, \tilde{\Theta} \setminus \sigma_\varepsilon', z) d\sigma_\varepsilon d\mathbf{x} \\ &= \int \rho(\mathbf{x}, \tilde{\Theta}' \setminus \sigma_\varepsilon' | z) \rho(\sigma_\varepsilon' | \mathbf{x}, \tilde{\Theta}' \setminus \sigma_\varepsilon', z) d\mathbf{x} \\ &= \int \rho(\mathbf{x}, \tilde{\Theta}' | z) d\mathbf{x} = \rho(\tilde{\Theta}' | z). \end{aligned}$$

It is thus proven that the partial marginalization technique can be applied to generate a Markov chain $\{\tilde{\Theta}^{(k)}\}$ that converges to its equilibrium distribution $\rho(\tilde{\Theta}|z)$.

To compare with the example of Fig. 1, our Gibbs-type scheme escapes from a local maximum configuration within an acceptable number of iterations, as Fig. 3 illustrates. The real configuration is reached after 20 iterations. Moreover, it is observed from the configurations obtained at the 18th and 19th iterations that our scheme is able to radically modify the amplitude vector \mathbf{x} in one single step.




Fig. 3. Sampled BG sequences obtained by our Gibbs-type scheme in the example of Fig. 1. The Markov chain escapes rapidly the initial configuration neighborhood.

Finally, it is worthwhile to mention that an efficient sampler should meet the two conflicting criteria [5, Chapter 7]:

- drawing each component conditional on others should be computationally simple (complexity per iteration);
- the induced Markov chain should converge reasonably fast to its equilibrium law (number of iterations).

The next section deals with the issue of the complexity per iteration, whereas the overall computational costs are compared in Sect. 4 using simulation tests.

3.3 Marginal posterior distribution

In this subsection, an efficient implementation of step 1 of Tab. 2 is proposed, inspired from the recursive method in [4] to evaluate $\rho(q_i = 0, 1 | \tilde{\Theta} \setminus q_i; \mathbf{z})$ sequentially. While the latter is based on the storage and update of an N_e^2 ($N_e = \sum_i q_i$) matrix, we finally propose an even less burdensome strategy based on the handling of its Cholesky factor only.

Let \mathbf{H} denote the $N \times M$ Toeplitz matrix such that (1) also reads $\mathbf{z} = \mathbf{H}\mathbf{x} + \varepsilon$. Then, the conditional posterior distribution of q_i takes the following form:

$$\begin{aligned} \rho(q_i | \tilde{\Theta} \setminus q_i; \mathbf{z}) &\propto |\mathbf{B}|^{-1/2} \exp\left(-\frac{1}{2} \mathbf{z}^t \mathbf{B}^{-1} \mathbf{z}\right) \left(\frac{\lambda}{1-\lambda}\right)^{q_i} \\ &\propto \exp(-f(q_i)/2) \end{aligned}$$

where

$$\begin{aligned} \mathbf{B} &= \mathbf{H} \text{diag}\{q\} \mathbf{H}^t + \sigma_\varepsilon^2 \mathbf{I}_N, \\ f(q_i) &= \mathbf{z}^t \mathbf{B}^{-1} \mathbf{z} + \log |\mathbf{B}| + 2q_i \log(1/\lambda - 1). \end{aligned}$$

After normalization, the marginal probability of q_i reads:

$$\rho(q_i = 1 | \tilde{\Theta} \setminus q_i; \mathbf{z}) = (1 + \exp(-(f(0) - f(1))/2))^{-1},$$

which reduces to the evaluation of $f(0) - f(1)$, using the following steps [4]:

$$\begin{aligned} \tilde{\tau}_i &= \delta_i + \mu \mathbf{h}_i^t \tilde{\mathbf{B}}_0^{-1} \mathbf{h}_i \\ \tilde{\mathbf{B}}_i^{-1} &= \tilde{\mathbf{B}}_0^{-1} - \mu \tilde{\mathbf{B}}_0^{-1} \mathbf{h}_i \tilde{\tau}_i^{-1} \mathbf{h}_i^t \tilde{\mathbf{B}}_0^{-1} \\ |\tilde{\mathbf{B}}_i^{-1}| &= |\tilde{\mathbf{B}}_0^{-1}| \delta_i \tilde{\tau}_i \end{aligned}$$

where $\tilde{\mathbf{B}} = \mathbf{B}/\sigma_\varepsilon^2$, $\mu = \sigma_\varepsilon^{-2}$ and $\delta_i = \pm 1$ depending on whether 1 is added or removed at q_i ; \mathbf{h}_i denotes the i th column of matrix \mathbf{H} while $\tilde{\mathbf{B}}_i$ and $\tilde{\mathbf{B}}_0$ differ only at q_i .

Further simplifications are also introduced in [4] by exploiting the sparse nature of $\tilde{\mathbf{B}}$ and then applying the matrix inversion lemma. Noticing that the rank of $\mathbf{H} \text{diag}\{q\} \mathbf{H}^t$ is only N_e , $\tilde{\mathbf{B}}$ takes the alternate form $\tilde{\mathbf{B}} = \mu \mathbf{G} \mathbf{G}^t + \mathbf{I}$, where $\mathbf{G} = \mathbf{H} \mathbf{D}$ is full rank, \mathbf{D} being made of the nonzero columns of $\text{diag}\{q\}$. Then, we have $\tilde{\mathbf{B}}^{-1} = \mathbf{I} - \mu \mathbf{G} \tilde{\mathbf{C}}^{-1} \mathbf{G}^t$, where $\tilde{\mathbf{C}} = \mu \mathbf{G}^t \mathbf{G} + \mathbf{I}$ is $N_e \times N_e$. Thus, only $\tilde{\mathbf{C}}^{-1}$ needs to be stored and updated instead of $\tilde{\mathbf{B}}^{-1}$, which amounts to $\mathcal{O}(N_e^2)$ operations instead of $\mathcal{O}(N^2)$. In the case of adding a pulse at q_i , the formula for the update of $\tilde{\mathbf{C}}_i^{-1}$ is

$$\tilde{\mathbf{C}}_i^{-1} = \begin{bmatrix} \tilde{\mathbf{C}}_0^{-1} + \tilde{\mathbf{b}} \tilde{\tau}_i \tilde{\mathbf{b}}^t & \tilde{\mathbf{b}} \\ \tilde{\mathbf{b}}^t & \tilde{\tau}_i^{-1} \end{bmatrix} \quad (7)$$

where $\tilde{\mathbf{b}} = -\mu \tilde{\tau}_i^{-1} \tilde{\mathbf{C}}_0^{-1} \mathbf{G}_0^t \mathbf{h}_i$. On the other hand, it can be shown that $f(1 - q_i) - f(q_i)$ takes the following form:

$$\begin{aligned} f(1 - q_i) - f(q_i) &= \\ &\log(\delta_i \tilde{\tau}_i) - \mu \tilde{\tau}_i^{-1} \sigma_\varepsilon^{-2} (\mathbf{z}^t \tilde{\mathbf{B}}_0^{-1} \mathbf{h}_i)^2 + 2\delta_i \log\left(\frac{1}{\lambda} - 1\right) \end{aligned} \quad (8)$$

where

$$\begin{aligned} \tilde{\tau}_i &= \delta_i + \mu \|\mathbf{h}_i\|^2 - \mu^2 \|\tilde{\mathbf{R}}_0 \mathbf{G}_0^t \mathbf{h}_i\|^2, \\ \mathbf{z}^t \tilde{\mathbf{B}}_0^{-1} \mathbf{h}_i &= \mathbf{z}^t \mathbf{h}_i - \mu (\tilde{\mathbf{R}}_0 \mathbf{G}_0^t \mathbf{z})^t (\tilde{\mathbf{R}}_0 \mathbf{G}_0^t \mathbf{h}_i). \end{aligned}$$

Last but not least, the computation and memory requirements can be still lowered, by updating and storing the upper triangular *Cholesky* factor $\tilde{\mathbf{R}}$ of $\tilde{\mathbf{C}}^{-1}$, instead of $\tilde{\mathbf{C}}^{-1}$. Since $(\mathbf{x} | \tilde{\Theta}, \mathbf{z})$ is Gaussian with covariance $\tilde{\mathbf{C}}^{-1}$, its Cholesky factor allows to solve step 2(a) in $\mathcal{O}(N_e^2)$ operations instead of $\mathcal{O}(N_e^3)$. In the mean time, $\tilde{\mathbf{R}}$ can be updated in $\mathcal{O}(N_e^2)$ by using a *Cholesky update* method, given that (7) also reads:

$$\tilde{\mathbf{C}}_i^{-1} = \begin{bmatrix} \tilde{\mathbf{C}}_0^{-1} & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} \tilde{\mathbf{b}} \sqrt{\tilde{\tau}_i} \\ 1/\sqrt{\tilde{\tau}_i} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{b}} \sqrt{\tilde{\tau}_i} \\ 1/\sqrt{\tilde{\tau}_i} \end{bmatrix}^t$$

Finally, step 1 of Tab. 2 can be summarized as follows:

- Evaluate $\rho(q_i = 0, 1 | \tilde{\Theta} \setminus q_i; \mathbf{z})$ using (8), and sample $q_i^{(k+1)}$ accordingly;
- Update $\tilde{\mathbf{R}}$ if needed, *i.e.*, if $q_i^{(k+1)} \neq q_i^{(k)}$.

4. SIMULATION RESULTS

A test scenario is designed to compare our proposed method with that of Cheng *et al.*'s sampler [1], in terms of robustness w.r.t. different random initial conditions, while the time-shift ambiguities are dealt with in both methods according to [7].

Brooks *et al.*'s convergence diagnostic [13] is based upon parallel chains. Let $\{\Phi_{jt}, j = 1, \dots, m; t = 1, \dots, n\}$ denote samples from m independent Markov chains of equal length n , and $\bar{\Phi}_j$ (and $\bar{\Phi}_\cdot$) the local mean of the j th chain (*resp.* the global mean). Let the intra-chain and inter-chain variances be defined as covariance matrix averages:

$$\begin{aligned} \mathbf{V}_{\text{intra}} &= \frac{1}{m(n-1)} \sum_{j=1}^m \sum_{t=1}^n (\Phi_{jt} - \bar{\Phi}_j)(\Phi_{jt} - \bar{\Phi}_j)^t \\ \mathbf{V}_{\text{inter}} &= \frac{1}{m-1} \sum_{j=1}^m (\bar{\Phi}_j - \bar{\Phi}_\cdot)(\bar{\Phi}_j - \bar{\Phi}_\cdot)^t. \end{aligned}$$

$\mathbf{V}_{\text{intra}}$ and $\mathbf{V}_{\text{inter}}$ allow to characterize the convergence behavior. Brooks and Gelman [13] proposed to evaluate:

$$\widehat{R} = \frac{n-1}{n} + \left(\frac{m+1}{m}\right) \lambda(\mathbf{V}_{\text{intra}}^{-1} \mathbf{V}_{\text{inter}})$$

where $\lambda(\cdot)$ denotes the largest eigenvalue, and to wait until \widehat{R} is close to one (e.g., $\widehat{R} < 1.2$).

Let us reexamine the example of Fig. 2. To concentrate on the convergence quality of the Bernoulli sequence, we take $\Phi_{ji} = q_{ji}$ for 10 independent Markov chains. Fig. 4 shows that while \widehat{R} is near but still above 1.2 after 10000 iterations for the reference Gibbs sampler, this threshold is reached after about 1000 iterations for the proposed sampler.




Fig. 4. Evolution of \widehat{R} for both methods in the case of SNR = 12.81 dB. \widehat{R} has been computed every 100 samples.

To further illustrate the robustness of the proposed sampler, deconvolution results are shown in Fig. 5 to compare with that of the reference sampler in Fig. 2 under identical initial conditions. Only one of the results is reported, as they are undistinguishable from each other (perfectly consistent estimations). We noticed that instead of calculating the posterior Gaussian mean based on $\{\widehat{\mathbf{q}}, \widehat{\mathbf{h}}, \mathbf{z}\}$, the estimator $\widehat{\mathbf{x}}$ can also be obtained by averaging the auxiliarily sampled $\mathbf{x}^{(k)}$ in step 2(a). In fact, as $\{\Theta^{(k)}\}$ is shown to converge to the distribution $\rho(\Theta | \mathbf{z})$ in Sect. 3.2, step 2(a) samples \mathbf{x} conditional on Θ and \mathbf{z} such that taken marginally $\mathbf{x}^{(k)}$ follows the distribution $\rho(\mathbf{x} | \mathbf{z})$.




Fig. 5. Estimation result using the proposed sampler, under the same conditions as in Fig. 2. The 10 independent Markov chains yield undistinguishable estimates.

Most importantly, it should also be noted that 1000 iterations of the Gibbs-type sampling take slightly less time than 2000 iterations of the reference Gibbs sampler [7]. We are thus able to conclude that our Gibbs-type sampler achieves a far better compromise between the two criteria [5] than the classical Gibbs sampler [1, 7]: in the displayed example, the time required to reach convergence is reduced by a factor of at least 5.

REFERENCES

- [1] Q. Cheng, R. Chen, and T.-H. Li, “Simultaneous wavelet estimation and deconvolution of reflection seismic signals”, *IEEE Trans. Geosci. Remote Sensing*, vol. 34, pp. 377–384, Mar. 1996.
- [2] S. Bourguignon and H. Carfantan, “Bernoulli-Gaussian spectral analysis of unevenly spaced astrophysical data”, in *IEEE Workshop Stat. Sig. Proc.*, Bordeaux, France, July 2005, pp. 811–816.
- [3] J. J. Kormylo and J. M. Mendel, “Maximum-likelihood seismic deconvolution”, *IEEE Trans. Geosci. Remote Sensing*, vol. GE-21, no. 1, pp. 72–82, Jan. 1983.
- [4] F. Champagnat, Y. Goussard, and J. Idier, “Unsupervised deconvolution of sparse spike trains using stochastic approximation”, *IEEE Trans. Signal Processing*, vol. 44, no. 12, pp. 2988–2998, Dec. 1996.
- [5] J. S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer Series in Statistics. Springer Verlag, New York, NY, 2001.
- [6] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, Springer Texts in Statistics. Springer Verlag, New York, NY, 2nd edition, 2004.
- [7] C. Labat and J. Idier, “Sparse blind deconvolution accounting for time-shift ambiguity”, in *Proc. IEEE ICASSP*, Toulouse, France, May 2006, pp. 616–619.
- [8] Y. Guan and S. M. Krone, “Small-world MCMC and convergence to multi-modal distributions: From slow mixing to fast mixing”, *J. Appl. Prob.*, vol. 17, no. 1, pp. 284–304, Mar. 2007.
- [9] N. Dobigeon, J.-Y. Tournet, and J. D. Scargle, “Joint segmentation of multivariate astronomical times series: Bayesian sampling with a hierarchical model”, *IEEE Trans. Signal Processing*, vol. 55, no. 2, pp. 414–423, Feb. 2007.
- [10] C. Y. Chi and J. M. Mendel, “Improved maximum-likelihood detection and estimation of Bernoulli-Gaussian processes”, *IEEE Trans. Inf. Theory*, vol. 30, pp. 429–435, 1984.
- [11] J. S. Liu, W. H. Wong, and A. Kong, “Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes”, *Biometrika*, vol. 81, pp. 27–40, 1994.
- [12] O. Cappé, A. Doucet, M. Lavielle, and E. Moulines, “Simulation-based methods for blind maximum-likelihood filter identification”, *Signal Processing*, vol. 73, no. 1-2, pp. 3–25, Feb. 1999.
- [13] S. P. Brooks and A. Gelman, “General methods for monitoring convergence of iterative simulations”, *J. Comput. Graph. Statist.*, vol. 7, no. 4, pp. 434–455, 1998.