



HAL
open science

Efficient Gaussian Sampling for Solving Large-Scale Inverse Problems using MCMC

Clément Gilavert, Saïd Moussaoui, Jérôme Idier

► **To cite this version:**

Clément Gilavert, Saïd Moussaoui, Jérôme Idier. Efficient Gaussian Sampling for Solving Large-Scale Inverse Problems using MCMC. IEEE Transactions on Signal Processing, Institute of Electrical and Electronics Engineers, 2015, 63 (01), pp.70-80. 10.1109/TSP.2014.2367457. hal-01059414

HAL Id: hal-01059414

<https://hal.archives-ouvertes.fr/hal-01059414>

Submitted on 31 Aug 2014

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.



Distributed under a Creative Commons Attribution| 4.0 International License

Efficient Gaussian Sampling for Solving Large-Scale Inverse Problems using MCMC Methods

Clément Gilavert, Saïd Moussaoui and Jérôme Idier

Ecole Centrale Nantes, IRCCyN, CNRS UMR 6597,
1 rue de la Noë, 44321, Nantes Cedex 3, France.

Corresponding author: said.moussaoui@ec-nantes.fr

August 31, 2014

Abstract

The resolution of many large-scale inverse problems using MCMC methods requires a step of drawing samples from a high dimensional Gaussian distribution. While direct Gaussian sampling techniques, such as those based on Cholesky factorization, induce an excessive numerical complexity and memory requirement, sequential coordinate sampling methods present a low rate of convergence. Based on the reversible jump Markov chain framework, this paper proposes an efficient Gaussian sampling algorithm having a reduced computation cost and memory usage. The main feature of the algorithm is to perform an approximate resolution of a linear system with a truncation level adjusted using a self-tuning adaptive scheme allowing to achieve the minimal computation cost. The connection between this algorithm and some existing strategies is discussed and its efficiency is illustrated on a linear inverse problem of image resolution enhancement.

This paper is under revision before publication in IEEE Transactions on Signal Processing

1 Introduction

A common inverse problem arising in several signal and image processing applications is to recover a hidden object $\mathbf{x} \in \mathbb{R}^N$ (e.g., an image or a signal) from a set of measurements $\mathbf{y} \in \mathbb{R}^M$ given an observation model [1, 2]. The most frequent case is that of a linear model between \mathbf{x} and \mathbf{y} according to

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \quad (1)$$

with $\mathbf{H} \in \mathbb{R}^{M \times N}$ the known observation matrix and \mathbf{n} an additive noise term representing measurement errors and model uncertainties. Such a linear model covers many real problems such as, for instance, denoising [3], deblurring [4], and reconstruction from projections [5, 6].

The statistical estimation of \mathbf{x} in a Bayesian simulation framework [7, 8] firstly requires the formulation of the posterior distribution $P(\mathbf{x}, \Theta | \mathbf{y})$, with Θ a set of unknown hyper-parameters. Pseudo-random samples of \mathbf{x} are then drawn from this posterior distribution. Finally, a Bayesian estimator is computed from the set of generated samples. Other quantities of interest, such as posterior variances, can be estimated likewise. Within the *standard Monte Carlo* framework, independent realizations of the posterior law must be generated, which is rarely possible in realistic cases of inverse problems. One rather resorts to *Markov Chain*

Monte Carlo (MCMC) schemes, where Markovian dependencies between successive samples are allowed. A very usual sampling scheme is then to iteratively draw realizations from the conditional posterior densities $P(\Theta|\mathbf{x}, \mathbf{y})$ and $P(\mathbf{x}|\Theta, \mathbf{y})$, according to a *Gibbs sampler* [9].

In such a context, when independent Gaussian models $\mathcal{N}(\boldsymbol{\mu}_y, \mathbf{R}_y)$ and $\mathcal{N}(\boldsymbol{\mu}_x, \mathbf{R}_x)$ are assigned to the noise statistics and to the unknown object distribution, respectively, the set of hyper-parameters Θ determines the mean and the covariance of the latter two distributions. This statistical model also covers the case of priors based on hierarchical or latent Gaussian models such as Gaussian scale mixtures [10, 11] and Gaussian Markov random fields [9, 12]. The additional parameters of such models are then included in Θ . According to this Bayesian modeling, the conditional posterior distribution $P(\mathbf{x}|\Theta, \mathbf{y})$ is also Gaussian, $\mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}^{-1})$, with a precision matrix \mathbf{Q} (*i.e.*, the inverse of the covariance matrix) given by

$$\mathbf{Q} = \mathbf{H}^t \mathbf{R}_y^{-1} \mathbf{H} + \mathbf{R}_x^{-1}, \quad (2)$$

and a mean vector $\boldsymbol{\mu}$ such that:

$$\mathbf{Q}\boldsymbol{\mu} = \mathbf{H}^t \mathbf{R}_y^{-1} (\mathbf{y} - \boldsymbol{\mu}_y) + \mathbf{R}_x^{-1} \boldsymbol{\mu}_x. \quad (3)$$

Let us remark that the precision matrix \mathbf{Q} generally depends on the hyper-parameter set Θ through \mathbf{R}_y and \mathbf{R}_x , so that \mathbf{Q} is a varying matrix along the Gibbs sampler iterations. Moreover, the mean vector $\boldsymbol{\mu}$ is expressed as the solution of a linear system where \mathbf{Q} is the normal matrix.

In order to draw samples from the conditional posterior distribution $P(\mathbf{x}|\Theta, \mathbf{y})$, a usual way is to firstly perform the Cholesky factorization of the covariance matrix [13, 14]. Since equation (2) yields the precision matrix \mathbf{Q} rather than the covariance matrix, Rue [15] proposed to compute the Cholesky decomposition of \mathbf{Q} , *i.e.*, $\mathbf{Q} = \mathbf{C}_q \mathbf{C}_q^t$, and to solve the triangular system $\mathbf{C}_q^t \mathbf{x} = \boldsymbol{\omega}$, where $\boldsymbol{\omega}$ is a vector of independent Gaussian variables of zero mean and unit variance. Moreover, the Cholesky factorization is exploited to calculate the mean $\boldsymbol{\mu}$ from (3) by solving two triangular systems sequentially. However, the Cholesky factorization of \mathbf{Q} generally requires $\mathcal{O}(N^3)$ operations. Spending such a numerical cost at each iteration of the sampling scheme rapidly becomes prohibitive for large values of N . In specific cases where \mathbf{Q} belongs to certain families of structured matrices, the factorization can be obtained with a reduced numerical complexity, *e.g.*, $\mathcal{O}(N^2)$ when \mathbf{Q} is Toeplitz [16] or even $\mathcal{O}(N \log N)$ when \mathbf{Q} is circulant [17]. Sparse matrices can be also factored at a reduced cost [15, 18]. Alternative approaches to the Cholesky factorization are based on using an iterative method for the calculation of the inverse square root matrix of \mathbf{Q} using Krylov subspace methods [19–21]. In practice, even in such favorable cases, the factorization often remains a burdensome operation to be performed at each iteration of the Gibbs sampler.

The numerical bottleneck represented by the Cholesky factorization can be removed by using alternative schemes that bypass the step of exactly sampling $P(\mathbf{x}|\Theta, \mathbf{y})$. For instance, a simple alternative solution is to sequentially sample each entry of \mathbf{x} given the other variables according to a scalar Gibbs scheme [22]. However, such a scalar approach reveals extremely inefficient when $P(\mathbf{x}|\Theta, \mathbf{y})$ is strongly correlated, since each conditional sampling step will produce a move of very small variance. As a consequence, a huge number of iterations will be required to reach convergence. A better trade-off between the numerical cost of each iteration and the overall convergence speed of the sampler must be found.

In this paper, we focus on a two-step approach named *Independent Factor Perturbation* in [12] and *Perturbation-Optimization* in [23] (see also [18, 24]). It consists in

- drawing a sample $\boldsymbol{\eta}$ from $\mathcal{N}(\mathbf{Q}\boldsymbol{\mu}, \mathbf{Q})$,
- solving the linear system $\mathbf{Q}\mathbf{x} = \boldsymbol{\eta}$.

It can be easily checked that, when the linear system is solved exactly, the new sample \mathbf{x} is distributed according to $\mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}^{-1})$. Hereafter, we refer to this method as *Exact Perturbation Optimization* (E-PO). However, the numerical cost of E-PO is typically as high as the Cholesky factorization of \mathbf{Q} . Therefore, an essential element of the Perturbation Optimization approach is to truncate the linear system solving by running a limited number of iterations of an iterative algorithm such as the conjugate gradient method

(CG) [12,23,24]. For the sake of clarity, let us call the resulting version *Truncated Perturbation Optimization* (T-PO).

Skipping from E-PO to T-PO allows to strongly reduce the numerical cost of each iteration. However, let us stress that no convergence analysis of T-PO exists, to our best knowledge. It is only argued that a well-chosen truncation level will induce a significant reduction of the numerical cost and a small error on \mathbf{x} . The way the latter error alters the convergence towards the target distribution remains a fully open issue, that has not been discussed in existing contributions. Moreover, how the resolution accuracy should be chosen in practice is also an open question.

A first contribution of the present paper is to bring practical evidence that the T-PO does not necessarily converge towards the target distribution (see Section 4). In practice, the implicit trade-off within T-PO is between the computational cost and the error induced on the target distribution, depending on the adopted truncation level. Our second contribution is to propose a new scheme similar to T-PO, but with a guarantee of convergence to the target distribution, whatever the truncation level. We call the resulting scheme *Reversible Jump Perturbation Optimization* (RJPO), since it incorporates an accept-reject step derived within the Reversible Jump MCMC (RJ-MCMC) framework [25,26]. Let us stress here that the numerical cost of the proposed test is marginal, so that RJPO has nearly the same cost per iteration as T-PO. Finally, we propose an unsupervised tuning of the truncation level allowing to automatically achieve a pre-specified overall acceptance rate or even to minimize the computation cost at a constant effective sample size. The resulting algorithm can be viewed as an adaptive MCMC sampler [27,28].

The rest of the paper is organized as follows: Section 2 introduces the global framework of RJ-MCMC and presents a general scheme to sample Gaussian vectors. Section 3 considers a specific application of the previous results, which finally boils down to the proposed RJPO sampler. Section 4 analyses the performance of RJPO compared to T-PO on simple toy problems and presents the adaptive RJPO which incorporates an automatic control of the truncation level. Finally, in section 5, an example of linear inverse problem, the unsupervised image resolution enhancement is presented to illustrate the applicability of the method. These results show the superiority of the RJPO algorithm over the usual Cholesky factorization based approaches in terms of computational cost and memory usage.

2 The reversible jump MCMC framework

The sampling procedure consists on constructing a Markov chain whose distribution asymptotically converges to the target distribution $P_{\mathbf{X}}(\cdot)$. Let $\underline{\mathbf{x}} \in \mathbb{R}^N$ be the current sample of the Markov chain.

2.1 General framework

In the constant dimension case, the Reversible Jump MCMC strategy [25,26] introduces an auxiliary variable $\mathbf{z} \in \mathbb{R}^L$, obtained from a distribution $P_{\mathbf{Z}}(\mathbf{z}|\underline{\mathbf{x}})$ and a deterministic move according to a differentiable transformation

$$\begin{aligned} \phi : (\mathbb{R}^N \times \mathbb{R}^L) &\mapsto (\mathbb{R}^N \times \mathbb{R}^L) \\ (\underline{\mathbf{x}}, \mathbf{z}) &\mapsto (\mathbf{x}, \mathbf{s}) \end{aligned}$$

This transformation must also be reversible, that is $\phi(\mathbf{x}, \mathbf{s}) = (\underline{\mathbf{x}}, \mathbf{z})$. The new sample $\bar{\mathbf{x}}$ is thereby obtained by submitting \mathbf{x} (resulting from the deterministic move) to an accept-reject step with an acceptance probability given by

$$\alpha(\underline{\mathbf{x}}, \mathbf{x}|\mathbf{z}) = \min \left(1, \frac{P_{\mathbf{X}}(\mathbf{x})P_{\mathbf{Z}}(\mathbf{s}|\mathbf{x})}{P_{\mathbf{X}}(\underline{\mathbf{x}})P_{\mathbf{Z}}(\mathbf{z}|\underline{\mathbf{x}})} |J_{\phi}(\underline{\mathbf{x}}, \mathbf{z})| \right),$$

with $J_{\phi}(\underline{\mathbf{x}}, \mathbf{z})$ the Jacobian determinant of the transformation ϕ at $(\underline{\mathbf{x}}, \mathbf{z})$.

Actually, the choice of the conditional distribution $P_{\mathbf{Z}}(\cdot)$ and the transformation $\phi(\cdot)$ must be adapted to the target distribution $P_{\mathbf{X}}(\cdot)$ and affects the resulting Markov chain properties in terms of correlation and convergence rate.

2.2 Gaussian case

To sample from a Gaussian distribution $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}^{-1})$, we generalize the scheme adopted in [29]. We set $L = N$ and take an auxiliary variable $\mathbf{z} \in \mathbb{R}^N$ distributed according to

$$P_{\mathbf{Z}}(\mathbf{z}|\underline{\mathbf{x}}) = \mathcal{N}(\mathbf{A}\underline{\mathbf{x}} + \mathbf{b}, \mathbf{B}), \quad (4)$$

where \mathbf{A} , \mathbf{B} and \mathbf{b} denote a $N \times N$ real matrix, a $N \times N$ real positive definite matrix and a $N \times 1$ real vector, respectively. The choice of the latter three quantities will be discussed later. The proposed deterministic move is performed using the transformation ϕ such that

$$\begin{pmatrix} \mathbf{x} \\ \mathbf{s} \end{pmatrix} = \begin{pmatrix} \phi_1(\underline{\mathbf{x}}, \mathbf{z}) \\ \phi_2(\underline{\mathbf{x}}, \mathbf{z}) \end{pmatrix} = \begin{pmatrix} -\underline{\mathbf{x}} + \mathbf{f}(\mathbf{z}) \\ \mathbf{z} \end{pmatrix}, \quad (5)$$

with functions $(\phi_1 : (\mathbb{R}^N \times \mathbb{R}^N) \mapsto \mathbb{R}^N)$ and $(\phi_2 : (\mathbb{R}^N \times \mathbb{R}^N) \mapsto \mathbb{R}^N)$ and $(\mathbf{f} : \mathbb{R}^N \mapsto \mathbb{R}^N)$.

Proposition 1. *Let an auxiliary variable \mathbf{z} be obtained according to (4) and a proposed sample \mathbf{x} resulting from (5). Then the acceptance probability is*

$$\alpha(\underline{\mathbf{x}}, \mathbf{x}|\mathbf{z}) = \min\left(1, e^{-\mathbf{r}(\mathbf{z})^t(\underline{\mathbf{x}}-\mathbf{x})}\right), \quad (6)$$

with

$$\mathbf{r}(\mathbf{z}) = \mathbf{Q}\boldsymbol{\mu} + \mathbf{A}^t\mathbf{B}^{-1}(\mathbf{z} - \mathbf{b}) - \frac{1}{2}(\mathbf{Q} + \mathbf{A}^t\mathbf{B}^{-1}\mathbf{A})\mathbf{f}(\mathbf{z}). \quad (7)$$

In particular, the acceptance probability equals one when $\mathbf{f}(\mathbf{z})$ is defined as the exact solution of the linear system

$$\frac{1}{2}(\mathbf{Q} + \mathbf{A}^t\mathbf{B}^{-1}\mathbf{A})\mathbf{f}(\mathbf{z}) = \mathbf{Q}\boldsymbol{\mu} + \mathbf{A}^t\mathbf{B}^{-1}(\mathbf{z} - \mathbf{b}). \quad (8)$$

Proof. See appendix A. □

Let us emphasize that \mathbf{b} is a dummy parameter, since the residual $\mathbf{r}(\mathbf{z})$ (and thus $\alpha(\underline{\mathbf{x}}, \mathbf{x})$) depends on \mathbf{b} through $\mathbf{z} - \mathbf{b}$ only. However, choosing a specific expression of \mathbf{b} jointly with \mathbf{A} and \mathbf{B} will lead to a simplified expression of $\mathbf{r}(\mathbf{z})$ in the next section.

Proposition 1 plays a central role in our proposal. When the exact resolution of (8) is numerically costly, it allows to derive a procedure where the resolution is performed only approximately, at the expense of a lowered acceptance probability. The conjugate gradient algorithm stopped before convergence, is a typical example of an efficient tool allowing to approximately solve (8).

Proposition 2. *Let an auxiliary variable \mathbf{z} be obtained according to (4), a proposed sample \mathbf{x} resulting from (5) and $\mathbf{f}(\mathbf{z})$ be the exact solution of (8). The correlation between two successive samples is zero if and only if matrices \mathbf{A} and \mathbf{B} are chosen such that*

$$\mathbf{A}^t\mathbf{B}^{-1}\mathbf{A} = \mathbf{Q}. \quad (9)$$

Proof. See Appendix B. □

Many couples (\mathbf{A}, \mathbf{B}) fulfill condition (9)

- Consider the Cholesky factorization $\mathbf{Q} = \mathbf{C}_q\mathbf{C}_q^t$ and take $\mathbf{A} = \mathbf{C}_q^t$, $\mathbf{B} = \mathbf{I}$. It leads to $\mathbf{z} = \mathbf{C}_q^t\underline{\mathbf{x}} + \mathbf{b} + \boldsymbol{\omega}$ with $\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$. According to (8), the next sample $\bar{\mathbf{x}} = -\underline{\mathbf{x}} + \mathbf{f}(\mathbf{z})$, will be obtained as

$$\begin{aligned} \bar{\mathbf{x}} &= -\underline{\mathbf{x}} + (\mathbf{C}_q\mathbf{C}_q^t)^{-1}(\mathbf{Q}\boldsymbol{\mu} + \mathbf{C}_q(\mathbf{z} - \mathbf{b})), \\ &= (\mathbf{C}_q^t)^{-1}(\mathbf{C}_q^{-1}\mathbf{Q}\boldsymbol{\mu} + \boldsymbol{\omega}). \end{aligned}$$

Such an update scheme is exactly the same as the one proposed by Rue in [15].

- The particular configuration

$$\mathbf{A} = \mathbf{B} = \mathbf{Q} \quad \text{and} \quad \mathbf{b} = \mathbf{Q}\boldsymbol{\mu}. \quad (10)$$

is retained in the sequel, since:

- i) $\mathbf{A}^\dagger \mathbf{B}^{-1} \mathbf{A} = \mathbf{Q}$ is a condition of Proposition 2;
- ii) $\mathbf{b} = \mathbf{Q}\boldsymbol{\mu}$ simplifies equation (8) to a linear system $\mathbf{Q}\mathbf{f}(\mathbf{z}) = \mathbf{z}$.

In particular, it allows to make a clear connection between our RJ-MCMC approach and the E-PO algorithm in the case of an exact resolution of the linear system.

3 Gaussian Sampling in the Reversible Jump MCMC framework

3.1 Sampling the Auxiliary Variable

According to (10), the auxiliary variable \mathbf{z} is distributed according to $\mathcal{N}(\mathbf{Q}\underline{\mathbf{x}} + \mathbf{Q}\boldsymbol{\mu}, \mathbf{Q})$. It can then be expressed as $\mathbf{z} = \mathbf{Q}\underline{\mathbf{x}} + \boldsymbol{\eta}$, $\boldsymbol{\eta}$ being distributed according to $\mathcal{N}(\mathbf{Q}\boldsymbol{\mu}, \mathbf{Q})$. Consequently, the auxiliary variable sampling step is reduced to the simulation of $\boldsymbol{\eta}$, which is the perturbation step in the PO algorithm. In [12,23], a subtle way of sampling $\boldsymbol{\eta}$ is proposed. It consists in exploiting equation (3) and *perturbing* each factor separately:

1. Sample $\boldsymbol{\eta}_y \sim \mathcal{N}(\mathbf{y} - \boldsymbol{\mu}_y, \mathbf{R}_y)$,
2. Sample $\boldsymbol{\eta}_x \sim \mathcal{N}(\boldsymbol{\mu}_x, \mathbf{R}_x)$,
3. Set $\boldsymbol{\eta} = \mathbf{H}^\dagger \mathbf{R}_y^{-1} \boldsymbol{\eta}_y + \mathbf{R}_x^{-1} \boldsymbol{\eta}_x$, a sample of $\mathcal{N}(\mathbf{Q}\boldsymbol{\mu}, \mathbf{Q})$.

It is important to notice that such a tricky method is interesting since matrices \mathbf{R}_y and \mathbf{R}_x have often a simple structure if not diagonal.

We emphasize that this perturbation step can be applied more generally for the sampling of any Gaussian distribution, for which a factored expression of the precision matrix \mathbf{Q} is available under the form $\mathbf{Q} = \mathbf{F}^\dagger \mathbf{F}$, with matrix $\mathbf{F} \in \mathbb{R}^{N' \times N}$. In such a case, $\boldsymbol{\eta} = \mathbf{Q}\boldsymbol{\mu} + \mathbf{F}^\dagger \boldsymbol{\omega}$, where $\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{N'})$.

3.2 Exact resolution case

As stated by proposition 1, the exact resolution of system (8) implies an acceptance probability of one. The resulting sampling procedure is thus based on the following steps:

1. Sample $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{Q}\boldsymbol{\mu}, \mathbf{Q})$,
2. Set $\mathbf{z} = \mathbf{Q}\underline{\mathbf{x}} + \boldsymbol{\eta}$,
3. Take $\bar{\mathbf{x}} = -\underline{\mathbf{x}} + \mathbf{Q}^{-1}\mathbf{z}$.

Let us remark that $\bar{\mathbf{x}} = -\underline{\mathbf{x}} + \mathbf{Q}^{-1}(\mathbf{Q}\underline{\mathbf{x}} + \boldsymbol{\eta}) = \mathbf{Q}^{-1}\boldsymbol{\eta}$, so the handling of variable \mathbf{z} can be skipped and Steps 2 and 3 can be merged to an equivalent but more direct step:

2. Set $\bar{\mathbf{x}} = \mathbf{Q}^{-1}\boldsymbol{\eta}$.

In the exact resolution case, the obtained algorithm is thus identical to the E-PO algorithm [23].

According to Proposition 2, E-PO enjoys the property that each sample is totally independent from the previous ones. However, a drawback is that the exact resolution of the linear system $\mathbf{Q}\mathbf{x} = \boldsymbol{\eta}$ often leads to an excessive numerical complexity and memory usage in high dimensions [12]. In practice, early stopping of an iterative solver such as the linear conjugate gradient algorithm is used, yielding the Truncated Perturbation Optimization (T-PO) version. The main point is that, up to our knowledge, there is no theoretical analysis of the efficiency of T-PO and of its convergence to the target distribution. Indeed, the simulation tests provided in Section 4 indicate that convergence to the target distribution is not guaranteed. As shown in the next subsection, two slight but decisive modifications of T-PO lead us to the RJPO version, which is a provably convergent algorithm.

3.3 Approximate resolution case

In the case (10), equation (7) reduces to

$$\mathbf{r}(\mathbf{z}) = \mathbf{z} - \mathbf{Q}\mathbf{f}(\mathbf{z}). \quad (11)$$

Therefore, a first version of the RJPO algorithm is as follows:

1. Sample $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{Q}\boldsymbol{\mu}, \mathbf{Q})$.
2. Set $\mathbf{z} = \mathbf{Q}\underline{\mathbf{x}} + \boldsymbol{\eta}$. Solve the linear system $\mathbf{Q}\mathbf{u} = \mathbf{z}$, in an approximate way. Let $\hat{\mathbf{u}}$ denote the obtained solution, $\mathbf{r}(\mathbf{z}) = \mathbf{z} - \mathbf{Q}\mathbf{u}$ and propose $\hat{\mathbf{x}} = -\underline{\mathbf{x}} + \hat{\mathbf{u}}$.
3. With probability $\min\left(1, e^{-r(\mathbf{z})^\dagger(\underline{\mathbf{x}} - \hat{\mathbf{x}})}\right)$, set $\bar{\mathbf{x}} = \hat{\mathbf{x}}$, otherwise set $\bar{\mathbf{x}} = \underline{\mathbf{x}}$.

An important point concerns the initialization of the linear solver in Step 2: in the case of an early stopping, the computed approximate solution may depend on the initial point \mathbf{u}_0 . On the other hand, $\mathbf{f}(\mathbf{z})$ must not depend on $\underline{\mathbf{x}}$, otherwise the reversibility of the deterministic move (5) would not be ensured. Hence, the initial point \mathbf{u}_0 must not depend on $\underline{\mathbf{x}}$ either. In the rest of the paper, $\mathbf{u}_0 = \mathbf{0}$ is the default choice.

A more compact and direct version of the sampler can be obtained by substituting $\mathbf{x} = \mathbf{f}(\mathbf{z}) - \underline{\mathbf{x}}$ in equation (11). The latter reduces to the solving of the system $\mathbf{Q}\mathbf{x} = \boldsymbol{\eta}$. Step 2 of the RJPO algorithm is then simplified to:

2. Solve the linear system $\mathbf{Q}\mathbf{x} = \boldsymbol{\eta}$ in an approximate way. Let $\hat{\mathbf{x}}$ denote the obtained solution and $\mathbf{r}(\mathbf{z}) = \boldsymbol{\eta} - \mathbf{Q}\hat{\mathbf{x}}$.

For the reason just discussed above, the initial point \mathbf{x}_0 of the linear solver must be such that $\mathbf{u}_0 = \mathbf{x}_0 + \underline{\mathbf{x}}$ does not depend on $\underline{\mathbf{x}}$. Hence, as counterintuitive as it may be, choices such as $\mathbf{x}_0 = \mathbf{0}$ or $\mathbf{x}_0 = \underline{\mathbf{x}}$ are not allowed, while $\mathbf{x}_0 = -\underline{\mathbf{x}}$ is the default choice corresponding to $\mathbf{u}_0 = \mathbf{0}$.

It is remarkable that both T-PO and the proposed algorithm (RJPO) rely on the approximate resolution of the same linear system $\mathbf{Q}\mathbf{x} = \boldsymbol{\eta}$. However, RJPO algorithm incorporates two additional ingredients that make the difference in terms of mathematical validity:

- RJPO relies on an accept-reject strategy to ensure the sampler convergence in the case of an approximate system solving,
- There is a constraint on the initial point \mathbf{x}_0 of the linear system solving: $\mathbf{x}_0 + \underline{\mathbf{x}}$ must not depend on $\underline{\mathbf{x}}$.

3.4 Implementation issues

There is no constraint on the choice of the linear solver, nor on the early stopping rule, except that they must not depend on the value of $\underline{\mathbf{x}}$. Indeed, any linear system solver, or any quadratic programming method could be employed. In the sequel, we have adopted the linear conjugate gradient algorithm for two reasons:

- Early stopping (*i.e.*, *truncating*) the conjugate gradient iterations is a very usual procedure to approximately solve a linear system, with well-known convergence properties towards the exact solution [30]. Moreover, a preconditioned conjugate gradient could well be used to accelerate the convergence speed.
- It lends itself to a matrix-free implementation with reduced memory requirements, as far as matrix-vector products involving matrix \mathbf{Q} can be performed without explicitly manipulating such a matrix.

On the other hand, we have selected a usual stopping rule based on a threshold on the relative residual norm:

$$\epsilon = \frac{\|\boldsymbol{\eta} - \mathbf{Q}\mathbf{x}\|_2}{\|\boldsymbol{\eta}\|_2}. \quad (12)$$

4 Performance analysis

The aim of this section is to analyze the performance of the RJPO algorithm and to discuss the influence of the relative residual norm (and hence, the truncation level of the conjugate gradient (CG) algorithm) on the performances of the proposed RJPO algorithm. A second part is dedicated to the analysis of the optimal tuning of the algorithm.

4.1 Incidence of the approximate resolution

Let us consider a Gaussian distribution with a precision matrix $\mathbf{Q} = \mathbf{R}^{-1}$ and a mean vector $\boldsymbol{\mu}$ defined by

$$R_{ij} = \sigma^2 \rho^{|i-j|}, \quad (\forall i, j = 1, \dots, N)$$

$$\mu_i \sim \mathcal{U}[0, 10], \quad (\forall i, \dots, N)$$

with $N = 20$, $\sigma^2 = 1$ and $\rho = 0.8$. Figure 1 shows the distribution of $K = 5000$ samples obtained by running the TPO algorithm for different truncation levels of a conjugate gradient algorithm. It can be noted that an early stopping, with $J < 5$, leads to a sample distribution different from the target one. One can also notice that the related acceptance probability expressed in the RJMCMC framework suggests at least $J = 6$ iterations to get samples with a nonzero acceptance probability. One can also see that an exact resolution is not needed since the acceptance probability is almost equal to one after $J > 10$ iterations. This result allows to conclude that the idea of truncating the system resolution is relevant, since it allows to avoid unnecessary calculations, but an acceptance-reject step must be added to ensure a correct behavior of the sampler.

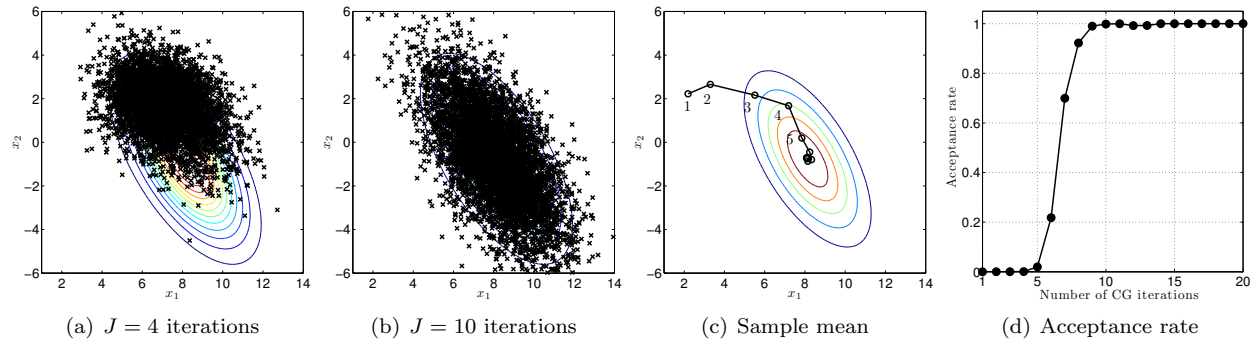


Figure 1: Influence of the truncation level on the distribution of $K = 10^5$ samples obtained by the TPO algorithm.

4.2 Acceptance rate

We focus in this experiment on a small size problem ($N = 16$) to discuss the influence of the truncation level on the numerical performance in terms of acceptance rate and estimation error. For the retained Gaussian sampling schemes, both RJPO and T-PO are run for a number of CG iterations allowing to reach a predefined value of relative residual norm (12). We also discuss the influence of the problem dimension on the best value of the truncation level leading to a minimal total number of CG iterations before convergence.

Figure 2 illustrates the average acceptance probability obtained over $n_{\max} = 10^5$ iterations of the RJPO sampler for different relative residual norm values. It can be noted that the acceptance rate is almost zero when the relative residual norm is larger than 10^{-2} and monotonically increases for higher resolution accuracies. Moreover, a relative residual norm lower than 10^{-5} leads to an acceptance probability almost equal to one. Such a curve indicates that the stopping criterion of the CG must be chosen carefully in order to run the RJPO algorithm efficiently and to get non-zero acceptance probabilities. Finally, note that

this curve mainly depends on the condition number of the precision matrix \mathbf{Q} . Even if the shape of the acceptance curve stays the same for different problems, it happens to be difficult to determine the value of the relative residual norm that corresponds to a given acceptance rate.

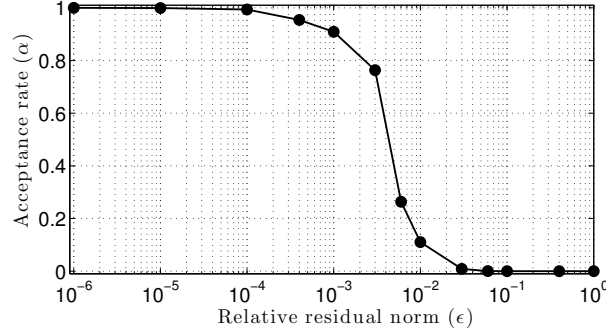


Figure 2: Acceptance rate α of the RJPO algorithm for different values of the relative residual norm in a small size problem ($N = 16$).

4.3 Estimation error

The estimation error is assessed as the relative mean square error (RMSE) on the estimated mean vector and covariance matrix

$$\begin{cases} \text{RMSE}(\boldsymbol{\mu}) = \frac{\|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\|_2}{\|\boldsymbol{\mu}\|_2} \\ \text{RMSE}(\mathbf{R}) = \frac{\|\mathbf{R} - \hat{\mathbf{R}}\|_F}{\|\mathbf{R}\|_F} \end{cases} \quad (13)$$

where $\|\cdot\|_F$ and $\|\cdot\|_2$ represent the Frobenius and the ℓ_2 norms, respectively. $\boldsymbol{\mu}$, \mathbf{R} , $\hat{\boldsymbol{\mu}}$, and $\hat{\mathbf{R}}$ are respectively the mean and the covariance matrix of the Gaussian vector, and their empirical estimates using the generated Markov chain samples according to

$$\begin{cases} \hat{\boldsymbol{\mu}} = \frac{1}{n_{\max} - n_{\min} + 1} \sum_{n=n_{\min}}^{n_{\max}} \mathbf{x}_n \\ \hat{\mathbf{R}} = \frac{1}{n_{\max} - n_{\min}} \sum_{n=n_{\min}}^{n_{\max}} (\mathbf{x}_n - \hat{\boldsymbol{\mu}})(\mathbf{x}_n - \hat{\boldsymbol{\mu}})^t \end{cases}$$

with n_{\min} iterations of burn-in and n_{\max} total iterations.

As expected, Figure 3 indicates that the estimation error is very high if the acceptance rate is zero (when the relative residual norm is lower than 10^{-2}), even for RJPO after $n_{\max} = 10^5$ iterations. This is due to the very low acceptance rate which slows down the chain convergence. However, as soon as new samples are accepted, RJPO leads to the same performance as when the system is solved exactly (E-PO algorithm). On the other hand, T-PO keeps a significant error for small and moderate resolution accuracies. Naturally, both methods present similar performance when the relative residual norm is very low since these methods tend to provide almost the same samples with an acceptance probability equal to one. This experimental result clearly highlights the deficiency of T-PO: the system must be solved with a relatively high accuracy to avoid an important estimation error. On the other hand, in the RJPO algorithm the acceptance rate is a good indicator whether the value of the relative residual norm threshold is appropriate to ensure a sufficient mixing of the chain.

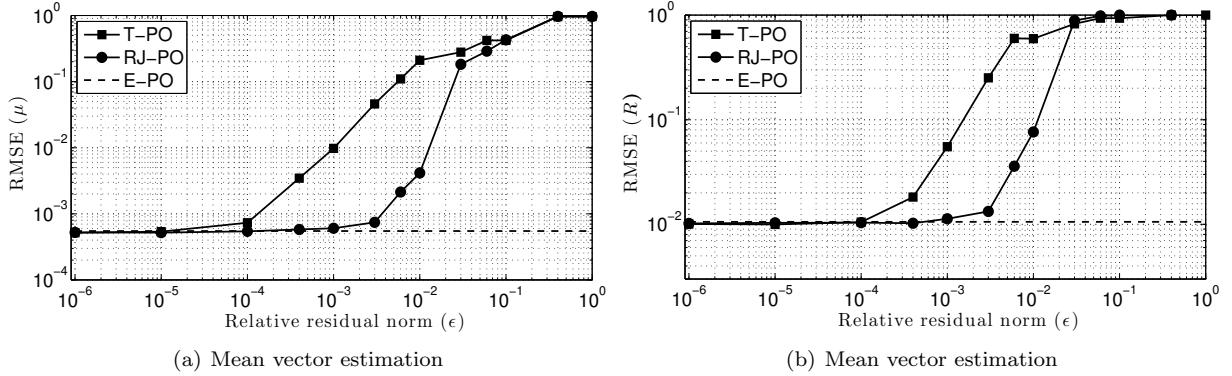


Figure 3: Estimation error for different values of the truncation level after $n_{\max} = 10^5$ iterations of E-PO, T-PO and RJPO algorithms: (a) mean vector, (b) covariance matrix.

4.4 Computation cost

Since the CG iterations correspond to the only burdensome task, the numerical complexity of the sampler can be expressed in terms of the total number J_{tot} of CG iterations to be performed before convergence and the number of required samples to get efficient empirical approximation of the estimators.

To assess the Markov chain convergence, we first use the Gelman-Rubin criterion based on multiple chains [31], which consists in computing a scale reduction factor based on the between and within-chain variances. In this experiment 100 parallel chains are considered. The results are summarized in Figure 4. It can be noted that a lower acceptance rate induces a higher number of iterations since the Markov chain converges more slowly towards its stationary distribution. One can also see that a minimal cost can be reached and, according to Figure 2, it corresponds to an acceptance rate of almost one. As the acceptance rate decreases, even a little, the computational cost rises very quickly. Conversely, if the relative residual is too small, the computation effort per sample will decrease but additional sampling iterations will be needed before convergence, which naturally increases the overall computation cost. The latter result points out the need to appropriately choose the truncation level to jointly avoid a low acceptance probability and a high resolution accuracy of the linear system since both induce unnecessary additional computations.

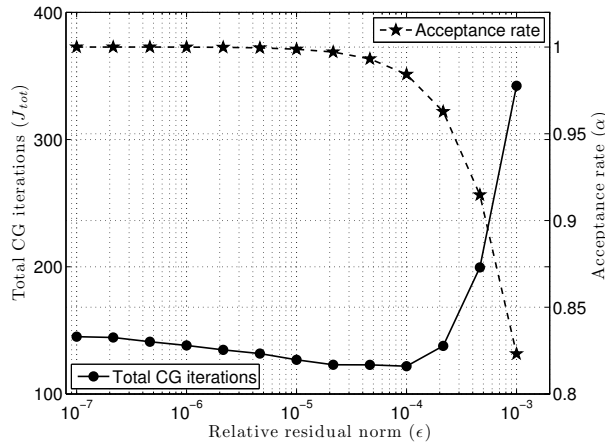


Figure 4: Number of CG iterations before convergence and acceptance probability of the RJPO algorithm for different values of relative residual norm for a small size problem ($N = 16$).

4.5 Statistical efficiency

The performance of the RJPO sampler can also be analyzed using the effective sample size (ESS) [32, p. 125]. This indicator gives the number of independent samples, n_{eff} , that would yield the same estimation variance of approximating the Bayesian estimator as n_{max} successive samples of the simulated chain [33]. It is related to the chain autocorrelation function according to

$$n_{\text{eff}} = \frac{n_{\text{max}}}{1 + 2 \sum_{k=1}^{\infty} \rho_k} \quad (14)$$

where ρ_k the autocorrelation coefficient at lag k . In the Gaussian sampling context, such a relation allows to define how many iterations n_{max} are needed for each resolution accuracy to get chains having the same effective sample size. Under the hypothesis of a first-order autoregressive chain, $\rho_k = \rho^k$, so (14) leads to the ESS ratio

$$\text{ESSR} = \frac{n_{\text{eff}}}{n_{\text{max}}} = \frac{1 - \rho}{1 + \rho}. \quad (15)$$

It can be noted that the ESSR is equal to one when the samples are independent ($\rho = 0$) and decreases as the correlation between successive samples grows. In the RJPO case, we propose to define the *computing cost per effective sample* (CCES) as

$$\text{CCES} = \frac{J_{\text{tot}}}{n_{\text{eff}}} = \frac{J}{\text{ESSR}} \quad (16)$$

where $J = J_{\text{tot}}/n_{\text{max}}$ is the average number of CG iterations per sample. Figure 5 shows the ESSR and the CCES in the case of a Gaussian vector of dimension $N = 16$. It can be seen that an early stopped CG algorithm induces a very small ESSR, due to a large sample correlation value, and thus a high effective cost to produce accurate estimates. On the contrary, a very precise resolution of the linear system induces a larger number of CG iterations per sample but a shorter Markov chain since the ESSR is almost equal to 1. The best trade-off is produced by intermediate values of the relative residual norm around $\epsilon = 2 \cdot 10^{-4}$.

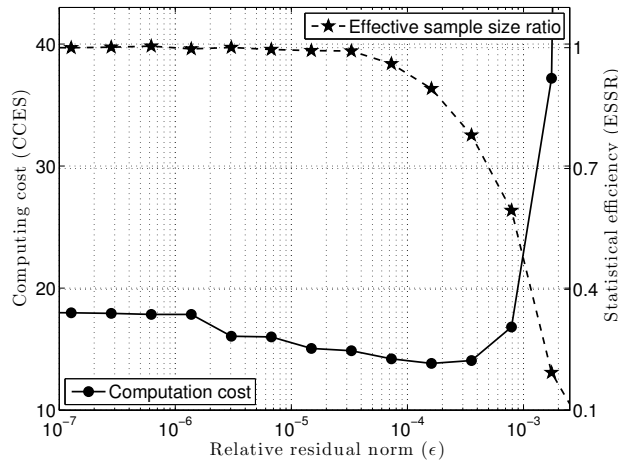


Figure 5: Computing cost per effective sample of the RJPO algorithm for different relative residual norm values on a small size problem ($N = 16$) estimated from $n_{\text{max}} = 10^4$ samples.

To conclude, the Gelman-Rubin convergence diagnostic and the ESS approach both confirm that the computation cost of the RJPO can be reduced by appropriately truncating the CG iterations. Although the Gelman-Rubin convergence test is probably more accurate, since it is based on several independent chains, the CCES based test is far simpler and provides nearly the same trade-off in the tested example. Such results motivate the development of an adaptive strategy to automatically adjust the threshold parameter ϵ by tracking the minimizer of the CCES. The proposed strategy is presented in Subsection 4.6.

4.5.1 Influence of the dimension

Figure 6 summarizes the optimal values of the truncation level ϵ that allows to minimize the CCES for different values of N . The best trade-off is reached for decreasing values of ϵ as N grows. More generally, the same observation can be made as the problem conditioning deteriorates. In practice, predicting the appropriate truncation level for a given problem is difficult. Fortunately, Figure 6 also indicates that the optimal setting is obtained for an acceptance probability that remains almost constant. The best trade-off is clearly obtained for an acceptance rate α lower than one ($\alpha = 1$ corresponds to $\epsilon = 0$, *i.e.*, to the exact solving of $Qx = \eta$). In the tested example, the optimal truncation level ϵ rather corresponds to an acceptance rate around 0.99. However, finding an explicit mathematical correspondence between ϵ and α is not a simple task. In the next subsection, we propose an unsupervised tuning strategy of the relative residual norm allowing either to achieve a predefined target acceptance rate, or even to directly optimize the computing cost per effective sample.




Figure 6: Influence of the problem dimension on the optimal values of the relative residual norm and the acceptance rate.

4.6 Adaptive tuning of the resolution accuracy

The suited value of the relative residual norm ϵ to achieve a desired acceptance rate α_t can be adjusted recursively using a Robbins-Monro type algorithm [34]. The result sampling approach relies on the family of controlled/adaptive MCMC methods [35]. See for instance [27] for a tutorial. Such an adaptive scheme is formulated in the stochastic approximation framework [36] in order to solve a non-linear equation of the form $g(\theta) = 0$ using an update

$$\theta_{n+1} = \theta_n + K_n [g(\theta_n) + \nu_n] \tag{17}$$

where ν is a random variable traducing the uncertainty on each evaluation of function $g(\cdot)$ and $\{K_n\}$ is a sequence of step-sizes ensuring stability and convergence [37]. Such a procedure has been already used for the optimal scaling of adaptive MCMC algorithms [27]. The use of Robbins-Monro procedure for the optimal scaling of some adaptive MCMC algorithms such as the Random walk Metropolis-Hastings (RWMH) algorithm. It is mainly shown that such procedure breakdown the Markovian structure of the chain but it does not alter its convergence towards the target distribution. For instance, it is used in [38, 39] to set adaptively the scale parameters of a RWMH algorithm in order to reach the optimal acceptance rate suggested by theoretical or empirical analysis [40, 41]. The same procedure was also used by [42] for the adaptive tuning of a Metropolis-adjusted Langevin algorithm (MALA) to reach the optimal acceptance rate proposed by [43].

4.6.1 Achieving a target acceptance rate

In order to ensure the positivity of the relative residual norm ϵ , the update is performed on its logarithm. At each iteration n of the sampler, the relative residual norm is **adjusted** according to

$$\log \epsilon_{n+1} = \log \epsilon_n + K_n [\alpha(\underline{\mathbf{x}}_n, \mathbf{x}_n) - \alpha_t]. \quad (18)$$

where α_t is a given target acceptance probability and $\{K_n\}$ is a sequence of step-sizes decaying to 0 as n grows in order to ensure the convergence of the Markov chain to the target distribution. As suggested in [27], the step-sizes are chosen according to $K_n = K_0/n^\beta$, with $\kappa \in]0, 1]$. We emphasize that more sophisticated methods, such as those proposed in [34] could be used to approximate the acceptance rate curve and to derive a more efficient adaptive strategy for choosing this parameter.

The adaptive RJPO is applied to the sampling of the previously described Gaussian distribution using the adopted step-size with parameters $K_0 = 1$ and $\kappa = 0.5$. Figure 7 presents the evolution of the average acceptance probability and the obtained relative residual norm for three different values of the target acceptance rate α_t . One can note that the average acceptance rate converges to the desired value. Moreover, the relative residual norm also converges to the expected values according to Figure 2 (for example, the necessary relative residual norm to get an acceptance probability $\alpha_t = 0.8$ is equal to $1.5 \cdot 10^{-3}$). In

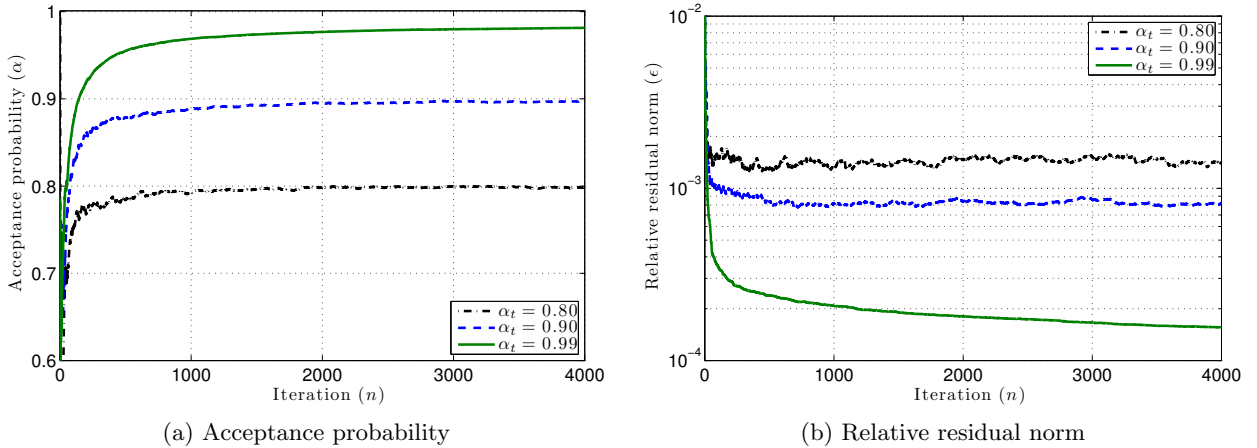


Figure 7: Behavior of the adaptive RJPO for 1000 iterations and three values of the target acceptance probability: (a) Evolution of the average acceptance probability and (b) Evolution of the computed relative residual norm.

practice, it remains difficult to *a priori* determine which acceptance rate should be targeted to achieve the faster convergence. The next subsection proposes to modify the target of the adaptive strategy to directly minimize the CCES (16).

4.6.2 Optimizing the numerical efficiency

A given threshold ϵ on the relative residual norm induces an average truncation level J and an ESSR value, from which the CCES can be deduced according to (16). Our goal is to adaptively adjust the threshold value ϵ in order to minimize the CCES. Let J_{opt} be the average number of CG iterations per sample corresponding to the optimal threshold value. In the plane (J, ESSR) , it is easy to see that J_{opt} is the abscissa of the point at which the tangent of the ESSR curve intercepts the origin (see Figure 8(a)).

The ESSR is expressed by (15) as a function of the chain correlation ρ , the latter being an implicit function of the acceptance rate α . For $\alpha = 1$, $\rho = 0$ according to Proposition 2. For $\alpha = 0$, $\rho = 1$ since no new sample can be accepted. For intermediate values of α , the correlation lies between 0 and 1, and it is typically decreasing. It can be decomposed on two terms:

- With a probability $1 - \alpha$, the accept-reject procedure produces identical (*i.e.*, maximally correlated) samples in case of rejection.
- In case of acceptance, the new sample is slightly correlated with the previous one, because of the early stopping of the CG algorithm.

While it is easy to express the correlation induced by rejection, it is difficult to find an explicit expression for the correlation between accepted samples. However, we have checked that the latter source of correlation is negligible compared to the correlation induced by rejection.

If we approximately assume that accepted samples are independent, we get $\rho = 1 - \alpha$, which implied

$$\text{ESSR} = \frac{2 - \alpha}{\alpha} \longrightarrow \text{CCES} = \frac{\alpha}{2 - \alpha}.$$

. Thus, by necessary condition, the best tuning of the relative residual norm leading to the lowest CCES is obtained by setting,

$$J \frac{d\alpha}{dJ} - \alpha + \frac{\alpha^2}{2} = 0.$$

Finally, the stochastic approximation procedure is applied to adaptively adjust the optimal value of the relative residual norm according to

$$\log \epsilon_{n+1} = \log \epsilon_n + K_n \left(J_n \frac{d\alpha_n}{dJ} - \alpha_n + \frac{\alpha_n^2}{2} \right), \quad (19)$$

where $\frac{d\alpha_n}{dJ}$ is evaluated numerically. Figure 8 illustrates that the proposed adaptive scheme efficiently adjusts ϵ to minimize the CCES, since J_{opt} is around 25 according to Figure 8(a).

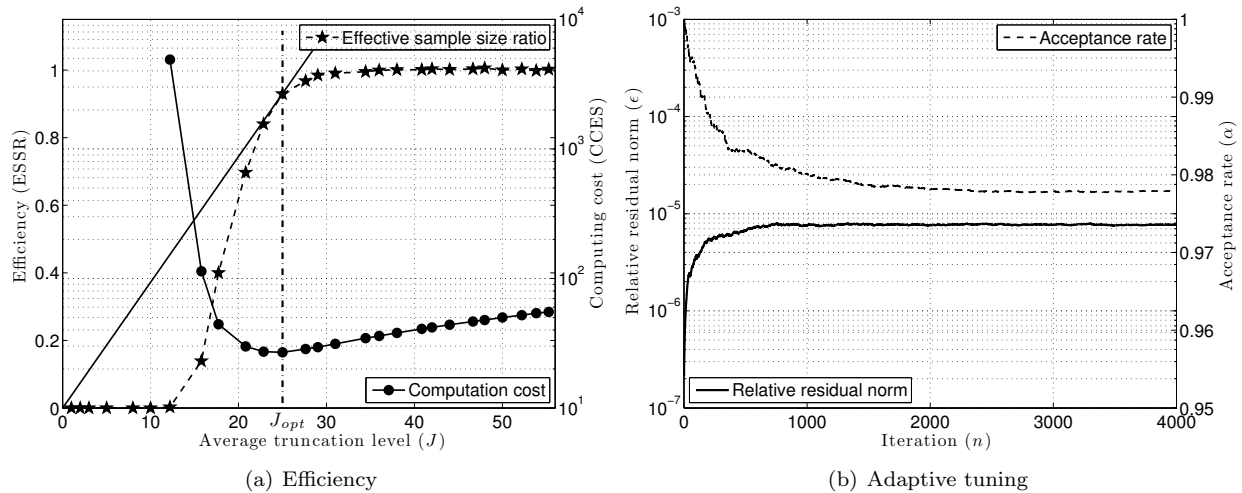


Figure 8: (a) Influence of the CG truncation level on the overall computation cost and the statistical efficiency of the RJPO for sampling a Gaussian of dimension $N = 128$. (b) Evolution of the relative residual norm and the acceptance rate for a Gaussian sampling problem of size $N = 128$. The adaptive algorithm leads to a relative residual norm $\epsilon_{\text{opt}} = 7.79 \cdot 10^{-6}$ leading to $\alpha_{\text{opt}} = 0.977$ and $J_{\text{opt}} = 26$.

5 Application to unsupervised super-resolution

In the linear inverse problem of unsupervised image super-resolution, several images are observed with a low spatial resolution. In addition, the measurement process presents a point spread function that introduces a blur on the images. The purpose is then to reconstruct the original image with a higher resolution using an unsupervised method. Such an approach allows to also estimate the model hyper-parameters and the PSF [23, 44, 45]. In order to discuss the relevance of the previously presented Gaussian sampling algorithms we apply a Bayesian approach and MCMC methods for solving this inverse problem.

5.1 Problem statement

The observation model is given by $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}$, where $\mathbf{H} = \mathbf{P}\mathbf{F}$, with $\mathbf{y} \in \mathbb{R}^M$ the vector containing the pixels of the observed images in a lexicographic order, $\mathbf{x} \in \mathbb{R}^N$ the sought high resolution image, \mathbf{F} the $N \times N$ circulant convolution matrix associated with the blur, \mathbf{P} the $M \times N$ decimation matrix and \mathbf{n} the additive noise.

Statistical modeling. The noise is assumed to follow a zero-mean Gaussian distribution with an unknown precision matrix $\mathbf{Q}_y = \gamma_y \mathbf{I}$. We also assume a zero-mean Gaussian distribution for the prior of the sought variable \mathbf{x} , with a precision matrix $\mathbf{Q}_x = \gamma_x \mathbf{D}^t \mathbf{D}$. \mathbf{D} is the circulant convolution matrix associated to a Laplacian filter. Non-informative Jeffrey's priors [46] are assigned to the two hyper-parameters γ_y and γ_x .

Bayesian inference. According to Bayes' theorem, the posterior distribution is given by

$$P(\mathbf{x}, \gamma_x, \gamma_y | \mathbf{y}) \propto \gamma_x^{(N-1)/2-1} \gamma_y^{M/2-1} \times e^{-\frac{1}{2}\gamma_y(\mathbf{y}-\mathbf{H}\mathbf{x})^t(\mathbf{y}-\mathbf{H}\mathbf{x}) - \frac{1}{2}\gamma_x \mathbf{x}^t \mathbf{D}^t \mathbf{D} \mathbf{x}}$$

To explore this posterior distribution, a Gibbs sampler iteratively draws samples from the following conditional distributions:

1. $\gamma_y^{(n)}$ from $P(\gamma_y | \mathbf{x}^{(n-1)}, \mathbf{y})$ given as

$$\mathcal{G}\left(1 + \frac{M}{2}, 2\|\mathbf{y} - \mathbf{H}\mathbf{x}^{(n-1)}\|^{-2}\right),$$

2. $\gamma_x^{(n)}$ from $P(\gamma_x | \mathbf{x}^{(n-1)})$ given as

$$\mathcal{G}\left(1 + \frac{N-1}{2}, 2\|\mathbf{D}\mathbf{x}^{(n-1)}\|^{-2}\right)$$

3. $\mathbf{x}^{(n)}$ from $P(\mathbf{x} | \gamma_x^{(n)}, \gamma_y^{(n)}, \mathbf{y})$ which is

$$\mathcal{N}\left(\boldsymbol{\mu}^{(n)}, [\mathbf{Q}^{(n)}]^{-1}\right)$$

with

$$\begin{aligned} \mathbf{Q}^{(n)} &= \gamma_y^{(n)} \mathbf{H}^t \mathbf{H} + \gamma_x^{(n)} \mathbf{D}^t \mathbf{D} \\ \mathbf{Q}^{(n)} \boldsymbol{\mu}^{(n)} &= \gamma_y^{(n)} \mathbf{H}^t \mathbf{y} \end{aligned}$$

The third step of the sampler requires an efficient sampling of a multivariate Gaussian distribution whose parameters change along the sampling iterations. In the sequel, direct sampling with Cholesky factorization [15] is firstly employed as a reference method. It yields the same results as the E-PO algorithm. For the inexact resolution case, the T-PO algorithm using a CG controlled by the relative residual norm, and the adaptive RJPO directly tuned with the acceptance probability are performed. For these two methods, the product matrix-vector used in the CG algorithm is done by exploiting the structure of the precision matrix \mathbf{Q} and thus only implies circulant convolutions, performed by FFT, and decimations.

5.2 Estimation results using MCMC

We consider the observation of five images of dimension 128×128 pixels ($M = 81920$) and we reconstruct the original one of dimension 256×256 ($N = 65536$). The convolution part \mathbf{F} has a Laplace shape with of full width at half maximum (FWHM) of 4 pixels. A white Gaussian noise is added to get a signal-to-noise ratio (SNR) equal to 20dB. The original image and one of the observations are shown in Figure 9.



Figure 9: Unsupervised super-resolution - image reconstruction with adaptive RJPO algorithm and $\alpha_t = 0.99$.

The Gibbs sampler is run for 1000 iterations and a burn-in period of 100 iterations is considered after a visual inspection of the chains. The performances are evaluated in terms of the mean and standard deviation of both hyper-parameters γ_y , γ_x and one randomly chosen pixel x_i of the reconstructed image. Table 1 presents the mean and standard deviation of the variable of interest. As we can see, the T-PO algorithm is totally inappropriate even with a precision of 10^{-8} . Conversely, the estimation from the samples given by the adaptive RJPO and Cholesky method are very similar, which demonstrates the correct behavior of the proposed algorithm.

| | γ_y | $\gamma_x \times 10^{-4}$ | x_i |
|---------------------------|--------------|---------------------------|--------------|
| Cholesky | 102.1 (0.56) | 6.1 (0.07) | 104.6 (9.06) |
| T-PO $\epsilon = 10^{-4}$ | 0.3 (0.06) | 45 (0.87) | 102.2 (3.30) |
| T-PO $\epsilon = 10^{-6}$ | 6.8 (0.04) | 32 (0.22) | 104.8 (2.34) |
| T-PO $\epsilon = 10^{-8}$ | 71.7 (0.68) | 21 (0.29) | 102.7 (2.51) |
| A-RJPO, $\alpha_t = 0.99$ | 101.2 (0.55) | 6.1 (0.07) | 101.9 (8.89) |

Table 1: Comparison between the Cholesky approach, the T-PO controlled by the relative residual norm and the A-RJPO tuned by the acceptance rate, in terms of empirical mean and standard deviation of hyper-parameters and one randomly chosen pixel.

Figure 10 shows the evolution of the average acceptance probability with respect to the number of CG iterations. We can notice that at least 400 iterations are required to have a nonzero acceptance rate. Moreover, more than 800 iterations seems unnecessary. For this specific problem, the E-PO algorithm needs theoretically $N = 65536$ iterations to have a new sample while the adaptive RJPO only requires around 700. Concerning the computation time, on a Intel Core i7-3770 with 8GB of RAM and a 64bit system, it took about 20.3s on average and about 6GB of RAM for the Cholesky sampler to generate one sample and only 15.1s and less than 200MB for the RJPO. This last result is due to the use of a conjugate gradient on which each matrix-vector product is performed without explicitly writing the matrix \mathbf{Q} . Finally, note that if we consider images of higher resolution, for instance $N = 1024 \times 1024$, the Cholesky factorization would require around 1TB of RAM and the Adaptive RJPO only about 3GB (when using double precision floating-point format).

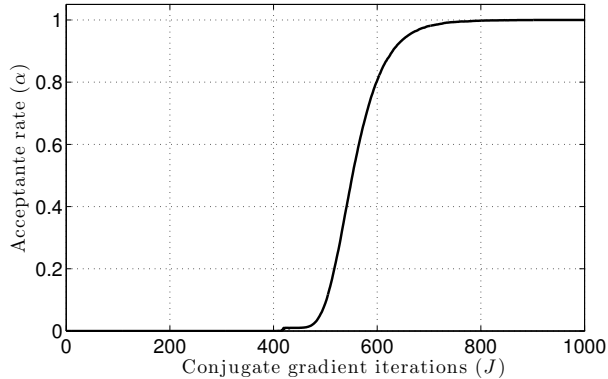


Figure 10: Evolution of the acceptance rate with respect to average conjugate gradient iterations for sampling a Gaussian of dimension $N = 65536$.

6 Conclusion

The sampling of high dimensional Gaussian distributions appears in the resolution of many linear inverse problems using MCMC methods. Alternative solutions to the Cholesky factorization are needed to reduce the computation time and to limit the memory usage. Based on the theory of reversible jump MCMC, we derived a sampling method allowing to introduce an approximate solution of a linear system during the sample generation step. The approximate resolution of a linear system was already adopted in methods like IFP and PO to reduce the numerical complexity, but without any guarantee of convergence to the target distribution. The proposed algorithm RJPO is based on an accept-reject step that is absent from the existing PO algorithms. Indeed, the difference between RJPO and existing PO algorithms is much comparable to the difference between the Metropolis-adjusted Langevin algorithm (MALA) [47] and a plainly discretized Langevin diffusion.

Our results pointed out that the required resolution accuracy in these methods must be carefully tuned to prevent a significant error. It was also shown that the proposed RJ-MCMC framework allows to ensure the convergence through the accept-reject step whatever the truncation level. In addition, thanks to the simplicity of the acceptance probability, the resolution accuracy can be adjusted automatically using an adaptive scheme allowing to achieve a pre-defined acceptance rate. We have also proposed a significant improvement of the same adaptive tuning approach, where the target is directly formulated in terms of minimal computing cost per effective sample.

Finally, the linear system resolution using the conjugate gradient algorithm offers the possibility to implement the matrix-vector products with a limited memory usage by exploiting the structure of the forward model operators. The adaptive RJPO has thus proven to be less consuming in both computational cost and memory usage than any approach based on Cholesky factorization.

This work opens some perspectives in several directions. Firstly, preconditioned conjugate gradient or alternative methods can be envisaged for the linear system resolution with the aim to reduce the computation time per iteration. Such an approach will highly depend on the linear operator and the ability to compute a preconditioning matrix. A second direction concerns the connection between the RJ-MCMC framework and other sampling methods such as those based on Krylov subspace [19, 20], particularly with appropriate choices of the parameters \mathbf{A} , \mathbf{B} , \mathbf{b} and $\mathbf{f}(\cdot)$ defined in section 2. Another perspective of this work is to analyze more complex situations involving non-gaussian distributions with the aim to be able to formulate the perturbation step and to perform an approximate optimization allowing to reduce the computation cost. Finally, the proposed adaptive tuning scheme allowing to optimize the computation cost per effective sample could be generalized to other Metropolis adjusted sampling strategies.

A Expression of the acceptance probability

According to the RJ-MCMC theory, the acceptance probability is given by

$$\alpha(\underline{\mathbf{x}}, \mathbf{x} | \mathbf{z}) = \min \left(1, \frac{P_{\mathbf{X}}(\mathbf{x}) P_{\mathbf{Z}}(\mathbf{s} | \mathbf{x})}{P_{\mathbf{X}}(\underline{\mathbf{x}}) P_{\mathbf{Z}}(\mathbf{z} | \underline{\mathbf{x}})} |J_{\phi}(\underline{\mathbf{x}}, \mathbf{z})| \right),$$

with $\mathbf{s} = \mathbf{z}$ and $\mathbf{x} = -\underline{\mathbf{x}} + \mathbf{f}(\mathbf{z})$. The Jacobian determinant of the deterministic move is $|J_{\phi}(\underline{\mathbf{x}}, \mathbf{z})| = 1$. Since

$$P_{\mathbf{X}}(\mathbf{x}) \propto e^{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^t \mathbf{Q}(\mathbf{x} - \boldsymbol{\mu})},$$

and

$$P_{\mathbf{Z}}(\mathbf{z} | \underline{\mathbf{x}}) \propto e^{-\frac{1}{2}(\mathbf{z} - \mathbf{A}\underline{\mathbf{x}} - \mathbf{b})^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{A}\underline{\mathbf{x}} - \mathbf{b})},$$

the acceptance probability can be written as

$$\alpha(\underline{\mathbf{x}}, \mathbf{x} | \mathbf{z}) = \min \left(1, e^{-\frac{1}{2}\Delta S} \right)$$

with $\Delta S = \Delta S_1 + \Delta S_2$ and

$$\begin{aligned} \Delta S_1 &= (\mathbf{x} - \boldsymbol{\mu})^t \mathbf{Q}(\mathbf{x} - \boldsymbol{\mu}) - (\underline{\mathbf{x}} - \boldsymbol{\mu})^t \mathbf{Q}(\underline{\mathbf{x}} - \boldsymbol{\mu}), \\ &= \mathbf{x}^t \mathbf{Q} \mathbf{x} - 2\mathbf{x}^t \mathbf{Q} \boldsymbol{\mu} - \underline{\mathbf{x}}^t \mathbf{Q} \underline{\mathbf{x}} + 2\underline{\mathbf{x}}^t \mathbf{Q} \boldsymbol{\mu}. \\ \Delta S_2 &= (\mathbf{z} - \mathbf{A}\mathbf{x} - \mathbf{b})^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{A}\mathbf{x} - \mathbf{b}) - (\mathbf{z} - \mathbf{A}\underline{\mathbf{x}} - \mathbf{b})^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{A}\underline{\mathbf{x}} - \mathbf{b}), \\ &= \mathbf{x}^t \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A} \mathbf{x} - 2\mathbf{x}^t \mathbf{A}^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{b}) - \underline{\mathbf{x}}^t \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A} \underline{\mathbf{x}} + 2\underline{\mathbf{x}}^t \mathbf{A}^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{b}). \end{aligned}$$

Since $\mathbf{x} = -\underline{\mathbf{x}} + \mathbf{f}(\mathbf{z})$, we get

$$\begin{aligned} \Delta S_1 &= (\mathbf{x} - \underline{\mathbf{x}})^t \mathbf{Q}(\mathbf{f}(\mathbf{z}) - 2\boldsymbol{\mu}) \\ \Delta S_2 &= (\mathbf{x} - \underline{\mathbf{x}})^t (\mathbf{A}^t \mathbf{B}^{-1} \mathbf{A} \mathbf{f}(\mathbf{z}) - 2\mathbf{A}^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{b})) \end{aligned}$$

Finally

$$\Delta S = (\mathbf{x} - \underline{\mathbf{x}})^t [(\mathbf{Q} + \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A}) \mathbf{f}(\mathbf{z}) - 2(\mathbf{Q} \boldsymbol{\mu} + \mathbf{A}^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{b}))] = 2(\underline{\mathbf{x}} - \mathbf{x})^t \mathbf{r}(\mathbf{z}).$$

Finally, when the system is solved exactly, $\Delta S = 0$ and thus $\alpha(\underline{\mathbf{x}}, \mathbf{x} | \mathbf{z}) = 1$.

B Correlation between two successive samples

Since

$$\bar{\mathbf{x}} = -\underline{\mathbf{x}} + 2(\mathbf{Q} + \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A})^{-1}(\mathbf{Q} \boldsymbol{\mu} + \mathbf{A}^t \mathbf{B}^{-1}(\mathbf{z} - \mathbf{b}))$$

and \mathbf{z} is sampled from $\mathcal{N}(\mathbf{A}\underline{\mathbf{x}} + \mathbf{b}, \mathbf{B})$, we have

$$\bar{\mathbf{x}} = -\underline{\mathbf{x}} + 2(\mathbf{Q} + \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A})^{-1}(\mathbf{Q} \boldsymbol{\mu} + \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A} \underline{\mathbf{x}} + \mathbf{A}^t \mathbf{B}^{-1} \boldsymbol{\omega}_{\mathbf{B}})$$

with $\boldsymbol{\omega}_{\mathbf{B}}$ totally independent of $\underline{\mathbf{x}}$. One can firstly check that $\mathbb{E}[\underline{\mathbf{x}}] = \mathbb{E}[\bar{\mathbf{x}}] = \boldsymbol{\mu}$. Consequently, the correlation between two successive samples is given by

$$\mathbb{E}[(\bar{\mathbf{x}} - \boldsymbol{\mu})(\underline{\mathbf{x}} - \boldsymbol{\mu})^t] = \left(2(\mathbf{Q} + \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A})^{-1} \mathbf{A}^t \mathbf{B}^{-1} \mathbf{A} - \mathbf{I} \right) \mathbf{Q}^{-1}$$

which is zero if and only if $\mathbf{A}^t \mathbf{B}^{-1} \mathbf{A} = \mathbf{Q}$.

References

- [1] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*. London, UK: IOP Publishing Ltd., 1998.
- [2] J. Idier, *Bayesian approach to Inverse problems*. ISTE Ltd and John Wiley & Sons Inc, 2008.
- [3] B. Frieden, “Image enhancement and restoration,” in *Picture Processing and Digital Filtering*, vol. 6 of *Topics in Applied Physics*, pp. 177–248, New York, NY, USA: Springer-Verlag, 1975.
- [4] G. Demoment, “Image reconstruction and restoration: Overview of common estimation structure and problems,” *IEEE Trans. Acoust. Speech, Signal Processing*, vol. 37, pp. 2024–2036, dec 1989.
- [5] G. T. Gordon, R. and Herman, “Reconstruction of pictures from their projections,” *Communications of the ACM*, vol. 14, no. 12, pp. 759–768, 1971.
- [6] R. M. Lewitt and S. Matej, “Overview of methods for image reconstruction from projections in emission computed tomography,” *Proceedings of the IEE*, vol. 91, no. 1, pp. 1588–1611, 2003.
- [7] W. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice*. London, UK: Chapman & Hall, 1999.
- [8] C. Robert, *The Bayesian Choice*. Springer-Verlag, 2nd ed., 2001.
- [9] S. Geman and D. Geman, “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, pp. 721–741, Nov. 1984.
- [10] D. Andrews and C. Mallows, “Scale mixtures of normal distributions,” *J. R. Statist. Soc. B*, pp. 99–102, 1974.
- [11] F. Champagnat and J. Idier, “A connection between half-quadratic criteria and EM algorithms,” *IEEE Signal Processing Lett.*, vol. 11, no. 9, pp. 709–712, 2004.
- [12] G. Papandreou and A. Yuille, “Gaussian sampling by local perturbations,” in *Proceedings of NIPS*, 2010.
- [13] E. M. Scheuer and D. S. Stoller, “On the generation of normal random vectors,” *Technometrics*, vol. 4, no. 2, pp. 278–281, 1962.
- [14] D. R. Barr and N. L. Slezak, “A comparison of multivariate normal generators,” *Commun. ACM*, vol. 15, pp. 1048–1049, Dec. 1972.
- [15] H. Rue, “Fast sampling of Gaussian Markov random fields,” *J. R. Statist. Soc. B*, vol. 63, no. 2, pp. 325–338, 2001.
- [16] W. F. Trench, “An algorithm for the inversion of finite Toeplitz matrices,” *J. Soc. Indust. Appl. Math.*, vol. 12, no. 3, pp. 515–522, 1964.
- [17] D. Geman and C. Yang, “Nonlinear image recovery with half-quadratic regularization,” *IEEE Trans. Image Processing*, vol. 4, no. 7, pp. 932–946, 1995.
- [18] P. Lalanne, D. Prévost, and P. Chavel, “Stochastic artificial retinas: algorithm, optoelectronic circuits, and implementation,” *Applied Optics*, vol. 40, no. 23, pp. 3861–3876, 2001.
- [19] A. Parker and C. Fox, “Sampling Gaussian distributions in Krylov spaces with conjugate gradients,” *SIAM J. Sci. Comput.*, vol. 34, no. 3, pp. B312–B334, 2012.
- [20] E. Aune, J. Eidsvik, and Y. Pokern, “Iterative numerical methods for sampling from high dimensional Gaussian distributions,” *Statist. and Comp.*, pp. 1–21, 2013.

- [21] E. Chow and Y. Saad, “Preconditioned Krylov subspace methods for sampling multivariate Gaussian distributions,” *SIAM J. Sci. Comput.*, vol. 36, no. 2, pp. A588–A608, 2014.
- [22] Y. Amit and U. Grenander, “Comparing sweep strategies for stochastic relaxation,” *Journal of Multivariate Analysis*, vol. 37, no. 2, pp. 197 – 222, 1991.
- [23] F. Orieux, O. Féron, and J. Giovannelli, “Sampling high-dimensional Gaussian distributions for general linear inverse problems,” *IEEE Signal Processing Lett.*, vol. 19, no. 5, p. 251, 2012.
- [24] X. Tan, J. Li, and P. Stoica, “Efficient sparse Bayesian learning via Gibbs sampling,” in *Proceedings of ICASSP*, pp. 3634–3637, 2010.
- [25] P. J. Green, “Reversible jump Markov chain Monte Carlo computation and Bayesian model determination,” *Biometrika*, vol. 82, pp. 711–732, 1995.
- [26] R. Waagepetersen and D. Sorensen, “A tutorial on reversible jump MCMC with a view toward applications in QTL-mapping,” *Int. Statist. Rev.*, vol. 69, no. 1, pp. 49–61, 2001.
- [27] C. Andrieu and J. Thoms, “A tutorial on adaptive MCMC,” *Statist. and Comp.*, vol. 18, no. 4, pp. 343–373, 2008.
- [28] Y. Atchade, G. Fort, E. Moulines, and P. Priouret, “Adaptive Markov chain Monte Carlo: theory and methods,” in *Bayesian Time Series Models*, pp. 33–53, Cambridge Univ. Press., 2011.
- [29] P. De Forcrand, “Monte Carlo quasi-heatbath by approximate inversion,” *Phys. Rev. D*, vol. 59, no. 3, pp. 3698–3701, 1999.
- [30] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. Society for Industrial and Applied Mathematics, 1994.
- [31] A. Gelman and D. B. Rubin, “Inference from iterative simulation using multiple sequences,” *Statist. and Comp.*, pp. 457–472, 1992.
- [32] J. S. Liu, *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics, Springer, 2nd ed., 2008.
- [33] J. Goodman and A. D. Sokal, “Multigrid Monte Carlo method. Conceptual foundations,” *Physical Review D*, vol. 40, no. 6, 1989.
- [34] B. Bercu and P. Fraysse, “A Robbins–Monro procedure for estimation in semiparametric regression models,” *Annals Statist.*, vol. 40, no. 2, pp. 666–693, 2012.
- [35] C. Andrieu and C. Robert, “Controlled MCMC for optimal sampling.” Tech. Report No. 0125, Cahiers de Mathématiques du Ceremade, Université Paris-Dauphine, 2001.
- [36] A. Benveniste, M. Métivier, and P. Priouret, *Adaptive algorithms and stochastic approximations*. Springer Publishing Company, Incorporated, 2012.
- [37] C. Andrieu, E. Moulines, and P. Priouret, “Stability of stochastic approximation under verifiable conditions,” *SIAM J. Control Optimization*, vol. 44, no. 1, pp. 283–312, 2006.
- [38] H. Haario, E. Saksman, and J. Tamminen, “An adaptive metropolis algorithm,” *Bernoulli*, vol. 7, no. 2, pp. 223–242, 2001.
- [39] Y. F. Atchadé and J. S. Rosenthal, “On adaptive Markov chain Monte Carlo algorithms,” *Bernoulli*, vol. 11, no. 5, pp. 815–828, 2005.

- [40] G. O. Roberts, A. Gelman, and W. R. Gilks, “Weak convergence and optimal scaling of random walk Metropolis algorithms,” *The Annals of Applied Probability*, vol. 7, no. 1, pp. 110–120, 1997.
- [41] A. Gelman, G. O. Roberts, and W. R. Gilks, “Efficient Metropolis jumping rules,” in *Bayesian statistics 5*, pp. 599–607, Oxford University Press, 1996.
- [42] Y. F. Achadé, “An adaptive version for the metropolis adjusted langevin algorithm with a truncated drift,” *Methodology and Computing in Applied Probability*, vol. 8, no. 2, pp. 235–254, 2006.
- [43] G. O. Roberts and J. S. Rosenthal, “Optimal scaling of discrete approximations to Langevin’s diffusions,” *J. R. Statist. Soc. B*, vol. 60, no. 1, pp. 255–268, 1998.
- [44] S. C. Park, M. K. Park, and M. G. Kang, “Super-resolution image reconstruction: a technical overview,” *IEEE Trans. Signal Processing*, vol. 20, no. 3, pp. 21–36, 2003.
- [45] G. Rochefort, F. Champagnat, G. Le Besnerais, and J.-F. Giovannelli, “An improved observation model for super-resolution under affine motion,” *IEEE Trans. Image Processing*, vol. 15, no. 11, pp. 3325–3337, 2006.
- [46] H. Jeffreys, “An invariant form for the prior probability in estimation problems,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 186, no. 1007, pp. 453–461, 1946.
- [47] G. O. Roberts and R. L. Tweedie, “Exponential convergence of Langevin distributions and their discrete approximations,” *Bernoulli*, vol. 2, no. 4, pp. 341–363, 1996.