Metropolized Randomized Maximum Likelihood
for sampling from multimodal distributions

Dean S. Oliver

Uni Research, Uni CIPR

Bergen, Norway

July 13, 2016

Abstract

This article describes a method for using optimization to derive efficient independent transition functions for Markov chain Monte Carlo (MCMC) simulations with application to sampling from a posterior density $\pi(\mathbf{x})$ which is multimodal. Although the proposals in the Randomized Maximum Likelihood method are placed in regions of high probability, the distribution of samples is only approximately correct. Introduction of auxiliary variables simplifies the computation of the proposal density, allowing the Metropolis-Hastings method to be applied. We restrict our attention to the special case for which the target density is the product of a multivariate Gaussian prior and a likelihood function for which the errors in observations are additive and Gaussian.

Introduction

Many geoscience inverse problems are characterized by nonlinearity in the relationship between parameters and observations [35, 7, 26]. In some cases, the posterior pdf for the model parameters after assimilation of observations is multimodal, although the presence or absence of multiple modes in large models is seldom verified. Examples of smaller models in which multiple modes have been verified include flow and transport problems in porous media, where the nonlinearity resulting from uncertainty in connections in layered media and observation operators that average over layers can both lead to multiple modes [9, 27]. Monte Carlo sampling is often the only viable method of quantifying uncertainty in the posterior distribution for problems of this type, but the cost of obtaining a substantial number of independent samples can be prohibitive.

In order to obtain a reasonably high probability of generating acceptable random-walk transitions in Metropolis sampling, the transitions must generally be small: for small enough transition distances, almost all proposals will be accepted. In that case, however, it

is possible to remain trapped in some modes for very long times before jumping to another mode and the time required to obtain useful Monte Carlo estimates can be impractical [15]. The problem of optimal scaling of the proposal distribution to maximize the efficiency of mixing of a random walk Metropolis algorithm has been solved for certain classes of target distributions [30, 31]. These scaling rules are not applicable in all contexts, and in particular are not appropriate for multimodal distributions [14]. The ability to design a proposal distribution that mixes well is key to the efficiency of the MCMC algorithm for multimodal distributions.

In this paper, we discuss an augmented variable independence Metropolis sampler that uses minimization to place proposals in regions of high probability. Augmented variable methods have been shown to be useful in MCMC to construct efficient sampling algorithms and in particular for sampling from multimodal distributions [4], where the auxiliary variables might allow for large jumps between modes [37]. Storvik [33] suggests that auxiliary variables can be used in Metropolis-Hastings (MH) either for generating better proposals or for easier calculation of acceptance probabilities and points out that the use of auxiliary variables allows flexibility in choosing the target distribution in the augmented space, as long as the marginal distribution for the variables of interest is unchanged. In this paper, the augmented variables are useful for obtaining a marginal proposal density that is close to the target marginal density, and for simplifying evaluation of the MH ratio.

Independence Metropolis samplers are Metropolis-Hastings samplers for which the transition probability $q(\mathbf{y}|\mathbf{x})$ does not depend on the current state of the chain ($q(\mathbf{y}|\mathbf{x}) = q(\mathbf{y})$) [36]. For independence sampling, it is useful to choose the proposal density $q(\mathbf{x})$ such that $\pi(\mathbf{x})/q(\mathbf{x})$ is bounded and as nearly constant as possible, in much the same way that optimal proposal densities would be chosen for importance sampling or for rejection sampling. Liu [20] compares the efficiency of the three methods, concluding that the independence Metropolis sampling is asymptotically comparable to rejection sampling, but that independence Metropolis sampling is simpler to implement as it does not require knowledge of the envelope constant.

A number of methods have been developed for MH that use either local gradient information to make transitions to regions of high probability [17, 22] or to allow better exploration of the target density than can be obtained by random-walk sampling [12, 5, 11]. In several methods, the transition is based on an optimization step. The Multiple-Try method [21] uses local minimization along a line in state space to propose candidate states for updating one of the current population of states. Alternatively, Tjelmeland and Hegstad [38] described a two-step transition in the Metropolis-Hastings method in which mode-jumping transitions alternated with mode-exploring transitions.

The Randomized Maximum Likelihood (RML) method [25] was developed as an independence Metropolis sampler for the special case in which the target density $\pi(\mathbf{x})$ is the product of a multivariate Gaussian prior $p(\mathbf{x})$ and a likelihood function $p(\mathbf{d}|\mathbf{x})$ for which

the errors in observations $\mathbf{d}$ are Gaussian. The prior distribution for $\mathbf{X}$ is assumed to be multivariate normal with mean $\boldsymbol{\mu}$ and covariance $\mathbf{C}_x$. Observations $\mathbf{d}_{\text{obs}} = \mathbf{g}(\mathbf{x}) + \boldsymbol{\epsilon}_d$ with $\boldsymbol{\epsilon}_d \sim N(0, \mathbf{C}_d)$ are assimilated resulting in a posterior density

$$\pi(\mathbf{x}) \propto \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^{\text{T}}\mathbf{C}_x^{-1}(\mathbf{x} - \boldsymbol{\mu}) - \frac{1}{2}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\text{obs}})^{\text{T}}\mathbf{C}_d^{-1}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\text{obs}})\right). \quad (1)$$

Candidate samples in RML are obtained by a two-step process. In the first step, randomized samples from the prior distribution for model variables and data variables are generated. In the second step, the candidate state is obtained by minimization of a stochastic objective function. The distribution of candidate states $q(\mathbf{x})$ depends on both the model and data variables so computation of the transition probability requires marginalization, $q(\mathbf{x}) = \int q(\mathbf{x}, \mathbf{d})\,d\mathbf{d}$. For problems in which the prior is Gaussian, the observation operator is linear, and errors in observations are additive and Gaussian, the proposal density $q(\mathbf{x})$ is equal to the target density $\pi(\mathbf{x})$ so that all independently proposed candidate states are accepted in Metropolis-Hastings. For problems in which the observation operator is nonlinear, proposed candidates are always placed in regions of high probability density, but evaluation of the Metropolis-Hastings acceptance test is difficult as it requires evaluation of the marginal density for model variables. Consequently, in practice the MH test is ignored for RML in large geoscience inverse problems [13, 16]. Bardsley et al. [2] showed that for non-linear problems in which the Jacobian of the observation operator is differentiable and monotonic, the MH acceptance test could be based on a Jacobian evaluated at the maximizer of Eq. 1. In this case, the probability of candidate states could be relatively easily computed and sampling was shown to be correct for several test problems.

Here we describe a modified Metropolization approach in which the randomized data are included in the state vector. The proposal of model variables is accomplished via a local optimization, but calibrated data variables are retained in the state space and used for evaluation of the probability of acceptance of the state. By doing this, the method allows sampling from multimodal distributions and the need to compute the marginal density for model variables is avoided. Correct sampling requires determination of the distribution of the proposals which is easier when the state space is augmented with data variables.

The paper is organized as follows. In Section 1 we summarize the original RML algorithm as an independence Metropolis sampler for which only model variables are included in the state of the Markov chain and discuss a pseudo-marginal MCMC approach to the sampling problem. In Section 2 we describe an augmented state version of the RML independence Metropolis sampler. For this version, computation of the Metropolis-Hastings ratio is simpler as it only requires the Jacobian of the transformation of the augmented state variables. Section 3 provides several examples for validation of the algorithm.

1 RML for model variables

In both the quasi-linear estimation method [19] and the randomized maximum likelihood method [25], samples are generated in high probability regions of the posterior pdf by the simple expedient of simultaneously minimizing the magnitude of misfit of the model with perturbed observations and the magnitude of the change in the model variables from an unconditional sample. Unconditional realizations of models variables ($\mathbf{x}_{\text{uc}}$) are sampled from the prior and realizations of observation ($\mathbf{d}_{\text{uc}}$) are sampled from the distribution of model noise,

$$\mathbf{x}_{\text{uc}} \sim N[\boldsymbol{\mu}, \mathbf{C}_x] \quad \text{and} \quad \mathbf{d}_{\text{uc}} \sim N[\mathbf{d}_{\text{obs}}, \mathbf{C}_d].$$

A sample from a high probability region is generated by minimizing a nonlinear least-squares cost function:

$$\mathbf{x}_* = \operatorname*{argmin}_{\mathbf{x}} \left[\frac{1}{2}(\mathbf{x} - \mathbf{x}_{\text{uc}})^{\text{T}} \mathbf{C}_x^{-1} (\mathbf{x} - \mathbf{x}_{\text{uc}}) + \frac{1}{2}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\text{uc}})^{\text{T}} \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\text{uc}}) \right]. \quad (2)$$

Although the samples generated in this way tend to be located in parts of parameter space with high posterior probability, the samples are not necessarily distributed according to the posterior distribution. The distribution of RML samples can, however, be computed if the inverse transformation from the calibrated samples to the prior samples is known. If the objective function (Eq. 2) is differentiable, then a necessary (but not sufficient) condition for $\mathbf{x}_*$ to be a minimizer is that

$$\mathbf{x}_{\text{uc}} = \mathbf{x}_* + \mathbf{C}_x \mathbf{G}^{\text{T}} \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{x}_*) - \mathbf{d}_{\text{uc}}) \quad (3)$$

where $\mathbf{G}^{\text{T}} = \nabla_x \mathbf{g}$. Eq. 3 provides the inverse transformation for the unconditional samples, restricted to the region of $(\mathbf{x}_*, \mathbf{d}_{\text{uc}})$ for which $\mathbf{x}_*$ is a minimizer of Eq. 2.

The joint distribution for $\mathbf{x}_{\text{uc}}$ and $\mathbf{d}_{\text{uc}}$ is

$$f(\mathbf{x}_{\text{uc}}, \mathbf{d}_{\text{uc}}) \propto \exp\left(-\frac{1}{2}(\mathbf{x}_{\text{uc}} - \boldsymbol{\mu})^{\text{T}} \mathbf{C}_x^{-1} (\mathbf{x}_{\text{uc}} - \boldsymbol{\mu}) - \frac{1}{2}(\mathbf{d}_{\text{uc}} - \mathbf{d}_{\text{obs}})^{\text{T}} \mathbf{C}_d^{-1} (\mathbf{d}_{\text{uc}} - \mathbf{d}_{\text{obs}})\right)$$

and hence the joint probability of proposing $(\mathbf{x}_*, \mathbf{d}_{\text{uc}})$ is

$$q(\mathbf{x}_*, \mathbf{d}_{\text{uc}}) = f(\mathbf{x}_{\text{uc}}(\mathbf{x}_*, \mathbf{d}_{\text{uc}}), \mathbf{d}_{\text{uc}}) |\mathbf{J}| \quad (4)$$

where $\mathbf{J}$ is the Jacobian of the inverse transformation from $(\mathbf{x}_*, \mathbf{d}_*)$ to $(\mathbf{x}_{\text{uc}}, \mathbf{d}_{\text{uc}})$. By choosing $\mathbf{d}_* = \mathbf{d}_{\text{uc}}$, the Jacobian determinant requires only computation of

$$|\mathbf{J}| = \left| \frac{\partial(\mathbf{x}_{\text{uc}}, \mathbf{d}_{\text{uc}})}{\partial(\mathbf{x}_*)} \right|.$$

The marginal probability of proposing $\mathbf{x}_*$ is then obtained by integrating $q(\mathbf{x}_*, \mathbf{d}_{\text{uc}})$ over the data space,

$$q(\mathbf{x}_*) = \int_D f(\mathbf{x}_{\text{uc}}(\mathbf{x}_*, \mathbf{d}_{\text{uc}}), \mathbf{d}_{\text{uc}}) |\mathbf{J}(\mathbf{x}_*, \mathbf{d}_{\text{uc}})| \, d\mathbf{d}_{\text{uc}}. \quad (5)$$

For an independence sampler in the Metropolis-Hastings algorithm [36], the probability of proposing a transition to state $\mathbf{x}_*$ is independent of the current state $\mathbf{x}$, then the proposed state $\mathbf{x}_*$ is accepted with probability

$$\alpha(\mathbf{x}, \mathbf{x}_*) = \min\left(1, \frac{\pi(\mathbf{x}_*)q(\mathbf{x})}{\pi(\mathbf{x})q(\mathbf{x}_*)}\right). \tag{6}$$

The probability density for the proposed model, $\pi(\mathbf{x}_*)$, is computed from Eq. 1. Note that the probability π is not based on the quality of the match to the perturbed data obtained in the minimization, but on the quality of the match to the prior model and the actual observed data. Algorithm 1 describes the method for using RML as a proposal mechanism in a Metropolis-Hastings procedure [25].

Algorithm 1 RML for model variables

1: Generate initial state $\mathbf{x}_0 \sim N[\boldsymbol{\mu}, \mathbf{C}_x]$
2: **for** $i \leq i_{\max}$ **do**
3:      **procedure** GENERATE CANDIDATE STATE
4:          Generate $\mathbf{x}_{uc} \sim N[\boldsymbol{\mu}, \mathbf{C}_x]$ and $\mathbf{d}_{uc} \sim N[\mathbf{d}_{\mathrm{obs}}, \mathbf{C}_d]$
5:          Compute
6:          $\mathbf{x}_* = \mathrm{argmin}_{\mathbf{x}}\left[\frac{1}{2}(\mathbf{x} - \mathbf{x}_{\mathrm{uc}})^{\mathrm{T}}\mathbf{C}_x^{-1}(\mathbf{x} - \mathbf{x}_{\mathrm{uc}}) + \frac{1}{2}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{uc}})^{\mathrm{T}}\mathbf{C}_d^{-1}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{uc}})\right]$
7:      **end procedure**
8:      Compute proposal density $q(\mathbf{x}_*) = \int_D q(\mathbf{x}_*, \mathbf{d}_{\mathrm{uc}})\, d\mathbf{d}_{\mathrm{uc}}$ using Eq. 4 for $q(\mathbf{x}_*, \mathbf{d}_{\mathrm{uc}})$.
9:      Compute $\alpha(\mathbf{x}_i, \mathbf{x}_*) = \min\left(1, \frac{\pi(\mathbf{x}_*)q(\mathbf{x}_i)}{\pi(\mathbf{x}_i)q(\mathbf{x}_*)}\right)$
10:     Generate u from $U(0,1)$
11:     **if** $u \leq \alpha(\mathbf{x}_i, \mathbf{x}_*)$ **then**
12:         $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_*$
13:     **else**
14:         $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i$
15:     **end if**
16: **end for**

When the relationship of model variables to observations is linear and observation errors are additive and Gaussian, it is straightforward to show that the minimizers of Eq. 2 are distributed as the posterior distribution for $\mathbf{x}$ [3, 23] and hence all proposed transitions are accepted when Algorithm 1 is used for Gauss-linear inverse problems.

When Algorithm 1 was applied to a problem for which the posterior distribution was bimodal (as in Example 1, below), the acceptance rate for proposed independent transitions was still very high (76%) [25]. Although the method was shown to sample correctly, the work required for computing the marginal distribution for the candidate states made application of the full method impractical in real problems. Consequently, when the method has been applied to large problems, the MH acceptance test was omitted [6, 8, 16], with the understanding that the sampling method is then only approximately correct.

The need to numerically compute the marginal density $q(\mathbf{x}_*)$ from $q(\mathbf{x}_*, \mathbf{d}_{\mathrm{uc}})$ in Algorithm 1 could potentially be avoided through application of the pseudo-marginal MCMC method [1, 32], which only requires an unbiased estimate of the marginal distribution. When the pseudo-marginal MCMC method as described in Algorithm 9.7 of [29] was applied to Example 1 (Section 3.1) using a using a single unbiased sample at each step, however, sampling from a Markov chain of length 10^5 was very poor. Results would almost certainly be improved by increasing the sample size used to represent the marginal distribution [32], but because of the dependence of x_* on d_{uc} in the RML method, it is not feasible to draw multiple d_{uc} for a given x_*.

2 Augmented state RML

Because the major challenge with the use of Metropolized RML as in Algorithm 1 is the computation of the marginal probability of proposing the candidate model variables, we introduce a modification that avoids the need for computation of the marginal proposal density by augmenting the state $\mathbf{x}_*$ with suitably defined data variables $\mathbf{d}_*$. Although auxiliary variables can improve sampling in several ways [33], the purpose of including auxiliary variables in this case is primarily to simplify computation of the MH acceptance criterion.

2.1 Posterior (target) distribution

Consider the joint space of model variables and data in the case for which the observation errors and the modelization errors can both be modeled as Gaussian. The prior distribution of model variables is assumed Gaussian with mean $\boldsymbol{\mu}$ and covariance $\mathbf{C}_x$; the prior distribution of observations is assumed to be Gaussian with mean $\mathbf{d}_{\mathrm{obs}}$ and covariance $(1 - \gamma)\mathbf{C}_d$; modelization errors are Gaussian with mean 0 and covariance $\gamma\mathbf{C}_d$. The posterior joint probability of $(\mathbf{x}, \mathbf{d})$, obtained by combining states of information is

$$\pi(\mathbf{x}, \mathbf{d}) \propto \exp\Big[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^{\mathrm{T}}\mathbf{C}_x^{-1}(\mathbf{x} - \boldsymbol{\mu}) - \frac{1}{2\gamma}(\mathbf{g}(\mathbf{x}) - \mathbf{d})^{\mathrm{T}}\mathbf{C}_d^{-1}(\mathbf{g}(\mathbf{x}) - \mathbf{d})$$
$$-\frac{1}{2(1 - \gamma)}(\mathbf{d} - \mathbf{d}_{\mathrm{obs}})^{\mathrm{T}}\mathbf{C}_d^{-1}(\mathbf{d} - \mathbf{d}_{\mathrm{obs}})\Big]. \quad (7)$$

After some rearrangement of terms, the joint probability can be factored into the product of the marginal posterior density for model variable $\mathbf{x}$ and the posterior density for $\mathbf{d}$ conditional to $\mathbf{x}$,

$$\pi(\mathbf{x}, \mathbf{d}) \propto \exp\Big[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^{\mathrm{T}}\mathbf{C}_x^{-1}(\mathbf{x} - \boldsymbol{\mu}) - \frac{1}{2}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{obs}})^{\mathrm{T}}\mathbf{C}_d^{-1}(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{obs}})\Big]$$
$$\exp\Big[-\frac{1}{2\gamma(1 - \gamma)}\left(\mathbf{d} - \mathbf{g}(\mathbf{x}) + \gamma(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{obs}})\right)^{\mathrm{T}}\mathbf{C}_d^{-1}\left(\mathbf{d} - \mathbf{g}(\mathbf{x}) + \gamma(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{obs}})\right)\Big]. \quad (8)$$

Although the introduction of Gaussian modelization error does not affect the posterior marginal distribution of $\mathbf{x}$ as long as it is compensated for by a reduction in observation

error, the posterior joint probability of $(\mathbf{x}, \mathbf{d})$ does depend strongly on the value of γ as discussed in Section 2.3 on the selection of γ.

2.2 Proposal distribution

As in the original version of RML, we draw samples from the normal distributions $\mathbf{x}_{\text{uc}} \sim N[\boldsymbol{\mu}, \mathbf{C}_x]$ and $\mathbf{d}_{\text{uc}} \sim N[\mathbf{d}_{\text{obs}}, \mathbf{C}_d]$ and denote the joint distribution as $p(\cdot, \cdot)$. Candidate transitions in a MH algorithm are obtained by minimizing a nonlinear least squares cost function

$$(\mathbf{x}_*, \mathbf{d}_*) = \underset{\mathbf{x}, \mathbf{d}}{\operatorname{argmin}} \left[\frac{1}{2} (\mathbf{x} - \mathbf{x}_{\text{uc}})^{\mathrm{T}} \mathbf{C}_x^{-1} (\mathbf{x} - \mathbf{x}_{\text{uc}}) + \frac{1}{2\rho} (\mathbf{g}(\mathbf{x}) - \mathbf{d})^{\mathrm{T}} \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{x}) - \mathbf{d}) \right.$$
$$\left. + \frac{1}{2(1 - \rho)} (\mathbf{d} - \mathbf{d}_{\text{uc}})^{\mathrm{T}} \mathbf{C}_d^{-1} (\mathbf{d} - \mathbf{d}_{\text{uc}}) \right]. \quad (9)$$

Implicit expressions for the relationship between $(\mathbf{x}_*, \mathbf{d}_*)$ and $(\mathbf{x}_{\text{uc}}, \mathbf{d}_{\text{uc}})$ are derived from requirement that the gradient at the minimum must vanish. At the minimum, the following relationships hold:

$$\mathbf{x}_* - \mathbf{x}_{\text{uc}} + \frac{1}{\rho} \mathbf{C}_x \mathbf{G}^{\mathrm{T}} \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{x}_*) - \mathbf{d}_*) = 0 \quad (10)$$

and

$$\mathbf{d}_* - \rho \, \mathbf{d}_{\text{uc}} - (1 - \rho) \mathbf{g}(\mathbf{x}_*) = 0. \quad (11)$$

Using Eq. 11 to eliminate $\mathbf{d}_*$ from Eq. 10 gives

$$\mathbf{x}_* - \mathbf{x}_{\text{uc}} + \mathbf{C}_x \mathbf{G}^{\mathrm{T}} \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{x}_*) - \mathbf{d}_{\text{uc}}) = 0 \quad (12)$$

which shows that the marginal distribution of $\mathbf{x}_*$ is independent of ρ. For Gauss-linear problems, $\mathbf{x}_*$ is distributed according to the posterior marginal distribution of $\mathbf{x}$.

The inverse mapping, obtained from Eqs. 10 and 11,

$$\begin{bmatrix} \mathbf{x}_{\text{uc}} \\ \mathbf{d}_{\text{uc}} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_* + \frac{1}{\rho} \mathbf{C}_x \mathbf{G}^{\mathrm{T}} \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{x}_*) - \mathbf{d}_*) \\ \frac{1}{\rho} \mathbf{d}_* - \left(\frac{1 - \rho}{\rho} \right) \mathbf{g}(\mathbf{x}_*) \end{bmatrix}. \quad (13)$$

is invertible if the domain of the mapping is restricted to the range of the mapping given by Eq. 9 which we denote by Γ. The distribution of candidate states, is then

$$q(\mathbf{x}_*, \mathbf{d}_*) = \begin{cases} p\left(\mathbf{x}_{\text{uc}}(\mathbf{x}_*, \mathbf{d}_*), \mathbf{d}_{\text{uc}}(\mathbf{x}_*, \mathbf{d}_*)\right) |\mathbf{J}(\mathbf{x}_*, \mathbf{d}_*)| & \text{if } (\mathbf{x}_*, \mathbf{d}_*) \in \Gamma \\ 0 & \text{if } (\mathbf{x}_*, \mathbf{d}_*) \notin \Gamma \end{cases} \quad (14)$$

where $\mathbf{x}_{\text{uc}}(\mathbf{x}_*, \mathbf{d}_*)$ and $\mathbf{d}_{\text{uc}}(\mathbf{x}_*, \mathbf{d}_*)$ are defined in Eq. 13 and the dependence of the expressions on ρ has been suppressed for clarity.

Because the state is composed of both model and data variables, the Jacobian matrix takes a block form with

$$\frac{\partial \mathbf{x}_{\text{uc}}^\alpha}{\partial \mathbf{x}_*^\beta} = I^{\alpha\beta} + \frac{1}{\rho} \sum_\gamma \sum_i \sum_j \mathbf{C}_x^{\alpha\gamma} \left[\mathbf{G}^{i\gamma} \left[\mathbf{C}_d^{-1} \right]^{ij} \mathbf{G}^{j\beta} + \frac{\partial \mathbf{G}^{i\gamma}}{\partial \mathbf{x}^\beta} \left[\mathbf{C}_d^{-1} \right]^{ij} (\mathbf{g}^j(\mathbf{x}_*) - \mathbf{d}_*^j) \right]$$

and

$$\mathbf{G}^{i\alpha} = \frac{\partial \mathbf{g}^i}{\partial \mathbf{x}^\alpha}.$$

Other derivatives required for the Jacobian can be compactly written in matrix notation.

$$\frac{\partial \mathbf{x}_{\mathrm{uc}}}{\partial \mathbf{d}_*} = -\frac{1}{\rho}\mathbf{C}_x \mathbf{G}^{\mathrm{T}} \mathbf{C}_d^{-1},$$

$$\frac{\partial \mathbf{d}_{\mathrm{uc}}}{\partial \mathbf{x}_*} = -\left(\frac{1-\rho}{\rho}\right)\mathbf{G}$$

and

$$\frac{\partial \mathbf{d}_{\mathrm{uc}}}{\partial \mathbf{d}_*} = \frac{1}{\rho}I.$$

For a linear Gaussian inverse problem, the relationship between the proposed states and the unconditional states is

$$\begin{bmatrix} \mathbf{x}_* \\ \mathbf{d}_* \end{bmatrix} = \begin{bmatrix} \mathbf{C}_d\left(\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}} + \mathbf{C}_d\right)^{-1} & \mathbf{C}_x\mathbf{G}^{\mathrm{T}}\left(\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}} + \mathbf{C}_d\right)^{-1} \\ (1-\rho)\mathbf{G}^{\mathrm{T}}\mathbf{C}_d\left(\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}} + \mathbf{C}_d\right)^{-1} & \left(\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}} + \rho\mathbf{C}_d\right)\left(\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}} + \mathbf{C}_d\right)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{\mathrm{uc}} \\ \mathbf{d}_{\mathrm{uc}} \end{bmatrix}$$

$$= \mathbf{A}\begin{bmatrix} \mathbf{x}_{\mathrm{uc}} \\ \mathbf{d}_{\mathrm{uc}} \end{bmatrix}.$$

The covariance of the proposed states $(\mathbf{x}_*, \mathbf{d}_*)$ is

$$\mathbf{C}_{\mathrm{post}} = \mathbf{A}\begin{bmatrix} \mathbf{C}_x & 0 \\ 0 & \mathbf{C}_d \end{bmatrix} \mathbf{A}^{\mathrm{T}}$$

$$= \begin{bmatrix} \mathbf{C}_{x'} & \mathbf{C}_{x'}\mathbf{G}^{\mathrm{T}} \\ \mathbf{G}\mathbf{C}_{x'} & \mathbf{G}\mathbf{C}_{x'}\mathbf{G}^{\mathrm{T}} + \rho^2\mathbf{C}_d\mathbf{A}_D^{-1}\mathbf{C}_d \end{bmatrix}$$

where $\mathbf{C}_{x'} = \left(\mathbf{C}_x^{-1} + \mathbf{G}^{\mathrm{T}}\mathbf{C}_d^{-1}\mathbf{G}\right)^{-1}$ and $\mathbf{A}_D = \left(\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}} + \mathbf{C}_d\right)$. Note that for $\rho = 0$, the covariance of proposed states is the same as the posteriori covariance for the joint model-data space of Tarantola [34], although the covariance is then rank deficient.

The joint density given by Eq. 14 can be used as the proposal density in an independence MH sampler for sampling of the joint posterior distribution (Eq. 7) with the objective of sampling from the posterior marginal distribution for model variables (Eq. 1). Algorithm 2 describes the use of minimization to provide candidate states.

2.3 Selection of ρ and γ

The efficiency of the independence Metropolis sampler depends on two factors. The first is requirement that the proposal distribution q (Eq. 14) be close to the target density π (Eq. 7) so that the ratio $\pi(\mathbf{x}_*, \mathbf{d}_*)/q(\mathbf{x}_*, \mathbf{d}_*)$ is as close to constant as possible [36]. The second requirement is loosely that the proposal density be positive wherever the target density is positive so that the Markov chain is irreducible.

Algorithm 2 Augmented state RML for model-data variables

1: Generate $\mathbf{x}_0 \sim N[\boldsymbol{\mu}, \mathbf{C}_x]$ and $\mathbf{d}_0 \sim N[\mathbf{d}_{\mathrm{obs}}, \mathbf{C}_d]$

2: **for** $i \leq i_{\max}$ **do**

3:     **procedure** GENERATE CANDIDATE STATE

4:        Generate $\mathbf{x}_{\mathrm{uc}} \sim N[\boldsymbol{\mu}, \mathbf{C}_x]$ and $\mathbf{d}_{\mathrm{uc}} \sim N[\mathbf{d}_{\mathrm{obs}}, \mathbf{C}_d]$

5:        Compute

6:        $(\mathbf{x}_*, \mathbf{d}_*) = \mathrm{argmin}_{\mathbf{x},\mathbf{d}} \left[\frac{1}{2}(\mathbf{x} - \mathbf{x}_{\mathrm{uc}})^{\mathrm{T}} \mathbf{C}_x^{-1}(\mathbf{x} - \mathbf{x}_{\mathrm{uc}}) \right.$

7:              $\left. + \frac{1}{2\rho}(\mathbf{g}(\mathbf{x}) - \mathbf{d})^{\mathrm{T}} \mathbf{C}_d^{-1}(\mathbf{g}(\mathbf{x}) - \mathbf{d}) + \frac{1}{2(1-\rho)}(\mathbf{d} - \mathbf{d}_{\mathrm{uc}})^{\mathrm{T}} \mathbf{C}_d^{-1}(\mathbf{d} - \mathbf{d}_{\mathrm{uc}}) \right]$

8:     **end procedure**

9:     Compute proposal density $q(\mathbf{x}_*, \mathbf{d}_*)$ using Eq. 14.

10:     Compute $\alpha(\mathbf{x}_i, \mathbf{d}_i, \mathbf{x}_*, \mathbf{d}_*) = \min\left(1, \frac{\pi(\mathbf{x}_*, \mathbf{d}_*)q(\mathbf{x}_i, \mathbf{d}_i)}{\pi(\mathbf{x}_i, \mathbf{d}_i)q(\mathbf{x}_*, \mathbf{d}_*)}\right)$

11:     Generate u from $U(0,1)$

12:     **if** $u \leq \alpha(\mathbf{x}_i, \mathbf{d}_i, \mathbf{x}_*, \mathbf{d}_*)$ **then**

13:        $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_*$ and $\mathbf{d}_{i+1} \leftarrow \mathbf{d}_*$

14:     **else**

15:        $\mathbf{x}_{i+1} \leftarrow \mathbf{x}_i$ and $\mathbf{d}_{i+1} \leftarrow \mathbf{d}_i$

16:     **end if**

17: **end for**

It is convenient to compare the proposal densities and the target densities based on the distribution of $\mathbf{d}$ given $\mathbf{x}$. From Eq. 8, the mode of the target density for $\mathbf{d}$ is located at

$$\mathbf{d} = \mathbf{g}(\mathbf{x}) - \gamma(\mathbf{g}(\mathbf{x}) - \mathbf{d}_{\mathrm{obs}}) \tag{15}$$

and the covariance of $\mathbf{d}$ given $\mathbf{x}$ in the target distribution is $\gamma(1 - \gamma)\mathbf{C}_d$. Similarly, it is possible to factor the proposal density Eq. 14 as the product of the target density for $\mathbf{x}$, a Gaussian distribution for $\mathbf{d}_*$ given $\mathbf{x}_*$, term that depends on $\mathbf{d}_*$, and the Jacobian determinant,

$$q(\mathbf{x}_*, \mathbf{d}_*) \propto \exp\left[-\frac{1}{2}(\mathbf{x}_* - \boldsymbol{\mu})^{\mathrm{T}} \mathbf{C}_x^{-1}(\mathbf{x}_* - \boldsymbol{\mu}) - \frac{1}{2}(\mathbf{g}(\mathbf{x}_*) - \mathbf{d}_{\mathrm{obs}})^{\mathrm{T}} \mathbf{C}_d^{-1}(\mathbf{g}(\mathbf{x}_*) - \mathbf{d}_{\mathrm{obs}})\right]$$

$$\times \exp\left[-\frac{1}{2\rho^2}\left(\mathbf{d}_* - \mathbf{g}(\mathbf{x}_*) + \rho\mathbf{V}^{-1}\boldsymbol{\eta}(\mathbf{x}_*)\right)^{\mathrm{T}} \mathbf{V}\left(\mathbf{d}_* - \mathbf{g}(\mathbf{x}_*) + \rho\mathbf{V}^{-1}\boldsymbol{\eta}(\mathbf{x}_*)\right)\right]$$

$$\times \exp\left[\frac{1}{2}\boldsymbol{\eta}(\mathbf{x}_*)^{\mathrm{T}}\mathbf{V}^{-1}\boldsymbol{\eta}(\mathbf{x}_*)\right] |\mathbf{J}(\mathbf{x}_*, \mathbf{d}_*)| \tag{16}$$

where

$$\mathbf{V} = \mathbf{C}_d^{-1} + \mathbf{C}_d^{-1}\mathbf{G}\mathbf{C}_x\mathbf{G}^{\mathrm{T}}\mathbf{C}_d^{-1} \tag{17}$$

and

$$\boldsymbol{\eta}(\mathbf{x}_*) = \mathbf{C}_d^{-1}\left[(\mathbf{G}\mathbf{x}_* - \mathbf{g}(\mathbf{x}_*)) - (\mathbf{G}\boldsymbol{\mu} - \mathbf{d}_{\mathrm{obs}})\right]. \tag{18}$$

Note that when the observation operator $\mathbf{g}$ is linear, $\boldsymbol{\eta}$, $\mathbf{V}$, and $\mathbf{J}$ are all independent of $\mathbf{x}_*$ and $\mathbf{d}_*$.

For linear observations, both the target density and the proposal density are Gaussian in which case it is sufficient to compare the modes and the variances of the two distributions for selection of ρ and γ. The modes of the target distribution and the mode of the proposal distribution in this case are equal when

$$-\rho \mathbf{V}^{-1} \boldsymbol{\eta} = \gamma (\mathbf{G}\mathbf{x}_* - \mathbf{d}_{\mathrm{obs}})$$

and the variances of the two distributions are equal when

$$\rho^2 \mathbf{V}^{-1} = \gamma(1 - \gamma)\mathbf{C}_d.$$

For linear g, it is therefore not possible to satisfy both relationships exactly, except by choosing $\rho = \gamma = 0$, in which case the distributions are degenerate. Figure 1 compares the joint distributions for several values of γ and ρ in the case of a linear observation. It is clear that the two distributions will only be similar if γ is small.

(a) Joint target pdf for three values of γ: 0.02 (black), 0.5 (dashed), 0.98 (red).

(b) Joint proposal pdf for three values of ρ: 0.05 (black), 0.5 (dashed), 0.95 (red).

Figure 1: Effect of the values of ρ and γ on the joint target and proposal densities for a linear observation.

For nonlinear problems, the mapping (Eq. 13) from $(\mathbf{x}_{\mathrm{uc}}, \mathbf{d}_{\mathrm{uc}})$ to $(\mathbf{x}_*, \mathbf{d}_*)$ is not invertible unless the domain of the mapping is restricted. In this case, it is necessary to select ρ and γ in such a way that the proposal density is concentrated in the region where the mapping is bijective (and hence the proposal density is positive). Insight can be obtained by considering the case for which a single model variable and a single data variable.

If $\partial^2 g / \partial x_*^2 \neq 0$, then solving for the location of the boundary, we obtain

$$d_* = g(x_*) + C_d \left(\frac{\partial^2 g}{\partial x_*^2}\right)^{-1} \left(\rho C_x^{-1} + G^{\mathrm{T}} C_d^{-1} G\right). \tag{19}$$

Increasing the value of ρ shifts the boundary away from $d_* = g(x_*)$ in the direction given by the sign of the second derivative of g.

The distance between the mode of the target density (Eq. 15) and the boundary of the region of positive proposal density (Eq. 19) is

$$\Delta d = C_d \left(\frac{d^2 g}{dx^2} \right) (\rho C_x^{-1} + G^{\mathrm{T}} C_d^{-1} G) + \gamma(g(x) - d_{\mathrm{obs}}).$$

The probability mass outside of the domain of validity is determined also by the variance of the distribution. The fraction of the cumulative distribution that is outside of the domain of $q(\mathbf{x}_*, \mathbf{d}_*)$ is less than ϵ if

$$\Phi \left(\frac{\Delta d}{\sqrt{\gamma(1 - \gamma)C_d}} \right) > 1 - \epsilon.$$

Thus for highly nonlinear problems, it is necessary to choose $\gamma \ll 1$ and $\rho \approx 1$ to satisfy the requirement that all states with significant probability density can be reached.

(a) One branch of $q(x, d)$ for $\rho = 0.5$.

(b) One branch of $q(x, d)$ for $\rho = 0.999$.

(c) One branch of $\pi(x, d)$ for two values of γ: 0.001 (black), 0.5 (red). Dashed blue contour shows boundary of proposal region.

Figure 2: Effect of the values of ρ and γ on the joint target and proposal densities for a quadratic observation.

Figure 2 illustrates the effect of the choice of ρ on the proposal density for the numerical example in Section 3.1. Note that increasing ρ from 0.5 (Fig. 2a) to 0.999 (Fig. 2b) increases the width of the proposal density but more importantly, it shifts the location of the boundary of the region with positive proposal density (dashed blue contour) away from the region of significant target density (Fig. 2c).

We note that although the RML method for generating proposals via the minimization of a cost function is somewhat inflexible, the target density for $(\mathbf{x}, \mathbf{d})$ can be chosen much more freely as long as the marginal density for $\mathbf{x}$ is correct. In the terminology of Papaspiliopoulos et al. [28], we have chosen to use a centered parameterization, but if

we had chosen $\tilde{\mathbf{d}} = \mathbf{d} - (1 - \gamma)\mathbf{g}(\mathbf{x})$, then the pair $(\mathbf{x}, \tilde{\mathbf{d}})$ would have been conditionally independent. That parameterization is similar to the choice made by [24, 39]. It does not alter the central difficulty, however, of ensuring that the target density is approximately zero when the proposal density is zero.

3   Examples

This section illustrates the characteristics of the Metroplolized RML method on problems with multiple modes. In each case, it is possible to compare the distribution of samples with the true posterior probability density for a small inverse problem. The first example consists of a single model variable and a single data variable so that the joint augmented distribution can be visualized easily. The posterior distribution on the model variable is bimodal but the probability in the region between modes is relatively large. In the second example, the probability mass is located in more than 100 widely separated modes and the efficiency of the Metropolized RML method is compared with the preconditioned Crank Nicholson method [10]. The third example illustrates application to a problem with nongaussian prior.

3.1   Example 1: Bimodal

In this example, the posterior distribution for the model variable is bimodal. A region of small but significant probability density spans the region between the peaks. This problem was originally used as to test the sampling of RML algorithm [25] where it was shown that correct sampling could be attained if the marginal distribution of transition proposals (Eq. 5) is used in the MH test. Computation of the marginal distribution was relatively difficult, however, even for this single-variable problem because the integration had to be limited to the region of the joint model-data space in which the mapping was invertible. Here we apply the augmented variable RML method (Algorithm 2) on the joint model-data space, without need for computation of the marginal distribution of proposals. The parameter γ that determines the relative contribution of modelization error vs observation error in the posterior is set at a small nonzero value (0.01) to prevent degeneracy of the posterior distribution.

We consider the problem of sampling from the following univariate distribution

$$\pi(x) = a \exp\left[-\frac{(x - \mu)^2}{2\sigma_x^2} - \frac{(g(x) - d_{\mathrm{obs}})^2}{2\sigma_d^2}\right], \tag{20}$$

where $\mu = 1.9$, $d_{\mathrm{obs}} = 0.8$, $\sigma_x^2 = 0.1$, $\sigma_d^2 = 0.01$, $g(x) = 1 - 9(x - 2\pi/3)^2/2$, and $a \approx 4.567$. The true target distribution is shown in Fig. 4a as a solid curve. The first term in Eq. (20) is the prior density for m. The second term represents the likelihood term in the Bayesian inverse problem. Because of the nonlinearity of the observation operator $g(\cdot)$, the posterior is bimodal. An optimal value of ρ was determined by trial and error, but in fact, the MH

(a) Target joint density ($\gamma = 0.005$).

(b) RML proposal density ($\rho = 0.995$).

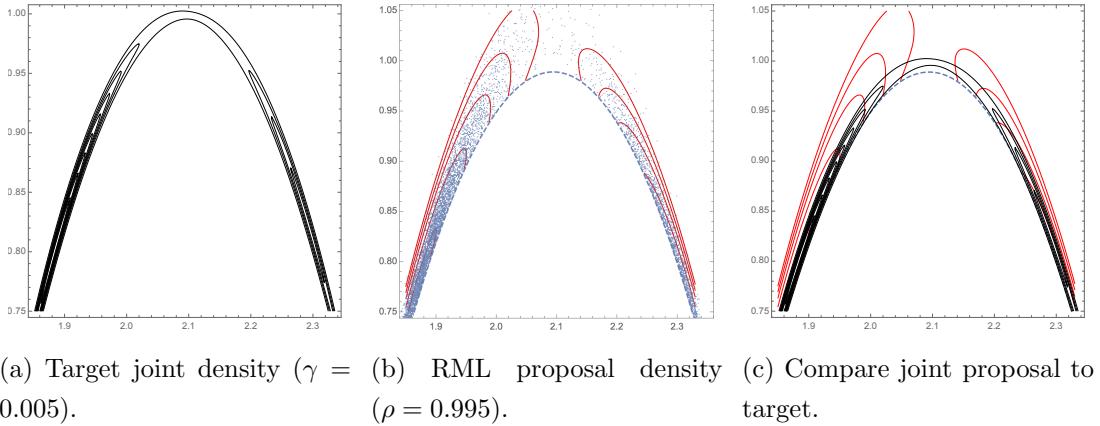(c) Compare joint proposal to target.

Figure 3: Comparison of RML proposal density with the target density for Example 1 (bimodal) sampling problem. Model variable (horizontal axis), data variable (vertical).

acceptance rate is not highly sensitive to the value of ρ. In the interval $0.5 \leq \rho \leq 0.8$ the acceptance rate for RML proposals is between 62% and 64 %. with a maximum acceptance rate of almost 0.64 when $\rho = 0.65$, which was the value we used. For the choice $\rho = 0.65$ and $\gamma = 0.01$, the joint proposal density for RML (Fig. 3a) is similar to the posterior joint density of x and d (Fig. 3b). Differences can more easily be seen in Fig. 3c where the two pdfs are superposed. The main deficiency in the proposal density is a reduced probability of proposals in the region between the two peaks.

Figure 4a shows the marginal posterior distribution of samples obtained using RML (Algorithm 2) compared with the target distribution (Eq. 20) for a chain of length 40,000. Figure 4b shows that the mixing of the chain is very good, as should be expected for an independence sampler with a high acceptance rate.
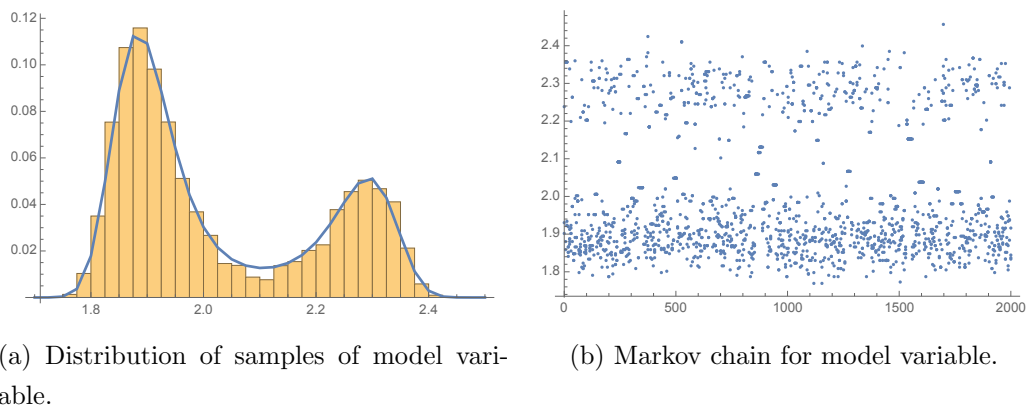

(a) Distribution of samples of model variable.

(b) Markov chain for model variable.

Figure 4: Metropolized RML sampling for Example 1 (bimodal pdf).

13

3.2 Example 2: Multimodal

In this example, the prior distribution for the two model variables is independent standard normal. The observations are related to model variables as

$$\mathbf{g}(x_1, x_2) = \begin{bmatrix} \sin(2\pi x_1) \\ \sin(2\pi x_2) \end{bmatrix}.$$

The observations are $\mathbf{d}_{\mathrm{obs}} = (0., 0.)$ with variance of observation error equal 0.04. The posterior pdf for $\mathbf{x}$ is composed of many widely separated peaks making random walk methods relatively inefficient (Fig. 5).

(a) Posterior pdf for model variables

(b) Posterior pdf for model variables conditional to $x_2 = 0$
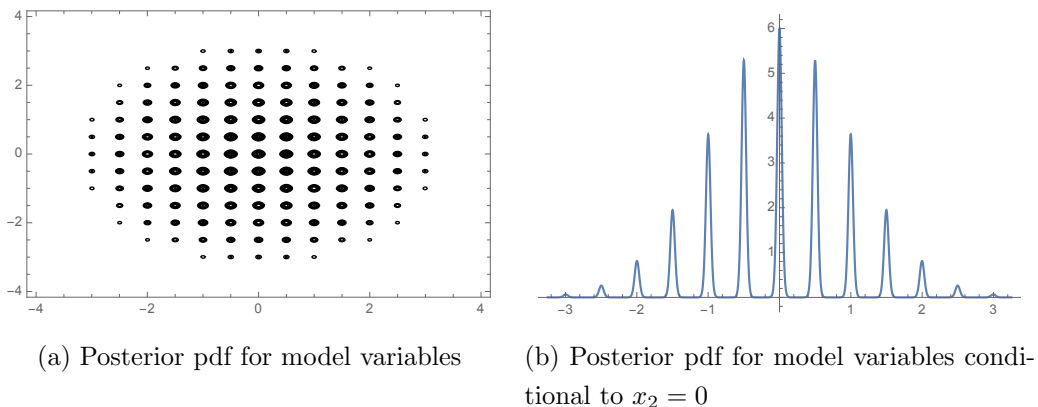
Figure 5: True posterior pdf for a problem with many modes.

The Metropolized RML method (Algorithm 2) was used for sampling with parameters $\rho = 0.995$ and $\gamma = 0.005$. Figure 6a shows the first 40000 elements of the Markov chain and Fig. 6b shows the distribution of samples from the same Markov chain in a more restricted region. Contours of the true pdf are plotted on top of the samples for comparison. For this particular example, the acceptance rate was essentially independent of ρ over a large range, varying between 87.1% for $\rho = 0.02$ to 87.4% for $\rho = 0.995$.

(a) All samples from MCMC length 40,000.

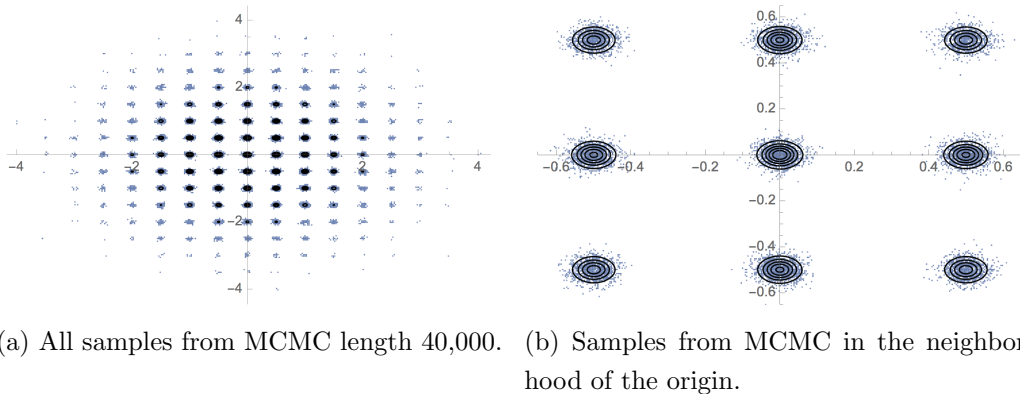(b) Samples from MCMC in the neighborhood of the origin.

Figure 6: True posterior pdf for a problem with many modes.

Because the acceptance rate for independent proposals is so high in this example, it

should be expected that the proposal density for the model variables is quite similar to the target density. Fig. 7 compares a section of the estimated pdf from the Metropolized RML to the true pdf (top) and a section of the estimated pdf from RML without an acceptance test to the true pdf (bottom). The estimated pdf was obtained using a Gaussian smoothing kernel with bandwidth 0.01. The differences between the two sampled distributions are clearly quite subtle.


Figure 7: Compare estimated pdf (black contours) to the true pdf (red contours) for Metropolized RML (top) and RML without an acceptance test (bottom). Both Markov chains were of length 40,000 but only about 4000 samples are within this window.

3.2.1 Computational cost

Although the acceptance rate for independent samples in Metropolized RML method is quite high, the cost of generating proposals is substantially higher than the cost of a standard MCMC method. Efficient implementation generally requires estimation of gradients of the cost function and a reasonable choice for the initialization of the minimization. The computational effort to generate a single proposal in RML for this numerical example is approximately 15 function evaluations using a modified Levenberg-Marquardt minimization with finite difference estimation of gradients (MINPACK subroutine LMDIF) using $\mathbf{x}_{uc}$ to initialize the minimization. In addition, computation of the Jacobian of the mapping requires an additional 5 function evaluations, hence the cost is approximately 20 function evaluations per element of the Markov chain. The acceptance rate for MH is 0.873 so the cost is approximately $20/0.873 \approx 23$ functions evaluations per independent sample from the target distribution.

For comparison, we also applied the preconditioned Crank Nicholson method (pCN) [10] to this example. The transition kernel for pCN is of the form

$$v^{(k)} = \sqrt{1-\beta^2}\, u^{(k)} + \beta\, \xi^{(k)}, \qquad \xi^{(k)} \sim N(0, C)$$

where β is a tunable parameter of the algorithm and $u^{(k)}$ is the state of the Markov chain at step k. If one uses $\beta = 0.055$ for an acceptance rate of 0.234 as recommended, a Markov chain of length 2×10^6 does not mix outside of a single mode. As an alternative, we selected an optimal value of β based on an estimate of effective sample size using the algorithm in Hoffman and Gelman [18] with 5 independent Markov chains. For this problem with multiple widely separated modes, it was determined that $\beta = 1$ was optimal. Note that in this case pCN is an independence sampler. The acceptance rate for the pCN method was 0.022, which is equivalent to 46 function evaluations per independent sample, thus Metropolized RML is approximately twice as efficient.

If the standard deviation of noise in the data is reduced from 0.2 to 0.1, the modes of the posterior pdf become narrower and more widely separated. The acceptance rate for the pCN method falls to 0.0057 which is equivalent to 177 function evaluations per independent sample. The acceptance rate for Metropolized RML is slightly higher (0.886) than when $\sigma_D = 0.2$ but the average number of function evaluations per iteration ($15 + 5$) is almost unchanged from that case. Thus, although the modes of the target pdf are more widely separated, the cost of using Metropolized RML is still approximately $20/0.886 \approx 23$ function evaluations per independent sample and the RML method becomes approximately 8.7 times as efficient as the pCN method.

For applications of RML to large-scale subsurface flow problems with many data, it has been reported that although Newton-like methods typically converge quickly, the computational cost for estimation of the Hessian at each iteration makes Newton-like methods inefficient compared to methods that only require computation of the gradient. In a comparative study, Zhang and Reynolds [40] found the limited memory BFGS minimizer to be more efficient than other methods, including Gauss-Newton, Modified Levenberg- Marquardt, preconditioned conjugate gradient, or BFGS. In all cases, the sample from the prior, $\mathbf{x}_{\mathrm{uc}}$, was used as the initial guess for minimization.

3.3 Example 3: Nongaussian prior

In this example, we look at a problem for which the prior distribution is not Gaussian, but for which a transformation to Gaussian variable can be introduced after which the previous methodology can be applied. The prior distribution for the model variable is exponential with mean and standard deviation equal to 1. A single measurement is made of ($g(x) = x + \epsilon$) with Gaussian observation error $\epsilon \sim N[0, 0.36]$. The observed value is $d_{\mathrm{obs}} = 1$. The prior and posterior distributions for the model variable are shown in Fig. 8.

Since the prior distribution in this example is not Gaussian, we define a new Gaussian variable z related to x as $z = F_z^{-1}(F_x(x))$, where F_x and F_z are the cdfs for x and z respectively. After transformation of the model variable, the joint prior distribution for model and data, the joint posterior distribution, and the joint proposal distribution are shown in Fig. 9. The joint proposal density that results from minimization with $\rho = 0.25$

Figure 8: Prior and posterior model variable distributions for Example 3.

is similar to the joint posterior density with $\gamma = 0.01$.

Figure 9: Joint prior, posterior, and proposal distributions for Gaussian model variable and data variable. Model variable on horizontal axis.

An inverse transform must be applied to each sample from the Markov chain to obtain samples in the original model space. Fig. 10a shows that the distribution of Monte Carlo samples is indistinguishable from the target posterior distribution. When this proposal density is used in the RML independence Meteropolis sampler (Algorithm 2) the acceptance rate is 74%. Because the proposals are independent and the acceptance rate is high, the mixing in the chain is very good. Figure 10b shows the first 2000 elements of the chain, with no evidence of a burn-in period or of nonstationary behavior. Recall that the candidates are first generated from the joint prior, then a minimization is performed to place the proposals in regions of high probability. Fig. 10c shows the mapping of the first 200 accepted candidates in the Markov chain from the joint prior to the proposal distribution. The weighting has already been used for this figure so that the 26% unsuccessful candidates are not shown.

4   Discussion

In this paper, we proposed an augmented variable independence Metropolis sampler that uses minimization to place proposals in regions of high probability. The computation

(a) Distribution of samples from Algorithm 2 (histogram) compared with posterior distribution (solid curve).

(b) First 2000 elements of the Markov chain for the model variable.

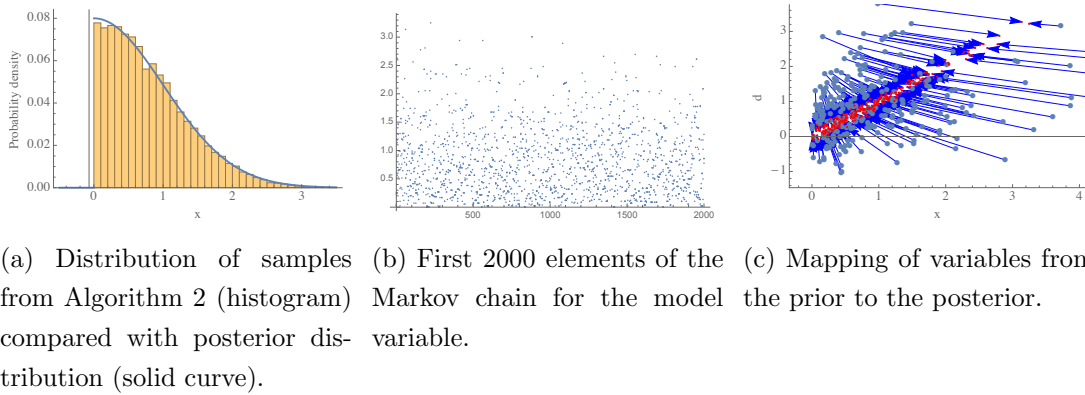(c) Mapping of variables from the prior to the posterior.

Figure 10: Summary plots for Example 3 in the original variable with an exponential prior distribution.

required for a single proposal is considerably higher than the computation required for a typical MH algorithm, but because the proposals are independent and the acceptance rate is high, the cost per independent sample from the target distribution can be low compared to alternative methods. The biggest limitation for application of the method to geoscience inverse problems will be the computation of the Jacobian determinant of the mapping that generates proposals.

The augmented variable RML method described in this paper contains two parameters ρ and γ that must be tuned for efficiency and for validity of sampling. For linear observations, the proposal density and the target density can be made close by setting both parameters to take small values. For multimodal posterior pdfs, however, the proposal density $q(\mathbf{x}_*, \mathbf{d}_*) \not> 0$ for all $\pi(\mathbf{x}_*, \mathbf{d}_*) > 0$ for arbitrary choices of ρ and γ so the Markov chain may not converge to $\pi(\cdot, \cdot)$ unless $\rho \ll 1$ and $\gamma \approx 1$. The acceptance rate for the numerical examples has relatively insensitive to the the value of γ over a fairly large range. Unlike typical applications of Metropolis random walk samplers, which often get stuck for long periods in a single mode of a multimodal distribution, the RML Metropolis sampler is more likely to remain in the regions between modes as the proposal density tends to undersample from those regions.

For cases in which the prior is not Gaussian, it can sometimes be possible to use anamorphosis to transform the variables so that the prior is Gaussian, in which cases the algorithm can be used for sampling. This is straightforward for single variable problems, but finding a transformation to multivariate Gaussian in high dimensions seems impractical.

5   Acknowledgements

as well as the Research Council of Norway (PETROMAKS).

References

[1] C. ANDRIEU AND G. O. ROBERTS, *The pseudo-marginal approach for efficient Monte Carlo computations*, Ann. Statist., 37 (2009), pp. 697–725.

[2] J. BARDSLEY, A. SOLONEN, H. HAARIO, AND M. LAINE, *Randomize-Then-Optimize: A method for sampling from posterior distributions in nonlinear inverse problems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1895–A1910.

[3] J. M. BARDSLEY, A. SOLONEN, A. PARKER, H. HAARIO, AND M. HOWARD, *An ensemble Kalman filter using the conjugate gradient sampler*, International Journal for Uncertainty Quantification, 3 (2013), pp. 357–370.

[4] J. BESAG AND P. J. GREEN, *Spatial statistics and Bayesian computation*, J. Royal Statistical Society B, 55 (1993), pp. 25–37.

[5] L. BONET-CUNHA, D. S. OLIVER, R. A. REDNAR, AND A. C. REYNOLDS, *A hybrid Markov chain Monte Carlo method for generating permeability fields conditioned to multiwell pressure data and prior information*, SPE Journal, 3 (1998), pp. 261–271.

[6] S. CALVERLEY, *Estimating uncertainty in radiographic analysis*, Discovery, The Scientific and Technology Journal of AWE, (2011), pp. 2–9.

[7] J. CARRERA, A. ALCOLEA, A. MEDINA, J. HIDALGO, AND L. J. SLOOTEN, *Inverse problem in hydrogeology*, Hydrogeology Journal, 13 (2005), pp. 206–222.

[8] Y. CHEN AND D. S. OLIVER, *History matching of the Norne full-field model with an iterative ensemble smoother*, SPE Reservoir Evaluation & Engineering, 17 (2014), pp. 244–256.

[9] M. CHRISTIE, V. DEMYANOV, AND D. ERBAS, *Uncertainty quantification for porous media flows*, J. Comput. Phys., 217 (2006), pp. 143–158.

[10] S. L. COTTER, G. O. ROBERTS, A. M. STUART, AND D. WHITE, *MCMC methods for functions: Modifying old algorithms to make them faster*, Statist. Sci., 28 (2013), pp. 424–446.

[11] P. DOSTERT, Y. EFENDIEV, T. Y. HOU, AND W. LUO, *Coarse-gradient Langevin algorithms for dynamic data integration and uncertainty quantification*, J. Comput. Phys., 217 (2006), pp. 123–142.

[12] S. DUANE, A. D. KENNEDY, B. J. PENDLETON, AND D. ROWETH, *Hybrid Monte Carlo*, Physics Letters B, 195 (1987), pp. 216–222.

[13] A. A. Emerick and A. C. Reynolds, *Investigation of the sampling performance of ensemble-based methods with a simple reservoir model*, Computat. Geosci., 17 (2013), pp. 325–350.

[14] G. Fort, *Sampling multimodal densities on large dimensional spaces*, International Conference '7th Journes de Statistique du Sud', Barcelone (Spain), June, 2014.

[15] G. Fort, B. Jourdain, E. Kuhn, T. Lelièvre, and G. Stoltz, *Efficiency of the Wang-Landau algorithm: A simple test case*, Applied Mathematics Research eXpress, 2014 (2014), pp. 275–311.

[16] G. Gao, M. Zafari, and A. C. Reynolds, *Quantifying uncertainty for the PUNQ-S3 problem in a Bayesian setting with RML and EnKF*, SPE Journal, 11 (2006), pp. 506–515.

[17] J. Geweke, *Bayesian inference in econometric models using Monte Carlo integration*, Econometrica, 57 (1989), pp. 1317–1339.

[18] M. D. Hoffman and A. Gelman, *The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo*, Journal of Machine Learning Research, 15 (2014), pp. 1593–1623.

[19] P. K. Kitanidis, *Quasi-linear geostatistical theory for inversing*, Water Resour. Res., 31 (1995), pp. 2411–2419.

[20] J. S. Liu, *Metropolized independent sampling with comparisons to rejection sampling and importance sampling*, Statistics and Computing, 6 (1996), pp. 113–119.

[21] J. S. Liu, F. Liang, and W. H. Wong, *The multiple-try method and local optimization in Metropolis sampling*, Journal of the American Statistical Association, 95 (2000), pp. 121–134.

[22] J. Martin, L. Wilcox, C. Burstedde, and O. Ghattas, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1460–A1487.

[23] D. S. Oliver, *On conditional simulation to inaccurate data*, Math. Geology, 28 (1996), pp. 811–817.

[24] ——, *Minimization for conditional simulation: Relationship to optimal transport*, Journal of Computational Physics, 265 (2014), pp. 1–15.

[25] D. S. Oliver, N. He, and A. C. Reynolds, *Conditioning permeability fields to pressure data*, in European Conference for the Mathematics of Oil Recovery, V, 1996, pp. 1–11.

[26] D. S. Oliver, A. C. Reynolds, and N. Liu, *Inverse Theory for Petroleum Reservoir Characterization and History Matching*, Cambridge University Press, Cambridge, first ed., 2008.

[27] D. S. Oliver, Y. Zhang, H. A. Phale, and Y. Chen, *Distributed parameter and state estimation in petroleum reservoirs*, Computers & Fluids, 46 (2011), pp. 70–77.

[28] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld, *A general framework for the parametrization of hierarchical models*, Statist. Sci., 22 (2007), pp. 59–73.

[29] S. Reich and C. Cotter, *Probabilistic Forecasting and Bayesian Data Assimilation*, Cambridge University Press, 2015.

[30] G. O. Roberts, A. Gelman, and W. R. Gilks, *Weak convergence and optimal scaling of random walk Metropolis algorithms*, Ann. Appl. Probab., 7 (1997), pp. 110–120.

[31] G. O. Roberts and J. S. Rosenthal, *Optimal scaling for various Metropolis-Hastings algorithms*, Statist. Sci., 16 (2001), pp. 351–367.

[32] C. Sherlock, A. H. Thiery, G. O. Roberts, and J. S. Rosenthal, *On the efficiency of pseudo-marginal random walk Metropolis algorithms*, Ann. Statist., 43 (2015), pp. 238–275.

[33] G. Storvik, *On the flexibility of Metropolis-Hastings acceptance probabilities in auxiliary variable proposal generation*, Scandinavian Journal of Statistics, 38 (2011), pp. 342–358.

[34] A. Tarantola, *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation*, Elsevier, Amsterdam, The Netherlands, 1987.

[35] A. Tarantola and B. Valette, *Generalized nonlinear inverse problems solved using the least squares criterion*, Reviews of Geophysics and Space Physics, 20 (1982), pp. 219–232.

[36] L. Tierney, *Markov chains for exploring posterior distributions (with discussion)*, Annals of Statistics, 22 (1994), pp. 1701–1762.

[37] H. Tjelmeland and J. Eidsvik, *On the use of local optimizations within Metropolis-Hastings updates*, Journal of the Royal Statistical Society Series B-Statistical Methodology, 66 (2004), pp. 411–427.

[38] H. Tjelmeland and B. K. Hegstad, *Mode jumping proposals in MCMC*, Scandinavian Journal of Statistics, 28 (2001), pp. 205–223.

[39] K. Wang, T. Bui-Thanh, and O. Ghattas, *A randomized maximum a posterior method for posterior sampling of high dimensional nonlinear Bayesian inverse problems*, arXiv preprint arXiv:1602.03658v1, (2016).

[40] F. Zhang and A. C. Reynolds, *Optimization algorithms for automatic history matching of production data*, in Proceedings of 8th European Conference on the Mathematics of Oil Recovery, 2002.