Uncertainty quantification of geological model parameters in 3D gravity inversion by Hessian informed Markov chain Monte Carlo

Zhouji Liang*, Florian Wellmann* and Omar Ghattas†

**Computational Geoscience and Reservoir Engineering(CGRE),*

RWTH Aachen University,

Aachen, Germany

†*Oden Institute for Computational Engineering and Sciences,*

Jackson School of Geosciences,

The University of Texas at Austin

Austin, Texas, USA

(December 16, 2021)

Running head: ***Hessian-informed Model-based Inversion***

ABSTRACT

Geological modeling has been widely adopted to investigate underground geometries. However, modeling processes inevitably have uncertainties due to scarcity of data, measurement errors and simplification of modeling method. Recent developments in geomodeling methods have introduced a Bayesian framework to constrain the model uncertainties by considering additional geophysical data into the modeling procedure. Markov chain Monte Carlo (MCMC) methods are normally used as tools to solve the Bayesian inference problem. To achieve a more efficient posterior exploration, advances in MCMC methods utilize derivative information. Hence, we introduce an approach to efficiently evaluate second-order

derivatives in geological modeling and introduce a Hessian-informed MCMC method, the generalized preconditioned Crank-Nicolson (gpCN), as a tool to solve the 3D model-based gravity Bayesian inversion problem. The result is compared with two other widely applied MCMC methods, random walk Metropolis-Hasting and Hamiltonian Monte Carlo, on a synthetic three-layer geological model. Our experiment demonstrates that superior performance is achieved by the gpCN, which has the potential to be generalized to more complex models.

INTRODUCTION

In many geoscience applications, inversion methods are used to estimate subsurface properties (e.g., structure, density and porosity) from observed geophysical data. Conventional geophysical inversion aims to find the best-fit parameter sets that minimize the error between observed geophysical data and simulation results. However, in practical cases, observational data have the general difficulties of sparsity and noise. The uncertainties in the data are combined with the geometry errors from the geological model constructions, leading to numerous possible solutions within the uncertainty range, and hence ill-posedness. Geoscientists are therefore interested in not only the best-fit model parameters but also a quantification of uncertainties associated with these parameters(e.g., Sen and Stoffa, 1996; Wellmann et al., 2018; Witter et al., 2019).

In this study, we focus on the uncertainty quantification in model-based inversion problems. Developments in geomodeling have provided us with tools to generate 3D structural geological models based on interface and orientation information from observed or inferred data (Caumon et al., 2009; Caumon, 2010; Wellmann and Caumon, 2018). In contrast to directly inverting the property field (e.g. density or thermal conductivity field), parametric geological models have the advantage of easy interpretability and low parameter dimensionality. Geophysical data can be used as an additional constrained to the developed model? The gravity method has been widely applied in geophysics to investigate subsurface geometries and properties (Nabighian et al., 2005). While conventional gravity inversion often suffers from ill-posedness due to insufficient data and intrinsic ambiguity (Boulanger and Chouteau, 2001), it is suitable as additional data to constrain a geological model.

Recent developments in structural geomodeling methods allow geoscientists to quantify

the uncertainties in geological models based on prior geological knowledge and additional geophysical data in a Bayesian inference approach (de la Varga and Wellmann, 2016). The Bayesian approach has long been used as a tool to quantitatively combine uncertainties from various sources in a probabilistic workflow (Tarantola and Valette, 1982; Mosegaard and Tarantola, 1995; Sambridge and Mosegaard, 2002; Sambridge et al., 2013). In a Bayesian inference problem, the parameters of interest (prior information) and the observational data (likelihood function) are treated probabilistically. The inference provides us the updated uncertainty range of the parameters of interest given the observational data. This output uncertainty range is known as the posterior, which is also a distribution. However, it is often not possible to calculate the posterior distribution analytically. The Markov chain Monte Carlo (MCMC) is used to sample from the posterior distribution by generating Markov chains of samples e.g. according to the Metroplis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970). The posterior distribution can then be approximated by the density distribution of the generated samples. The samples can also provide useful statistical information about the posterior probability density function (PDF) (e.g., mean and standard deviation).

The random walk Metropolis-Hasting (RMH) algorithm is likely the most popular MCMC method. It is the simplest algorithm and has been widely adopted from the pioneering applications in traditional geophysical inversion (Mosegaard and Tarantola, 1995; Malinverno and Leaney, 2000; Malinverno, 2002) through the recent applications of geological modeling inversion (de la Varga and Wellmann, 2016; Wellmann et al., 2018; de la Varga et al., 2019). While the classical RMH algorithm is simple in its intuition and implementation, it is often inefficient for complex posterior distributions. Complexity may result from high dimensionality or strong nonlinearities of the model. This leads to slow conver-

gence of the MCMC chain (Ruggeri et al., 2015). The inefficiency results primarily from the mismatch between the proposal distribution used in the random walk method and the actual posterior distribution. The recent work of Scalzo et al. (2019) used a preconditioned Crank–Nicolson (pCN) method to address the high dimensionality problem in the context of the joint geophysical inversion problems.

In cases where the gradient of the negative logarithm of the posterior with respect to the parameters can be obtained easily and efficiently, gradient information can accelerate an MCMC method by biasing samples toward higher probability regions. The Metropolis adjusted Langevin algorithm (MALA) (Rossky et al., 1978; Roberts et al., 1996), which involves Langevin diffusions, utilizes gradient information to confine the induced pseudo-dynamics (Betancourt, 2019). MALA can significantly increase the efficiency of MCMC. However, its performace deteriorates for strongly anisotropic distributions since the gradient may no longer point in a globally useful direction. A recent example of the application of this algorithm to the seismic inversion problem can be found in the work of Mosser et al. (2020). Another popular gradient-informed algorithm, the Hamiltonian Monte Carlo (HMC) method, also known as hybrid Monte Carlo, was first introduced by Duane et al. (1987). HMC analogizes the Hamiltonian flow of a trajectory in a potential field, which helps to draw more independent samples, at the cost of a leapfrog integrator for the Hamiltonian dynamics at each iteration (Betancourt, 2019). In addition, the parameters for the leapfrog iterations require careful tuning, and the computational cost of leapfrog steps can be substantial (Girolami and Calderhead, 2011). Recent studies have shown the applications of HMC in geophysical inversion (Muir and Tkalcic, 2015; Fichtner et al., 2019).

The use of higher-order derivative information in MCMC, in particular the Hessian of the negative log-posterior has been shown to lead to more rapid exploration of the posterior

by adapting to the posterior curvature and anisotropy (Geweke and Tanizaki, 1999; Qi and Minka, 2002). Explicitly constructing the Hessian for high dimensional inverse problems is prohibitive. Instead, for many ill-posed inverse problems, the Hessian of the negative log-likelihood admits a low rank approximation, and this can be efficiently computed using a randomized eigensolver along with higher order adjoints (Martin et al., 2012; Petra et al., 2014; Isaac et al., 2015). A number of advanced Hessian-based MCMC methods for large-scale inverse problems have appeared in the past decade, for example, the stochastic Newton MCMC method (Martin et al., 2012), dimension-independent, likelihood-informed MCMC (Cui et al., 2016) and adaptive Gaussian process emulated geometric Monte Carlo (Lan et al., 2016). These Hessian-based methods overcome the slow convergence of gradient-only MCMC for ill-posed problems and are applicable in high-dimensional parameter space.

Developing first and higher order adjoints for complex simulations can be challenging. As an alternative, automatic differentiation (AD) can be employed. The earlier development of geological modeling methods by de la Varga et al. (2019) utilized automatic differentiation (AD) to make gradient evaluations tractable. Recent work of Güdük et al. (2021) presented the successful application of this technique to model-based probabilistic inversion using HMC. In this work, we introduce the Hessian-based MCMC method as a tool to conduct efficient uncertainty quantification in model-based geophysical simulation. We follow the implicit geological modeling method and forward geophysical simulation introduced by De La Varga et al. (2015); de la Varga et al. (2019) and adapted the Hessian-based MCMC algorithm developed by Villa et al. (2018). Here, we first find the maximum a posteriori (MAP) point by using Adam, a adaptive gradient-based optimization method (Kingma and Ba, 2014). Then we construct the Laplace approximation of the posterior PDF by equating the posterior covariance to the inverse of the Hessian of the negative-log poste-

rior at the MAP point. A generalized preconditioned Crank-Nicolson proposal (Rudolf and Sprungk, 2018) is then applied to conduct the MCMC sampling. Most importantly, to efficiently calculate the second-order derivatives, we further extend the geomodeling framework GemPy (de la Varga et al., 2019) to an implementation in TensorFlow (Abadi et al., 2016), a machine-learning framework that allows fast derivative evaluation through AD. The second-order derivative is calculated by a forward-over-backward propagation to achieve time and memory efficiency. We establish a conceptual 3D fold model and simulate synthetic gravity as the observational data. We compare the efficiency of gpCN with the RMH and HMC methods for this inverse problem. The results show the potential of Hessian-based MCMC methods in real-world geophysical inversion applications.

PRELIMINARIES

In this section, we briefly introduce the forward model, including the construction of the structural geological model and the gravity forward simulations, and formulate the Bayesian inverse problem. We review two commonly used MCMC algorithms namely the RMH algorithm and HMC algorithm, which were implemented in our study as a comparison.

Implicit modeling

To generate the 3D geological models, we applied a implicit surface representation method, specifically the universal co-kriging method, originally introduced by Lajaunie et al. (1997), which has been implemented in the recently developed open-source software Gempy (de la Varga et al., 2019). A geological model can be characterized by a series of interface points and orientations. Compared with explicitly defining the geometry of the model in an explicit

method, the implicit method handles the unknown area and complex geometries automatically using an interpolation function, thereby simplifying the modeling procedure in some cases (Wellmann and Caumon, 2018). Here, we will briefly describe the universal co-kriging method.

A 3D geological model can usually be described by several geological units bounded by interfaces or faults. These interfaces and faults are typically the result of changing sedimentary environments or tectonic events during geological history. The general idea of the modeling method is to interpolate a scalar field $Z(\mathbf{x})$ over the entire model in the continuous 3D space $(x, y, z) \in R^3$ based on the spatial relationship of the interfaces and orientations, where $\mathbf{x} = \{x, y, z\}$ in Euclidean coordinates. Two parameters of the interfaces are considered to characterize the scalar field namely the position of the layer interface points and their orientations. We assume that points set J_K at the same interface K has the same scalar value $Z(x_j) \forall j \in J_K$, the gradients of the scalar field ∇Z^K are perpendicular to the isocurve of the scalar field and parallel to the orientations of interface point G. These constraints can be formulated as follows (Lajaunie et al., 1997):

$$\begin{cases} \nabla Z^K(\mathbf{x}_i) = G_i & \forall i \in I \\ \left\langle \nabla Z^K(\mathbf{x}_{i'}), \tau_{i'} \right\rangle = 0 & \forall i' \in I' \\ Z^K(\mathbf{x}_j) - Z^K(\mathbf{x}_j) = 0 & \forall j, j' \in J_k \end{cases}$$

Solving the co-kriging system based on the interface variance and covariance information and the above constraints provides a continuous scalar field. A 2D scalar field example is given in Figure 1. Explicit representation of geological surfaces can be extracted as isosurfaces in the scalar field. The full implementation is described in de la Varga et al. (2019).

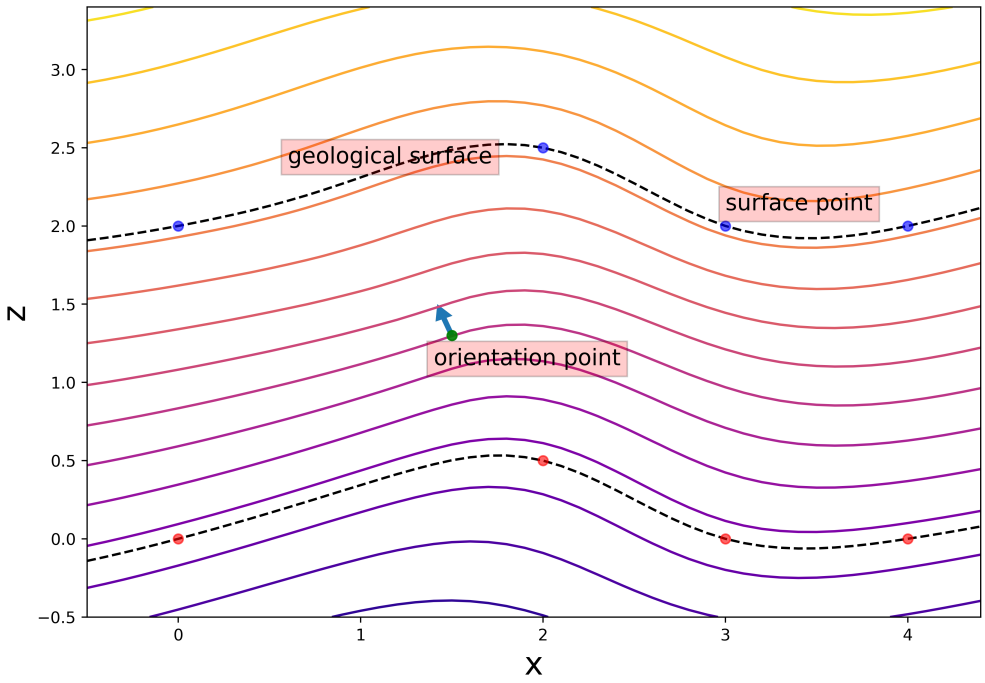

Figure 1: Interpolated scalar filed using co-kriging of interface points and orientation values. Isocurves represents the scalar values based on the interpolation of the surface points and orientations.

Gravity simulation

Gravity method, to date still one of the most widely used geophysical techniques in exploration (Nabighian et al., 2005) can be used to integrate geophysical information into geomdeling procedures. One of the most popular methods to include geophysical data is through inversion (Tarantola, 2005). By including additional geophysical information, we can validate the existing model or reduce its uncertainty range. A forward simulation of the gravity field is required to conduct the inversion. To calculate the numerical solution of the gravity field, we first discretize the model into cells and then assign densities to each of these cells. Subsequently, we adapt the formulation presented by Nagy (1966) and implementation by de la Varga et al. (2019) to simulate the gravity attraction from these rectangular prisms.

$$F_z = \rho \cdot t_z$$

$$t_z = |||x \ln(y+r) + y \ln(x+r) - z \arctan\left(\frac{xy}{zr}\right)|_{x_1}^{x_2}|_{y_1}^{y_2}|_{z_1}^{z_2}$$

where x,y, and z are the Cartesian components of the targeting prism, ρ is density and r stands for the Euclidian distance from the center of the cell to the receiver points. In this work we use virtual receivers on the ground surface, which is usually the top of the model. By summing up the attractions of all the cells, the gravity at that receiver point can be simulated.

The resolution of the model, which reflects the size of the cells, has an impact on the simulation. Low-resolution cells introduce more uncertainties into the model. Additionally, receivers close to the boundary of the model will have boundary effects. Therefore, it is

necessary to further extend the model to minimize the boundary effects. A model with a small cell size as well as a large model extent would certainly be ideal; however, our computational resources are limited. The attraction force decays exponentially with respect to distance r from the gravitation equation:

$$\Delta F = G\rho\frac{\Delta v}{r^2}$$

Hence, instead of using a regular grid (which means all cells have the same size), we apply a irregular grid (termed as centered grid in this paper) (see also (Güdük et al., 2021,p.5)). To simplify the calculation, we make this grid isotropic around a center position where the gravity response is calculated (illustrated in Figure 2). Both the height and widths of the prism cells grow exponentially with distance from the center. The scalar field is queried in these extended regions. For each receiver at the surface, a corresponding centered grid is created.

Bayesian Inference and Markov-chain Monte Carlo

The geophysical observations $\boldsymbol{d}_{obs} \in D$ in any forward problem can be described by $\boldsymbol{d}_{obs} = \boldsymbol{f}(\boldsymbol{m}) + \boldsymbol{e}_d$, where $\boldsymbol{m} \in M$ is the model parameters, $\boldsymbol{f}$ is the parameter-to-observation map that represents the modeling or simulation procedure, and $\boldsymbol{e}_d$ is due to both the error in modeling and the noise in the data. Geophysical inversion problems are often ill-posed (Tikhonov and Arsenin, 1977) because of sparse and erroneous data, imperfect forward modeling and loss of information in the parameter-to-observation map. Therefore, in many cases we are interested not only in the best-fit parameters set but also other plausible solutions that fit the observations within an uncertainty range. Solving such inversion problems then amounts to exploring the distribution of plausible model parameters. Bayesian infer-

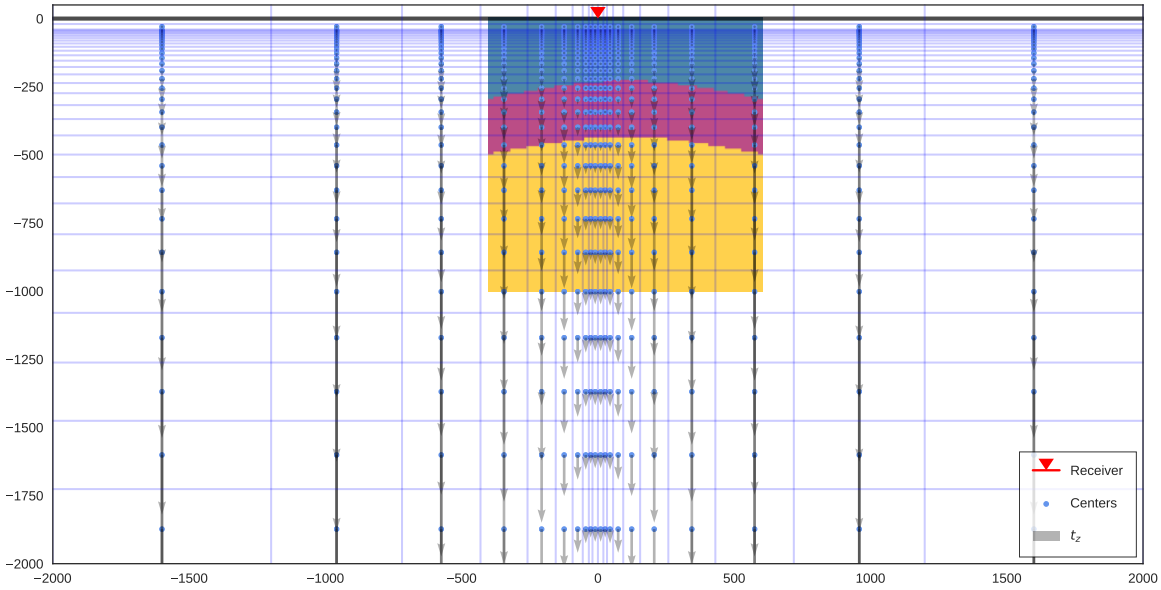

Figure 2: Illustration of centered grid with respect to one receiver. The grids are centered around the receiver, and extended further to alleviate the boundary effect.

ence, provides a systematic framework to quantify uncertainties of the inferred parameters, given uncertainties in the forward model, the observational data, and any prior information on the model parameters.

Bayesian inference starts from Bayes' theorem,

$$p\left(\boldsymbol{m} \mid \boldsymbol{d}_{obs}\right) = \frac{p\left(\boldsymbol{d}_{obs} \mid \boldsymbol{m}\right) p(\boldsymbol{m})}{p\left(\mathbf{d}_{obs}\right)}$$

where $p\left(\boldsymbol{d}_{obs} \mid \boldsymbol{m}\right)$ is known as the likelihood of the data $\boldsymbol{d}_{obs}$ given model parameters $\boldsymbol{m}$, and $p\left(\boldsymbol{m}\right)$ represents the prior probability of the model parameters $\boldsymbol{m}$. The posterior is normalized by the denominator $p\left(\boldsymbol{d}_{obs}\right)$, which is called the marginal likelihood or evidence, so that the integral of the right-hand side is one. However, the denominator, marginal likelihood, is well known to be intractable in most cases because of the infinite possibility of

observations. Sampling methods, such as MCMC, are therefore often used to approximate the integrals by sampling the posterior space. In the random walk Metropolis-Hasting algorithm, which is the most well-known and intuitively simple MCMC algorithm, a chain of samples is drawn based on a Gaussian proposal distribution, and by an accept/reject step, the sampling chain can move from the previous state to a new state based on an acceptance probability a. The accept/reject step restores the balance of the reversibility condition; therefore, the target distribution can be represented by the chain of samples. The steps of the Metropolis-Hasting algorithm can be expressed as follows:

Algorithm 1: METROPOLIS-HASTING ALGORITHM

while $k <$ *number of samples* **do**
  Set $k = 0$ and initiate with $\boldsymbol{m}^{(0)}$
  Propose $\boldsymbol{m}^{cand} = \boldsymbol{m}^k + \beta\boldsymbol{\xi}^k$, $\boldsymbol{\xi} \sim \mathcal{N}(0, C)$
  Set $\boldsymbol{m}^{k+1} = \boldsymbol{m}^{cand}$ with acceptance probability a
  Set $\boldsymbol{m}^{k+1} = \boldsymbol{m}^k$ otherwise
end

where acceptance probability $a(\boldsymbol{m}^{cand}, \boldsymbol{m}^k) = \min\left(1, \frac{p(\boldsymbol{m}^k|\boldsymbol{d}_{obs})}{p(\boldsymbol{m}^{cand}|\boldsymbol{d}_{obs})}\right)$, β is the step size, and C is the covariance matrix of the proposal distribution.

The random-walk method has limitations in cases of high dimensionality and high correlations. Gradient information is therefore employed to assist in the posterior exploration. The most popular gradient-based method, the HMC algorithm, is designed to draw independent samples, and therefore efficiently explore the state space (Duane et al., 1987; Neal, 1993; Chen et al., 2014; Betancourt, 2017). HMC introduces auxiliary momentum variables

$\boldsymbol{r}$, and therefore we can write the Hamiltonian of a particle as:

$$H(\boldsymbol{m}, \boldsymbol{r}) = K(\boldsymbol{r}) + V(\boldsymbol{m})$$

where K and V are kinetic and potential energies respectively and are defined as:

$$K(\boldsymbol{r}) = \frac{1}{2} \sum_{i,j=1}^{n} r_i \boldsymbol{M}_{ij}^{-1} r_j$$

$$V(\boldsymbol{m}) = -\log p(\boldsymbol{m} \mid \boldsymbol{d})$$

Here, the mass matrix $\boldsymbol{M}$ is often set as the identity matrix (Chen et al., 2014)

Hamiltonian dynamics can be simulated by following Hamilton's equations with an artificially introduced time variable t:

$$\begin{aligned} \frac{\mathrm{d}\boldsymbol{m}}{\mathrm{d}t} &= \frac{\partial H}{\partial \boldsymbol{r}} \\ \frac{\mathrm{d}\boldsymbol{r}}{\mathrm{d}t} &= -\frac{\partial H}{\partial \boldsymbol{m}} \end{aligned} \tag{1}$$

The HMC algotithm requires that Hamilton's equations are integrated using a symplectic integrator. In practice, the "leapfrog" integrator is often used. The algorithm can be

summarized as follows (Chen et al., 2014):

Algorithm 2: Hamiltonian Monte Carlo algorithm

while $k <$ *number of samples* **do**
  Set $k = 0$ and initiate with $\boldsymbol{m}^0$ and step size ϵ
  Sample momentum $\boldsymbol{r}^k \sim \mathcal{N}(0, \boldsymbol{M})$
  Simulate discretization of Hamiltonian dynamics in Eq. 1:
  $\boldsymbol{r}_0 = \boldsymbol{r}_0 - \frac{\epsilon}{2}\nabla V(\boldsymbol{m}_0)$
  for *i=1,i< leapfrog steps n* **do**
    $\boldsymbol{m}_i = \boldsymbol{m}_{(i-1)} + \epsilon \boldsymbol{M}^{-1}\boldsymbol{r}_{i-1}$
    $\boldsymbol{r}_i = \boldsymbol{r}_{i-1} - \epsilon \nabla V(\boldsymbol{m}_i)$
  end
  Set $\boldsymbol{r}_n = \boldsymbol{r}_n - \frac{\epsilon}{2}\nabla V(\boldsymbol{m}_n)$
  Propose $(\boldsymbol{m}^{cand}, \boldsymbol{r}^{cand}) = (\boldsymbol{m}_n, \boldsymbol{r}_n)$
  Set $\boldsymbol{m}^{k+1} = \boldsymbol{x}^{cand}$ with acceptance probability a
  Set $\boldsymbol{m}^{k+1} = \boldsymbol{m}^k$ otherwise
end

where the acceptance probability $a(\boldsymbol{m}^{cand}, \boldsymbol{m}^k) = \min\left(1, e^{-H(\boldsymbol{m}^{cand}|r^{cand}) + H\left(\boldsymbol{m}^k|r^k\right)}\right)$

RMH is simple to implement and gradient-free, while HMC uses gradient information to obtain low autocorrelated samples; however, both methods are popular and widely applied. In this study, we will compare the efficiency of the generalized preconditioned Crank-Nicolson (gpCN) with both these methods and discuss their applicability.

METHODS

In this work, we conducted an end-to-end procedure from the model construction to the geophysical simulation and finally to Bayesian inference. We first generated a geological model by using the implicit modeling method. A synthetic gravity field was then simulated based on the model. Finally, we apply three different MCMC algorithms, namely RMH, HMC, and gpCN to solve the Bayesian inference on the same synthetic model. In this section, we describe the Automatic differentiation method which links the geological modeling and Bayesian inference framework and makes the high-order derivative evaluation tractable, and present, how the high-order derivative information can be used to accelerate MCMC.

gpCN

Although it is intuitively simple, the MH algorithm suffers from the problem of dimensionality (Cotter et al., 2013; Betancourt, 2017). As the dimension increases, the naive symmetric proposal used in the MH algorithm will result in most of the samples being rejected, and thereby resulting in an inefficient exploration of the posterior. A decent volumetric explanation was established by (Betancourt, 2017). Thus, researchers have been searching for an alternative to the MH algorithm. Cotter et al. (2013) introduced the preconditioned Crank–Nicolson (pCN) proposal, which is a slight modification to the original RMH algorithm but provides its scalability to high-dimensional problems.

In recent years, the idea of employing the geometry of the posterior to accelerate the exploration has attracted many researchers (Girolami and Calderhead, 2011; Martin et al., 2012; Law, 2014; Cui et al., 2016), where the first-order or second-order derivative information or either both (Martin et al., 2012) are utilized. Among those studies, Rudolf and

Sprungk (2018) introduced the gpCN as an extension to the pCN algorithm. The key idea of this algorithm is to employ the covariance of the posterior C_v to construct a proposal distribution that adapts the posterior geometry Figure 3. The algorithm can be summarized as follows:

Algorithm 3: GENERALIZED PRECONDITIONED CRANK–NICOLSON ALGORITHM

ALGORITHM

1 Set $k = 0$ and initiate with $\boldsymbol{m}^{(0)}$

2 Propose $\boldsymbol{m}^{cand} = \boldsymbol{m}_\nu + \sqrt{(1-\beta^2)}(\boldsymbol{m}^k - \boldsymbol{m}_\nu) + \beta\boldsymbol{\xi}^k$, $\boldsymbol{\xi} \sim \mathcal{N}(0, C_v)$

3 Set $\boldsymbol{m}^{k+1} = \boldsymbol{m}^{cand}$ with acceptance probability a

4 Set $\boldsymbol{m}^{k+1} = \boldsymbol{m}^k$ otherwise

where the acceptance probability $a(\boldsymbol{m}^{cand}, \boldsymbol{m}^k) = \min\left\{1, \exp\left(\Delta\left(\boldsymbol{m}^k\right) - \Delta\left(\boldsymbol{m}^{cand}\right)\right)\right\}$,

and $\Delta(\boldsymbol{m}) = \Phi\left(\boldsymbol{m}, \boldsymbol{d}_{obs}\right) + \frac{1}{2}\left\|\boldsymbol{m} - \boldsymbol{m}_{prior}\right\|^2_{C^{-1}_{prior}} - \frac{1}{2}\left\|\boldsymbol{m} - \boldsymbol{m}_\nu\right\|^2_{C^{-1}_\nu}$,

$\Phi(\boldsymbol{m}, \boldsymbol{d}_{obs})$ is the negative log likelihood function:

$$\Phi(\boldsymbol{m}, \boldsymbol{d}_{obs}) = \frac{1}{2}\left\|\boldsymbol{f}(\boldsymbol{m}) - \boldsymbol{d}_{obs}\right\|^2_{\Gamma^{-1}_{noise}}$$

Subsequently, Villa et al. (2018) introduced the method to evaluate the posterior covariance by using the Laplacian approximation with the Hessian of the log likelihood $\mathcal{H}_{misfit}$ at the MAP point $\boldsymbol{m}_\nu$ and covariance of the prior C_{prior}

$$C_\nu = \left(\mathcal{H}_{misfit}\left(\boldsymbol{m}_\nu\right) + C^{-1}_{prior}\right)^{-1}$$

$$\boldsymbol{m}_\nu := \arg\min \mathcal{J}(\boldsymbol{m}) := \left(\frac{1}{2}\left\|\boldsymbol{f}(\boldsymbol{m}) - \mathbf{d}_{obs}\right\|^2_{\Gamma^{-1}_{noise}} + \frac{1}{2}\left\|\boldsymbol{m} - \boldsymbol{m}^{prior}\right\|^2_{C^{-1}_{prior}}\right)$$

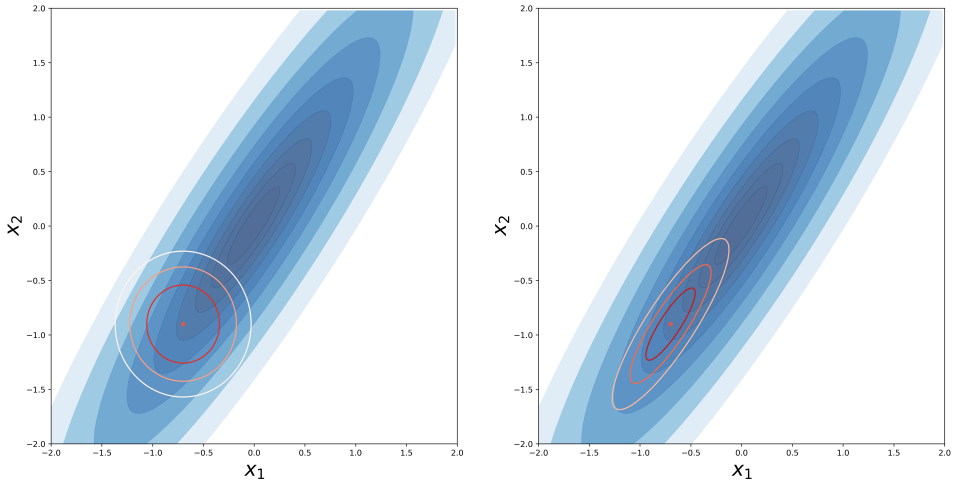

Figure 3: Different proposal distribution illustrated on the 2D Gaussian posterior distribution $\boldsymbol{d}|\boldsymbol{m} \sim \mathcal{N}(\mu, I), \mu := x_1 - 0.7x_2$, where $x_i \overset{iid}{\sim} \mathcal{N}(0, I)$. Left shows the Gaussian proposal without preconditioning. Right shows the proposal distribution with Laplacian approximation at MAP.

Here, we applied the recently developed first-order optimization algorithm Nadam (Nesterov-accelerated adaptive moment estimation)(Dozat, 2016) to find the MAP. The update rule of Nadam is given as follows:

$$\boldsymbol{m}_{t+1} = \boldsymbol{m}_t - \frac{\alpha}{\sqrt{\hat{v}_t} + \epsilon}\hat{b}_t$$

where α is the step size. $\hat{b}_t$ and $\hat{v}_t$ are the bias-corrected first and second momentum estimates of the gradients, respectively. The momentum term averages the past gradient and thus accelerates the convergence rate compared with the standard gradient descent algorithm, while the second momentum which averages the past squared gradients adapts its learning rate.

Automatic Differentiation and Evaluation of Hessian

Implementing derivative-informed MCMC methods such as HMC and gpCN in geomodeling require the derivative information in several steps such as finding the MAP point and constructing the posterior covariance approximation. This derivative refers to the derivative of the posterior (or negative log posterior) with respect to the model input parameters (e.g., surface points and orientation points). By tracing the derivative from the Bayesian inference through the geophysical simulation, geological modeling of the input parameters is nontrivial. An analytical solution of the derivatives is difficult or even impossible to find. Using methods such as finite difference (FD) is not only computationally costly but also can suffer from numerical inaccuracy. Therefore, we adopted the automatic differentiation (AD) technique, which is widely applied in the field of artificial intelligence and is critical to the success of training neural networks. Here we briefly introduce how AD works and

how gradient and higher-order derivatives can be evaluated efficiently using AD.

Consider an arbitrary function $\mathbf{y} = F(\mathbf{x})$, where F is the mapping function maps the input parameter $\mathbf{x} = x_1, x_2, ...x_i, (\mathbf{x} \in R^{D_0})$ to the cost function $\mathbf{y} = y_1, y_2, ..., y_j, (\mathbf{y} \in R^{D_1})$. The first-order partial derivative, also known as the Jacobian matrix, is given as follows:

$$(\mathcal{J}_F)^i_j(x) = \frac{\partial (F)^i}{\partial x^j}(x)$$

by constructing the composite function $F = F_N \circ F_{N-1} \circ \cdots \circ F_2 \circ F_1$ and ignoring i, the Jacobian matrix can be represented by iteratively applying the chain rule:

$$\mathcal{J}_F(x) = \frac{\partial (F)}{\partial x^j}(x) = \frac{\partial (F_N)}{\partial (F_{N-1})} \cdots \frac{\partial (F_2)}{\partial (F_1)} \frac{\partial (F_1)}{\partial x}(x)$$

AD records the derivative of each fundamental operator in a program and builds a computational graph to record the dependency of each operator. The derivative of the output with respect to the input parameter of interest is evaluated by a forward or a backward propagation using the chain rule based on the computational graph. Forward mode and backward mode both have their advantages and disadvantages, by combining these two methods, we get an efficient method for second order derivative evaluation. The approach is summarized in the following.

To evaluate the gradient in the forward mode AD, a tangent vector v is selected at the evaluation point x,

$$\mathcal{J}_F(x) \cdot v = \left(\frac{\partial (F_N)}{\partial (F_{N-1})} \cdots \left(\frac{\partial (F_2)}{\partial (F_1)} \left(\frac{\partial (F_1)}{\partial x}(x) \cdot v \right) \right) \right)$$

For example, $\frac{\partial(F)^i}{\partial(x_1)}$ can be evaluated by a tangent vector $v = (1, 0, \ldots, 0), v \in R^i$. Thus, the forward mode AD provides the directional derivative.

In contrast, backward mode AD evaluation is based on a fixed dependent variable, and the quantity of interest is the adjoint

$$\mathcal{J}_F^i(x) = \left(\left(\left(\frac{\partial(F_N)}{\partial(F_{N-1})}\right)\cdots\frac{\partial(F_2)}{\partial(F_1)}\right)\frac{\partial(F_1)}{\partial x}(x)\right)$$

Evidently, for problem $\mathbb{R}^{D_I} \to \mathbb{R}^{D_O}$, where D_I denotes the dimension of input and D_o denotes the dimension of output, if $D_I \gg D_O$, backward propagation is more efficient, which matches the cases of many machine learning problems and the Bayesian inference in our study, where the output is a single cost function, so $D_O = 1$. In contrast, forward propagation is more efficient when $D_I \ll D_O$, requiring less memory. An efficient second-order derivate calculation for multivariate input and single output problem is given by the combination of a forward-over-backward propagation. Applying backward propagation gives us the Jacobian matrix $\mathcal{J}_F(x) = [\frac{\partial F}{x_1}, \frac{\partial F}{x_2}, \ldots, \frac{\partial F}{x_j}]$. The second derivative then evaluates the partial derivative over the Jacobian, which is an $\mathbb{R}^{D_I} \to \mathbb{R}^{D_I}$ problem. Considering the memory efficiency, we can evaluate the Hessian vector product efficiently by an additional forward sweep over the backward propagation. Each forward-over-backward iteration will return a column of the Hessian matrix; therefore, this method has an $\mathcal{O}(n)$ complexity, where n is the dimension of the parameter of interest. It is significantly more efficient than evaluating the Hessian by the FD method, which has a complexity of $\mathcal{O}(n^2)$. A comprehensive review of AD and its implementation can be found in Betancourt (2018) and Margossian (2019).

Programming the geological model by adopting AD, a second-order derivative can be efficiently evaluated. The numerical comparison of the computational efficiency is discussed

in the following sections.

RESULTS

When the available data are limited, uncertainties are inevitable in 3D geological models. We can constrain the uncertainties by additional geophysical data. To configure such a problem in a Bayesian inference framework as described above, both the prior information and the likelihood must be expressed in terms of a probability distribution. As an example we constructed a 3D geological model using the implicit method described above. The numerical implementation is based on the previously developed modeling methods of the GemPy software (`https://www.gempy.org/`) and specific extension in the differentiable programming framework TensorFlow (`https://www.tensorflow.org/`).

The example model has a simple dome geometry, replicating a conceptual setting which is often observed in real-world geology. It consists of three layers, representing three distinct lithologies. The model has an extent from 0 m to 1,000 m for all three axes. The positions (x-y coordinates) of the interface points are fixed, and the depth z values of the interface points are the uncertain parameters in this study. Both the layout of the interface points and pseudo-gravity receivers is shown in Figure 4. The ground truth z value of the interface point follows a sinusoidal function. The ground truth 3D model and interface point layout are shown in Figure 5. To simplify the problem, we assume a constant distance between the two surfaces at configuration interfaces positions and only vary the upper interface points. In other words, the lower surface moves parallel in dependence of the upper surface; therefore, a constant thickness of the middle layer is expected. Interpolation could introduce some variation to the thickness at locations where no interface points are given. Sixteen pseudo-gravity receivers are evenly located at the ground surface, which is the top surface

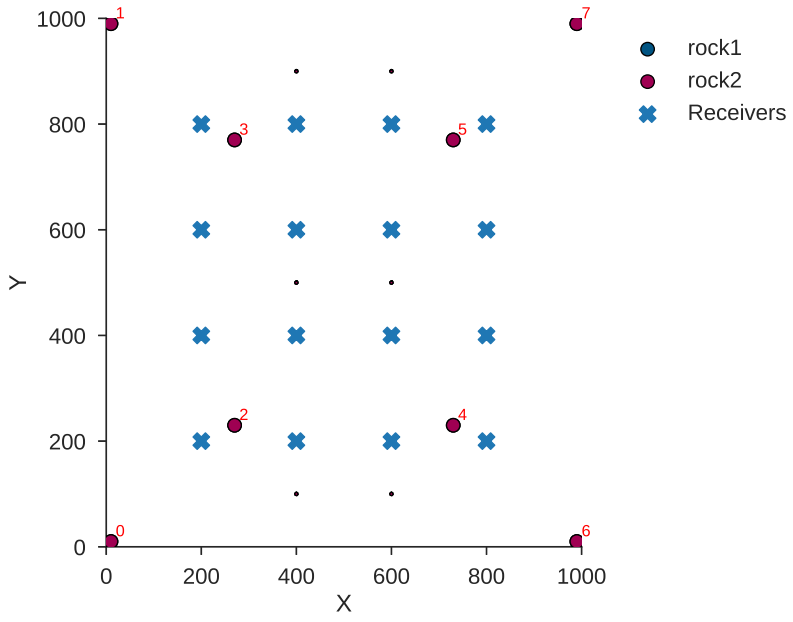

Figure 4: Top view of model parameter layout and receivers position. The interfaces point for lower rock is overlap by the top layer. Interface points 2 and 3 are set close to each other to simulate correlations.

of the model. The grid size for gravity calculation was selected to be $10 \times 10 \times 30$ to have a better resolution in the vertical axis because the parameters of interest vary along the z axis. The simulated ground truth gravity data is shown in Figure 6. Our goal is to estimate the uncertainties of the depth position of the interface points of the upper surface considering the additional gravity information.

We assume a simple Gaussian prior to all eight interface points with the mean at the ground truth location m_0 and a standard deviation (std) of 25 m. Therefore, the prior distribution can be expressed as a multivariate normal distribution centered at m_0 and a diagonal covariance matrix assuming no correlations between the parameters in our prior

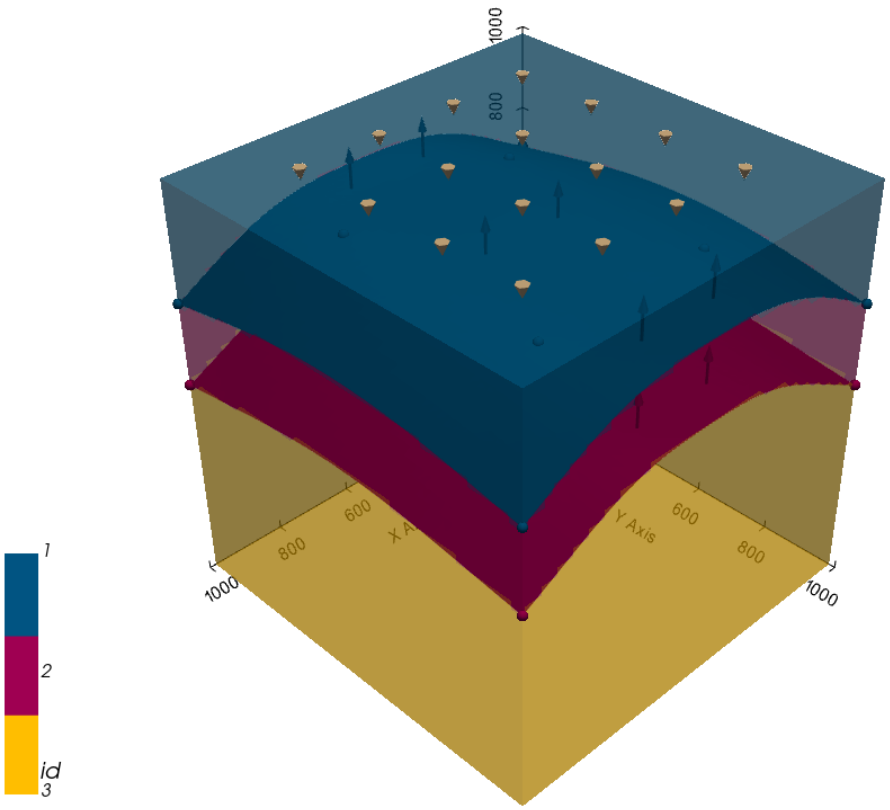

Figure 5: 3D plot of the base-case geological model. Gravity receivers placed at the top.

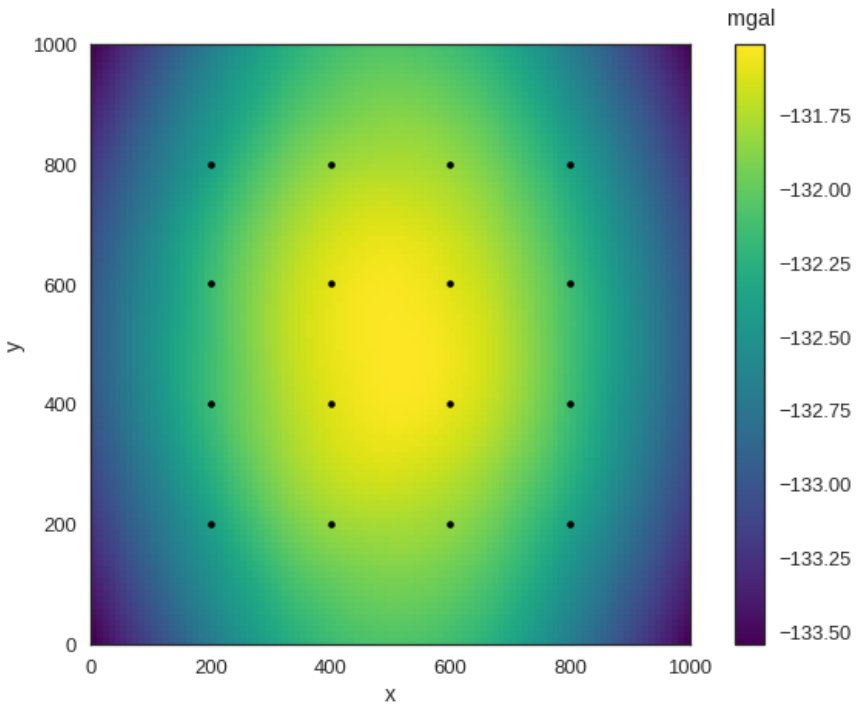

Figure 6: Simulated base-case at top surface.

information. Pseudo-noisy gravity observation data are described by another multivariate normal distribution as the likelihood function. Pseudo-noisy gravity observation data was generated by seven forward gravity simulations with an independent variation of 7 m.

To solve the Bayesian inference by MCMC, specifically by using the gpCN, we first search for the Maximum a Posteriori (MAP) point m_{MAP} in the posterior space. m_{MAP} can be sufficiently found by Nadam with the assistance of the gradient of the target negative log posterior with respect to the variable (depth value). The gradient is evaluated through AD. The time cost for each gradient evaluation is within 1 s depending on the model size. To avoid the optimization to get stuck at local minima in the parameter space, an initial status is generated by randomly sampling the prior distribution. The optimization can converge normally within 200 iterations, depending on the model complexity, step size, and the initial status. The full Hessian matrix of the target negative log posterior with respect to the variable is evaluated at the MAP point through the forward-over-backward propagation described above.

To compare the performance of gpCN, we implemented the two other most commonly used methods, namely RMH and HMC, to solve the same inference problem. The simulations were run on a single Tesla P100 GPU. The resulting sampling chain, posterior, and efficiency are compared below. Figure 7 shows the trace plot of the same two interface points for each algorithm. We observe that better-mixed chains are obtained by both HMC and gpCN, but the chain of RMH is poorly mixed, and far from convergent. To quantitatively evaluate the chain, we plotted the autocorrelation of the results from all three methods in Figure 8. A similar low autocorrelation is obtained by HMC and gpCN, and RMH samples are highly autocorrelated with the lowest acceptance rate. One can always tune the step size of MCMC to achieve fewer autocorrelated samples, but this will also lead to a lower

acceptance rate.

Although equivalent or even better sampling chain is obtained by HMC, gpCN is more efficient in terms of computational cost. In HMC, the leapfrog integrator is usually chosen to solve the numerical integration in the Hamiltonian equation, which is a computationally demanding process. The choice of step size and the number of steps used for the leapfrog integration are essential for the performance and efficiency of HMC. Too few steps will result in a bad integration, and therefore, a lower acceptance rate, while increased number of steps will increase the computational cost. Because HMC is intended to find an uncorrelated point by using several evaluations, the theoretical acceptance rate of HMC is 100%, but slightly lower due to the imperfection of numerical integration. Thus, in addition to the acceptance rate and the autocorrelation, we also compared the effective sample size n_{eff} (Liu, 2008,p. 125) per unit time. The effective sample size n_{eff} is defined as follows:

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

where ρ_k is the autocorrelation coefficient at lag k. A higher number of n_{eff} means a better mixing Markov chain. Our numerical experiments indicate that gpCN outperformed HMC by two orders of magnitude of the $AvgN_{eff}$ (Table 1), while preserving the computational efficiency, even though we only used three leapfrog steps in the experiments for HMC. Although we demonstrate only a simple example, we can expect a better performance of the gpCN in a higher dimension and more correlated case.

The main reason for the superior performance of gpCN is the highly correlated posterior space. If we use the sample of HMC as a benchmark and plot the adopted proposal distribution used in the gpCN (Figure 9), we can see that the proposal distribution captured the

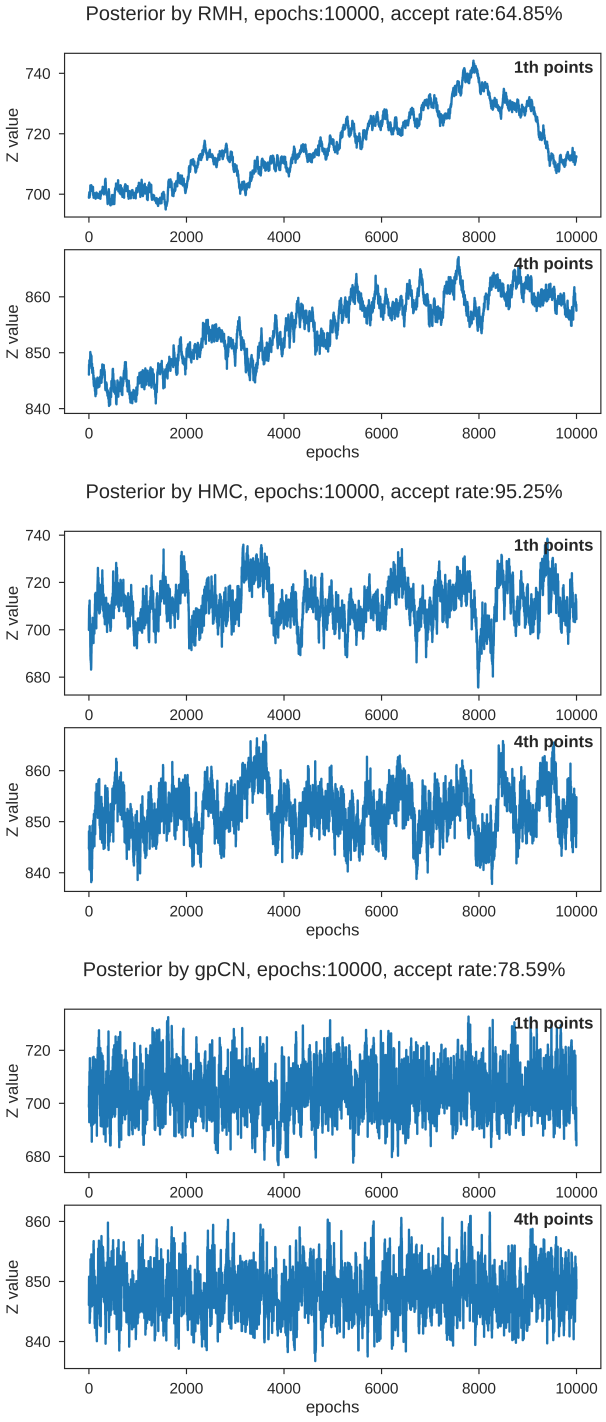

Figure 7: Trace plot of RMH, HMC and gpCN. RMH has a poor chain mixing and lowest acceptance ratio. Without considering computational efficiency, HMC has the best chain mixing and acceptance ratio.

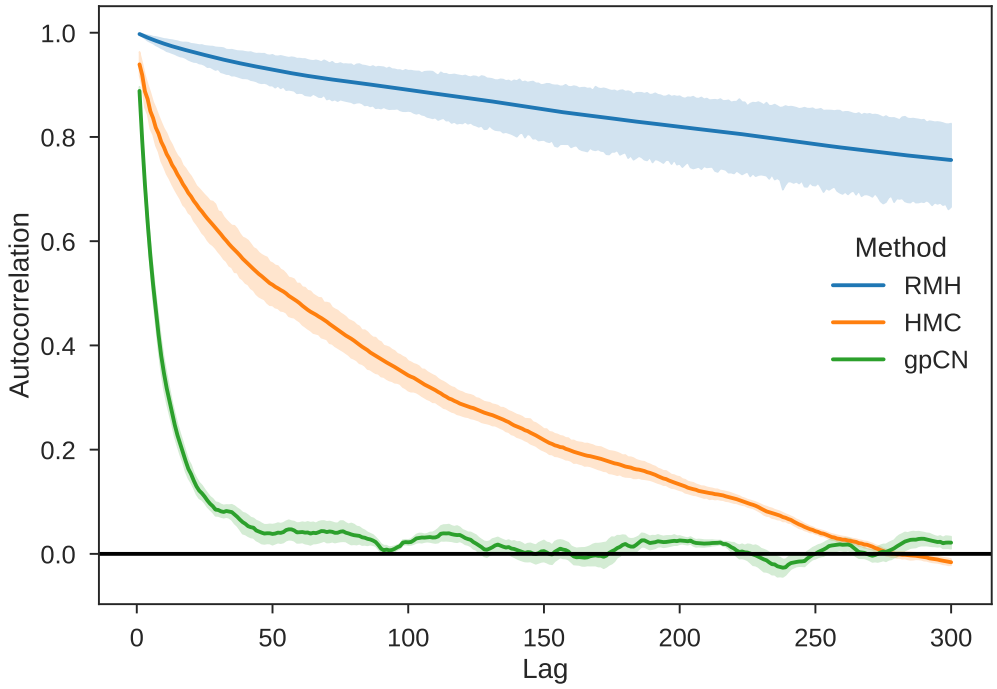

Figure 8: Autocorrelation plot of the MCMC chain of all three methods. The solid lines represent the mean autocorrelation among different variables, and the shadow represents the 95% confidence interval of the autocorrelation on different dimensions.

Methods	Compuation time for 10000 samples	n_{eff}	$AvgN_{eff}$per 100s
RMH	191s	4.0	2.1
HMC	1851s	60.8	3.3
gpCN	123s	385.5	312.5

Table 1: Computational efficiency of different methods

posterior geometry.

Finally, we represent the uncertainties following the information entropy method introduced by Wellmann and Regenauer-Lieb (2012). In Figure 10, we compared the uncertainties of the Prior and Posterior. We generated 1000 samples from the prior distribution and build the geological models based on the sampled parameters. The samples drawn by the gpCN MCMC are used to generate geological models for posterior representation. By converting the lithology changes into information entropy, we can see that the uncertainty range of the position of the upper interface is significantly reduced by additional gravity data.

To consolidate our conclusion, an additional example with less ideal prior information is conducted and compared. In this case, the prior is chosen to be a uniform distribution with the same variation but a uniform mean at $z = 780\ m$. Analogue to the previous example, pseudo noisy data is generated by tuning the base-case model, and the inference results are compared with three different MCMC methods. Even with relatively poorer prior information, the posterior distribution of the surface can be sufficiently explored by

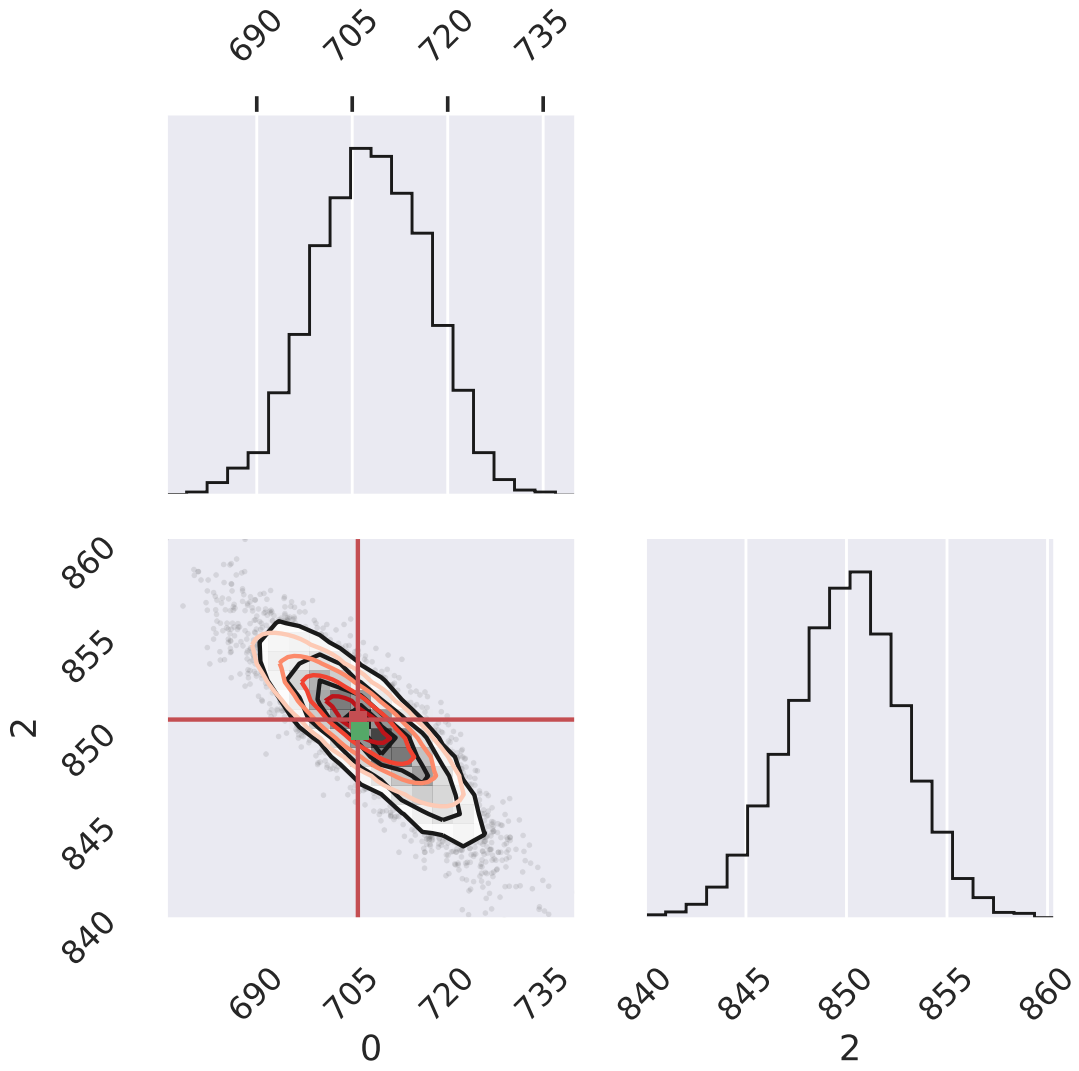

Figure 9: Posterior sampling of posint 0 and posint 2 by HMC and the proposal distribution used in gpCN in red contour. The red square represent the found MAP.

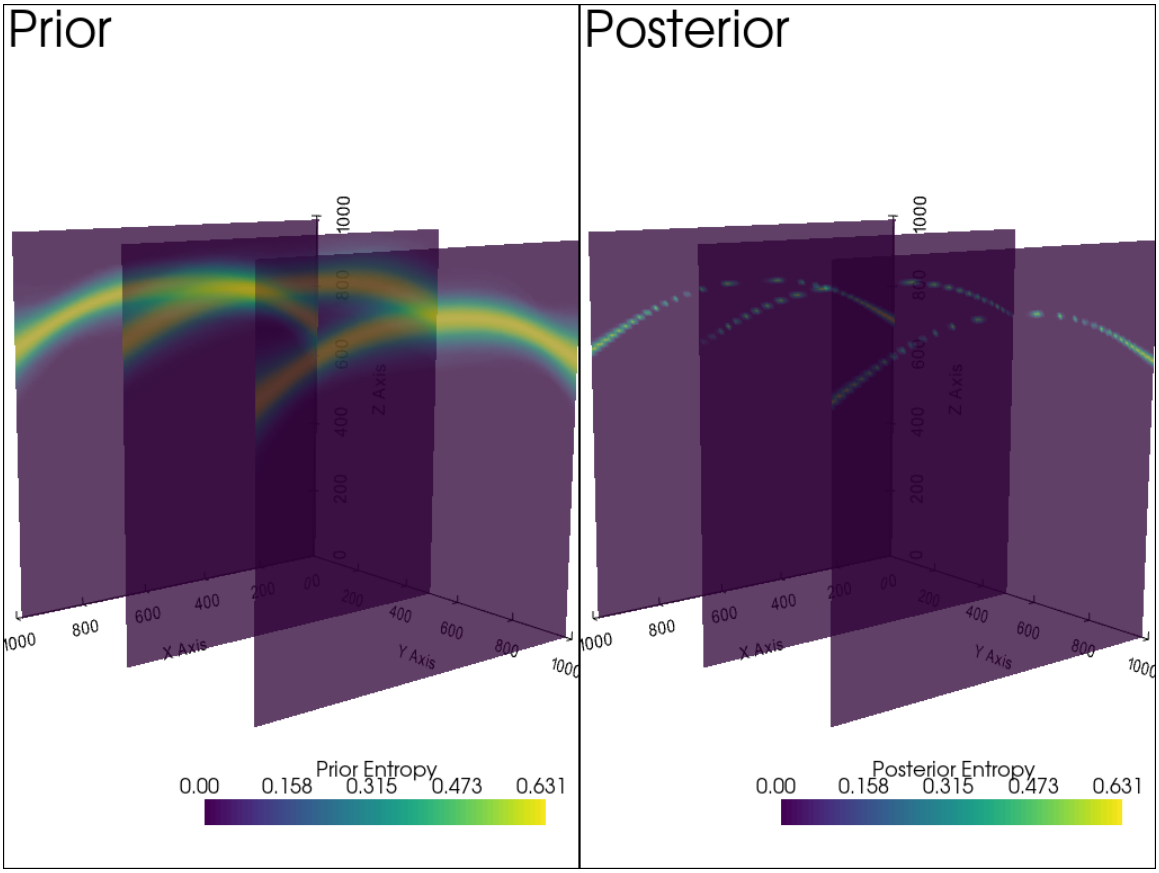

Figure 10: Information entropy plot of the upper boundary of Prior and Posterior. High entropy represents high uncertainties.

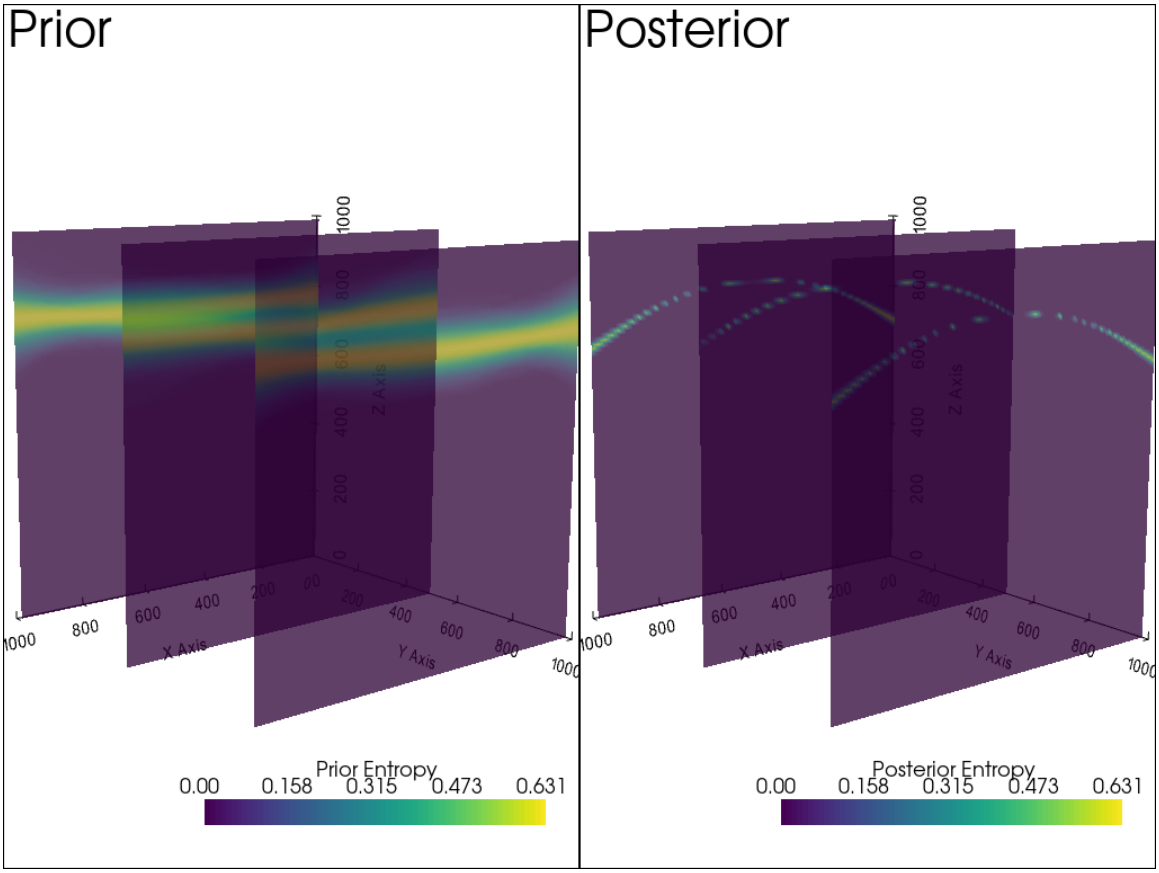

Figure 11: Inference results with uniform Prior distribution. Information entropy plot of the upper boundary of Prior and Posterior. High entropy represents high uncertainties.

the gpCN (Figure 11). Similar inference efficiency can be achieved in this case, with slightly more cost in the MAP searching process. This is because our initial guess drawn from the Prior distribution could deviate more from the target set. Similar posterior results are found in this example because of the low variance in the observation data, which dominates the prior information in the weights of posterior probability.

DISCUSSION

The gpCN method we adopted in this study is a MAP-based method. These MAP-based methods utilize a global Hessian evaluated at the MAP point. The primary advantage of such methods compared with other derivative methods such as HMC is that the Hessian matrix is only required to be evaluated once. It works the best in cases where the posterior is Gaussian-like. In addition to the adoption of pCN and the adapted proposal, gpCN outperforms RMH in high-dimension problems and a highly correlated posterior. Additionally, finding the MAP and starting the MCMC from the MAP saves the burn-in computation normally required in MCMC, because the MAP is already in the target distribution.

However,this global Hessian method is also limited when the model is highly nonlinear and the posterior becomes non-Gaussian, such as a banana distribution (e.g. Lan et al., 2016). Hessian evaluated at the MAP point are not as representative, and therefore, such bad adaptation could lead to a less efficient Markov chain. Such cases can be extended into the state-dependent local approximation with additional computational costs (Petra et al., 2014; Rudolf and Sprungk, 2018). Another limitations is the ability to explore the multi-modes posterior distribution. In these cases, the posterior could be non-convex, and the global minimum is normally not easily found by simple gradient-based optimization methods such as Nadam used in this study. This is a common issue in optimization in many applications (Nocedal and Wright, 2006). Second-derivative optimization methods such as Newton's methods, could be potentially be used to assess the MAP exploration(Villa et al., 2018).

The high computational cost of the high-order derivative is the main obstacle for the application of advanced MCMC methods involving high-order derivatives. The efficient

computation of high-order derivatives remains an open research problem (Laue et al., 2018; Nilsen et al., 2019; Margossian, 2019). In our work, the computation time required for a single Hessian calculation with the current implementation depends on the complexity of the computational graph used in AD and the number of parameters of interest. Although the evaluation of the full Hessian matrix is still more expensive than the gradient evaluation using a single backward propagation in AD, it is significantly faster than computing Hessian by the FD method, regardless of the precision. The numerical analysis is provided in Figure 12. As described previously, AD has an $\mathcal{O}(n)$ complexity, while FD has a complexity of $\mathcal{O}(n^2)$. A full Hessian matrix requires $\frac{n^2+n}{2} \times 4 \times t$ by using FD, where n represents the dimension of the parameter of interest, and t is the time cost for a single forward simulation.

In addition to the computational efficiency, the computational graph employed in AD also allows any variable in the graph to be traced with a derivative with minor modifications to the computer program.

This study is the first step towards the application of Hessian-informed MCMC in geological modeling. With a more efficient implementation of the Hessian evaluation, we can expect a good performance of these methods in higher-dimensional problems. This work has constructed the Hessian explicitly, which is appropriate for the size of problems considered. However this is not necessary, as shown for large-scale problems in Villa et al. (2021); Isaac et al. (2015); Bui-Thanh et al. (2013). The data misfit Hessian (H_{misfit}) admits a low rank approximation (due to ill-posedness), and this approximation can be obtained with $\mathcal{O}(r)$ Hessian-vector products using the randomized SVD algorithm, where r is the effective rank of H_{misfit}. Once this low rank approximation is constructed, the Sherman-Morrison formula can be used to efficiently draw samples from the Laplace approximation, as the gpCN algorithm requires. Thus only $\mathcal{O}(r)$ Hessian-vector products are required, each of

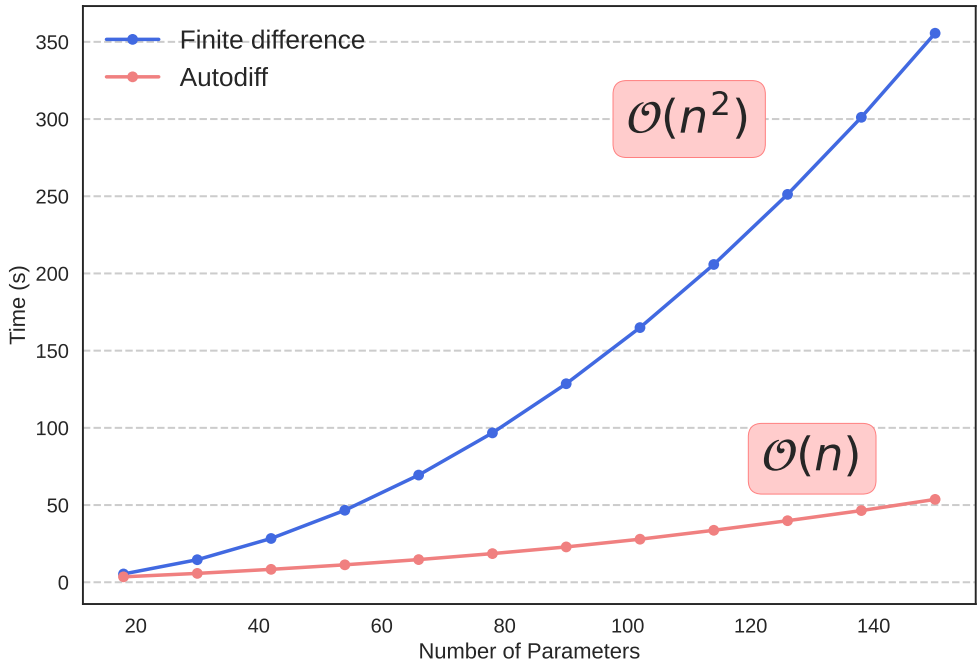

Figure 12: Comparison of computation time of a full Hessian matrix by Finite Difference(FD) method and Auto-Diff(AD). FD computation time is approximated by extrapolating the computation time for a single forward simulation. AD computation time is simulated with the same configuration used in the previous example.

which requires the solution of a pair of forward/adjoint models. This is in contrast with explicitly building the Hessian, which requires $\mathcal{O}(n)$ forward model solutions. The rank r depends on how informative the data are, i.e. how ill-posed the inverse problem is. For many problems, $r << n$, and r is independent of the parameter dimension. Therefore one can benefit greatly from Hessian information, even for high-dimensional parameters.

CONCLUSION

In summary, this study extended the previous development of stochastic geological modeling methods . We used the automatic differential technique implemented in TensorFlow to allow second-order derivative information to be efficiently evaluated in geological models. We applied the recently developed Hessian-informed MCMC, the generalized preconditioned Crank-Nicolson (gpCN), to solve the Bayesian inference problem on a synthetic three-layer geological model. We compared gpCN with the other two commonly used MCMC methods, including RMH and gradient-based method HMC. The results demonstrate with a single Hessian evaluation, gpCN outperforms RMH while preserving the computational efficiency, with no additional computational cost at each step, which has the potential to be generalized to more complex models.

REFERENCES

Abadi, M., P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., 2016, Tensorflow: A system for large-scale machine learning: 12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16), 265–283.

Betancourt, M., 2017, A conceptual introduction to hamiltonian monte carlo: arXiv preprint arXiv:1701.02434.

———, 2018, A geometric theory of higher-order automatic differentiation: arXiv preprint arXiv:1812.11592.

———, 2019, The convergence of markov chain monte carlo methods: from the metropolis method to hamiltonian monte carlo: Annalen der Physik, **531**, 1700214.

Boulanger, O., and M. Chouteau, 2001, Constraints in 3d gravity inversion: Geophysical prospecting, **49**, 265–280.

Bui-Thanh, T., O. Ghattas, J. Martin, and G. Stadler, 2013, A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion: SIAM Journal on Scientific Computing, **35**, A2494–A2523.

Caumon, G., 2010, Towards stochastic time-varying geological modeling: Mathematical Geosciences, **42**, 555–569.

Caumon, G., P. Collon-Drouaillet, C. L. C. De Veslud, S. Viseur, and J. Sausse, 2009, Surface-based 3d modeling of geological structures: Mathematical Geosciences, **41**, 927–945.

Chen, T., E. Fox, and C. Guestrin, 2014, Stochastic gradient hamiltonian monte carlo: International conference on machine learning, 1683–1691.

Cotter, S. L., G. O. Roberts, A. M. Stuart, and D. White, 2013, Mcmc methods for functions: modifying old algorithms to make them faster: Statistical Science, 424–446.

Cui, T., K. J. Law, and Y. M. Marzouk, 2016, Dimension-independent likelihood-informed mcmc: Journal of Computational Physics, **304**, 109–137.

de la Varga, M., A. Schaaf, and F. Wellmann, 2019, Gempy 1.0: open-source stochastic geological modeling and inversion: Geoscientific Model Development.

De La Varga, M., F. Wellmann, and R. Murdie, 2015, Adding geological knowledge to improve uncertain geological models: a bayesian perspective: Geotectonic Research, **97**, 18–20.

de la Varga, M., and J. F. Wellmann, 2016, Structural geologic modeling as an inference problem: A bayesian perspective: Interpretation, **4**, SM1–SM16.

Dozat, T., 2016, Incorporating nesterov momentum into adam.

Duane, S., A. D. Kennedy, B. J. Pendleton, and D. Roweth, 1987, Hybrid monte carlo: Physics letters B, **195**, 216–222.

Fichtner, A., A. Zunino, and L. Gebraad, 2019, Hamiltonian monte carlo solution of tomographic inverse problems: Geophysical Journal International, **216**, 1344–1363.

Geweke, J., and H. Tanizaki, 1999, On markov chain monte carlo methods for nonlinear and non-gaussian state-space models: Communications in Statistics-Simulation and Computation, **28**, 867–894.

Girolami, M., and B. Calderhead, 2011, Riemann manifold langevin and hamiltonian monte carlo methods: Journal of the Royal Statistical Society: Series B (Statistical Methodology), **73**, 123–214.

Güdük, N., M. de la Varga, J. Kaukolinna, and F. Wellmann, 2021, Model-based probabilistic inversion using magnetic data: A case study on the kevitsa deposit: Geosciences,

11, 150.

Hastings, W. K., 1970, Monte carlo sampling methods using markov chains and their applications.

Isaac, T., N. Petra, G. Stadler, and O. Ghattas, 2015, Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the antarctic ice sheet: Journal of Computational Physics, **296**, 348–368.

Kingma, D. P., and J. Ba, 2014, Adam: A method for stochastic optimization: arXiv preprint arXiv:1412.6980.

Lajaunie, C., G. Courrioux, and L. Manuel, 1997, Foliation fields and 3d cartography in geology: principles of a method based on potential interpolation: Mathematical Geology, **29**, 571–584.

Lan, S., T. Bui-Thanh, M. Christie, and M. Girolami, 2016, Emulation of higher-order tensors in manifold monte carlo methods for bayesian inverse problems: Journal of Computational Physics, **308**, 81–101.

Laue, S., M. Mitterreiter, and J. Giesen, 2018, Computing higher order derivatives of matrix and tensor expressions: Advances in Neural Information Processing Systems, 2750–2759.

Law, K. J., 2014, Proposals which speed up function-space mcmc: Journal of Computational and Applied Mathematics, **262**, 127–138.

Liu, J. S., 2008, Monte carlo strategies in scientific computing: Springer Science & Business Media.

Malinverno, A., 2002, Parsimonious bayesian markov chain monte carlo inversion in a non-linear geophysical problem: Geophysical Journal International, **151**, 675–688.

Malinverno, A., and S. Leaney, 2000, A monte carlo method to quantify uncertainty in the

inversion of zero-offset vsp data, *in* SEG Technical Program Expanded Abstracts 2000: Society of Exploration Geophysicists, 2393–2396.

Margossian, C. C., 2019, A review of automatic differentiation and its efficient implementation: Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, **9**, e1305.

Martin, J., L. C. Wilcox, C. Burstedde, and O. Ghattas, 2012, A stochastic newton mcmc method for large-scale statistical inverse problems with application to seismic inversion: SIAM Journal on Scientific Computing, **34**, A1460–A1487.

Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, 1953, Equation of state calculations by fast computing machines: The journal of chemical physics, **21**, 1087–1092.

Mosegaard, K., and A. Tarantola, 1995, Monte carlo sampling of solutions to inverse problems: Journal of Geophysical Research: Solid Earth, **100**, 12431–12447.

Mosser, L., O. Dubrule, and M. J. Blunt, 2020, Stochastic seismic waveform inversion using generative adversarial networks as a geological prior: Mathematical Geosciences, **52**, 53–79.

Muir, J. B., and H. Tkalcic, 2015, Probabilistic joint inversion of lowermost mantle p-wave velocities and core mantle boundary topography using differential travel times and hierarchical hamiltonian monte-carlo sampling: AGUFM, **2015**, S14A–03.

Nabighian, M. N., M. Ander, V. Grauch, R. Hansen, T. LaFehr, Y. Li, W. Pearson, J. Peirce, J. Phillips, and M. Ruder, 2005, Historical development of the gravity method in exploration: Geophysics, **70**, 63ND–89ND.

Nagy, D., 1966, The gravitational attraction of a right rectangular prism: Geophysics, **31**, 362–371.

Neal, R. M., 1993, Bayesian learning via stochastic dynamics: Advances in neural informa-

tion processing systems, 475–482.

Nilsen, G. K., A. Z. Munthe-Kaas, H. J. Skaug, and M. Brun, 2019, Efficient computation of hessian matrices in tensorflow: arXiv preprint arXiv:1905.05559.

Nocedal, J., and S. Wright, 2006, Numerical optimization: Springer Science & Business Media.

Petra, N., J. Martin, G. Stadler, and O. Ghattas, 2014, A computational framework for infinite-dimensional bayesian inverse problems, part ii: Stochastic newton mcmc with application to ice sheet flow inverse problems: SIAM Journal on Scientific Computing, **36**, A1525–A1555.

Qi, Y., and T. P. Minka, 2002, Hessian-based markov chain monte-carlo algorithms.

Roberts, G. O., R. L. Tweedie, et al., 1996, Exponential convergence of langevin distributions and their discrete approximations: Bernoulli, **2**, 341–363.

Rossky, P. J., J. Doll, and H. Friedman, 1978, Brownian dynamics as smart monte carlo simulation: The Journal of Chemical Physics, **69**, 4628–4633.

Rudolf, D., and B. Sprungk, 2018, On a generalization of the preconditioned crank–nicolson metropolis algorithm: Foundations of Computational Mathematics, **18**, 309–343.

Ruggeri, P., J. Irving, and K. Holliger, 2015, Systematic evaluation of sequential geostatistical resampling within mcmc for posterior sampling of near-surface geophysical inverse problems: Geophysical Journal International, **202**, 961–975.

Sambridge, M., T. Bodin, K. Gallagher, and H. Tkalčić, 2013, Transdimensional inference in the geosciences: Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, **371**, 20110547.

Sambridge, M., and K. Mosegaard, 2002, Monte carlo methods in geophysical inverse problems: Reviews of Geophysics, **40**, 3–1.

588 Scalzo, R., D. Kohn, H. Olierook, G. Houseman, R. Chandra, M. Girolami, and S. Cripps,
589 2019, Efficiency and robustness in monte carlo sampling for 3-d geophysical inversions
590 with obsidian v0. 1.2: setting up for success: Geoscientific Model Development, **12**, 2941–
591 2960.
592 Sen, M. K., and P. L. Stoffa, 1996, Bayesian inference, gibbs' sampler and uncertainty
593 estimation in geophysical inversion 1: Geophysical Prospecting, **44**, 313–350.
594 Tarantola, A., 2005, Inverse problem theory and methods for model parameter estimation:
595 SIAM.
596 Tarantola, A., and B. Valette, 1982, Generalized nonlinear inverse problems solved using
597 the least squares criterion: Reviews of Geophysics, **20**, 219–232.
598 Tikhonov, A. N., and V. Y. Arsenin, 1977, Solutions of ill-posed problems: New York, 1–30.
599 Villa, U., N. Petra, and O. Ghattas, 2018, hippylib: An extensible software framework for
600 large-scale inverse problems: Journal of Open Source Software, **3**, 940.
601 ———, 2021, hippylib: An extensible software framework for large-scale inverse problems
602 governed by pdes: Part i: Deterministic inversion and linearized bayesian inference: ACM
603 Transactions on Mathematical Software (TOMS), **47**, 1–34.
604 Wellmann, F., and G. Caumon, 2018, 3-d structural geological models: Concepts, methods,
605 and uncertainties: **59**, 1–121.
606 Wellmann, J. F., M. De La Varga, R. E. Murdie, K. Gessner, and M. Jessell, 2018, Un-
607 certainty estimation for a geological model of the sandstone greenstone belt, western
608 australia–insights from integrated geological and geophysical inversion in a bayesian in-
609 ference framework: Geological Society, London, Special Publications, **453**, 41–56.
610 Wellmann, J. F., and K. Regenauer-Lieb, 2012, Uncertainties have a meaning: Information
611 entropy as a quality measure for 3-d geological models: Tectonophysics, **526**, 207–216.

612 Witter, J. B., W. J. Trainor-Guitton, and D. L. Siler, 2019, Uncertainty and risk evaluation
613 during the exploration stage of geothermal development: A review: Geothermics, **78**,
614 233–242.