A tutorial introduction to the Hamiltonian Monte Carlo solution of weakly nonlinear inverse problems

Andreas Fichtner[1], Andrea Zunino[2], and Lars Gebraad[1]

[1]ETH Zürich. Institute of Geophysics, Sonneggstrasse 5. 8092, Zürich
[2]Niels Bohr Institute, University of Copenhagen, Denmark

Correspondence to: Andreas Fichtner (andreas.fichtner@erdw.ethz.ch)

Abstract. We present an introduction to Hamiltonian Monte Carlo (HMC) sampling of high-dimensional model spaces with focus on linear and weakly nonlinear inverse problems. This includes the theoretical foundations, intuitive examples, as well as applications to linear and nonlinear, adjoint-based traveltime tomography.

HMC rests on the construction of an artificial Hamiltonian system where a model is treated as a high-dimensional particle moving along a trajectory in an extended model space. Using derivatives of the forward equations, HMC is able to make long-distance moves from the current towards a new model, thereby promoting sample independence while maintaining high acceptance rates. Though HMC is well-known and appreciated in many branches of the physical and computational sciences, it is only starting to be explored for the solution of geophysical inverse problems.

The first part of this article is therefore dedicated to an introduction to HMC using common geophysical terminology. Based on a brief review of equilibrated random walks, we provide a self-contained proof of the HMC algorithm. This leads us to consider simple examples with linear forward problems and Gaussian priors. Having analytical solutions, these problems are not the primary target of Monte Carlo methods. However, they allow us to gain intuition and to study critical tuning issues, the independence of successive samples, the tractable number of dimensions, and limitations of the algorithm.

The second part covers applications to linear traveltime tomography, reflector imaging, and nonlinear traveltime tomography. The latter rests on a fast-sweeping method combined with adjoint techniques for the computation of derivatives. Using HMC allows us to solve the tomographic problem fully probabilistically for several thousand unknowns, and without the requirement of supercomputing resources.

1 Introduction

The inference of physical properties that are not directly observable, such as elastic moduli inside the Earth or future atmospheric conditions, is the central theme of geophysics. It has defined the science since its entry into the quantitative era, when instrumental observations and predictive mathematical theories became available (e.g. Mallet, 1861; von Rebeur-Paschwitz, 1889; Oldham, 1906). The work of Backus and Gilbert (1968, 1970) provided an early formalisation of indirect geophysical inference, contributing to the establishment of inverse theory as a science in its own right.

1.1 Bayesian inference and model space sampling

Geophysical inversion is a delicate balancing act between elegant universality and computational pragmatism. Sparse and erroneous data combined with imperfections of the forward modelling theory make inverse problems ill-posed in the sense that an ensemble of plausible models explains the data to within their observational uncertainties. Solving an inverse problem means to explore and characterise the ensemble (Backus and Gilbert, 1968), which can be achieved most comprehensively by computing the posterior probability density in model space using Bayes' theorem (Bayes and Price, 1763; Jaynes, 2003; Tarantola, 2005; Sen and Stoffa, 2013). While being beautifully simple, Bayes' theorem mostly requires sampling of the model space to approximate the posterior. Depending on its dimension and the numerical cost of solving the forward problem, sampling can be prohibitively expensive. Therefore, the majority of inverse problems is simplified, often to the estimation of the maximum-likelihood model and the posterior covariances, which may, however, not represent the complete plausible ensemble. Possible consequences are the interpretation of meaningless models, and unrealistic uncertainties that hamper decision making and the comparison with earlier results.

Since the early work of Keilis-Borok and Yanovskaya (1967) and Press (1968), the number of geophysical inverse problems that can be solved through Bayesian inference has grown continuously, thanks to increasing computational power and the development of advanced sampling strategies that improve the classical Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970; Mosegaard and Tarantola, 1995). The latter include genetic algorithms (e.g. Holland, 1975; Gallagher et al., 1991; Sen and Stoffa, 1992; Sambridge and Drijkoningen, 1992), parallel tempering (e.g. Geyer, 1991; Marinari and Parisi, 1992; Geyer and Thompson, 1995; Sambridge, 2014), the neighbourhood algorithm (Sambridge, 1999a, b), and the reversible-jump algorithm used for transdimensional inversion (e.g. Green, 1995; Sambridge et al., 2006; Bodin and Sambridge, 2009; Sambridge et al., 2013).

A landmark in the development of sampling methods was the recognition, formalised in a series of 'No-Free-Lunch theorems' (e.g. Wolpert and Macready, 1997; Mosegaard, 2012), that all algorithms are from the outset equally inefficient. Superiority of an algorithm can result only from the injection of prior knowledge in a broad sense. This may include, for instance, knowledge on the nature of the forward problem or the likely topology of the misfit surface, which may enable specific tuning strategies. Conversely, specific algorithms may be developed to work efficiently for inverse problems that share certain properties.

Within this context, Hamiltonian Monte Carlo (HMC) was developed for problems where the derivative of the target probability density with respect to the model parameters can be computed quickly (e.g. Neal, 2011; Betancourt, 2017). Using derivatives, in contrast, for instance, to the Metropolis-Hastings algorithm, helps to focus on those regions of model space that are more plausible, thus wasting less samples. Since its first formulation for applications in lattice quantum chromodynamics (Duane et al., 1987), HMC has found widespread use in neural networks and machine learning (Neal, 1996; Bishop, 2006; Gutmann and Hyvärinen, 2012), molecular simulations (e.g. Dubbledam et al., 2016), nuclear physics (e.g. Elhatisari et al., 2015), genomics (e.g. Honkela et al., 2015), signal processing (e.g. Wei et al., 2015), and quantum mechanics (e.g. Seah et al., 2015).

Numerous geophysical inverse problems for which first derivatives are easy to compute also fall within the range of applicabil-

ity of HMC. These include, for instance, traveltime tomography (e.g. Aki et al., 1976; Iyer and Hirahara, 1993; Nolet, 2008), gravity inversion (e.g. Blakely, 2002; Li and Oldenburg, 1998), seismic source inversion (e.g. Dziewoński et al., 1981; Ekström et al., 2012; Lee et al., 2014), and inversions for the Gauss coefficients of the geomagnetic field (e.g. Blakely, 2002; Jackson et al., 2000). Nevertheless, the potential of HMC for geophysical inversion is only slowly being recognised (Muir and Tkalčić, 2015; Biswas and Sen, 2017; Sen and Biswas, 2017).

1.2 Objectives and outline

The objectives of this manuscript are twofold: First, we aim to provide a detailed and pedagogical introduction to HMC with focus on linear or mildly nonlinear inverse problems. Despite the availability of literature in other fields, we believe that the use of geophysical terminology and the link to concrete problems will facilitate future applications of HMC in tomography. Second, we will provide various examples, ranging from simple linear problems with Gaussian priors to a novel, fully nonlinear traveltime tomography. The examples are intended to highlight the potential of HMC to solve high-dimensional problems and to expose critical limitations.

Throughout this work, our focus will be on inverse problems where the forward problem is nearly linear, mostly because efficient tuning can be derived for such cases. This does not mean, however, that HMC cannot work in a more general setting. The term 'nearly linear' is deliberately diffuse because the boundary between efficient and inefficient tuning is unlikely to be sharp, and may be shifted, as so often, by some level of trial and error.

This manuscript is organised as follows: Section 2 establishes the necessary notation and background, as well as some key motivations for the development of HMC. This is followed, in section 3, by the description and proof of the HMC algorithm, complemented by remarks on tuning and HMC variants. A series of examples is provided in section 4. These are intended to illustrate key properties, the potential, and critical limitations of HMC. Finally, the discussion in section 5 is focused on the potentials and limitations of HMC, also in the context of the No-Free-Lunch Theorems.

2 Preliminaries

To set the stage for the description of Hamiltonian Monte Carlo and for pedagogical reasons, we provide a brief review of Bayesian inference, random walks and the Metropolis rule. Readers familiar with these topics may wish to skip this section.

2.1 Bayesian inference

Bayesian inference describes information on model parameters $\mathbf{m} \in \mathbb{M}$ and data $\mathbf{d} \in \mathbb{D}$ in terms of probability density functions (pdfs) ρ (e.g. Tarantola and Valette, 1982; Tarantola, 2005), where $\mathbb{M}$ and $\mathbb{D}$ denote the model and data spaces, respectively. Prior information in model space, $\rho_0(\mathbf{m})$, captures all constraints on model parameters available prior to the analysis of any data, and the prior in data space, $\rho_d(\mathbf{d}|\mathbf{m})$, is the likelihood of a computed data vector $\mathbf{d}$ given some $\mathbf{m}$. The data space prior or likelihood function thus contains the forward problem equations, linking the model and data spaces. Based on the priors $\rho_0(\mathbf{m})$

and $\rho_d(\mathbf{d}|\mathbf{m})$, Bayes' theorem gives the posterior pdf in model space, $\rho_m(\mathbf{m}|\mathbf{d})$, which encapsulates all available information on the model parameters that can be deduced from prior knowledge,

$$\rho_m(\mathbf{m}|\mathbf{d}) = k\,\rho_d(\mathbf{d}|\mathbf{m})\rho_0(\mathbf{m})\,. \tag{1}$$

The constant k normalises $\rho_m(\mathbf{m}|\mathbf{d})$ so that $\int_{\mathbb{M}} \rho_m(\mathbf{m}|\mathbf{d})\,d\mathbf{m} = 1$. Since the potentially high-dimensional posterior may be difficult to interpret, we are mostly interested in lower-dimensional quantities computed from $\rho_m(\mathbf{m}|\mathbf{d})$ via an integral over model space. These may include the probability $P(\mathbf{m} \in \mathbb{M}_{\text{sub}}) = \int_{\mathbb{M}_{\text{sub}}} \rho_m(\mathbf{m}|\mathbf{d})\,d\mathbf{m}$ of $\mathbf{m}$ being inside a model space subvolume $\mathbb{M}_{\text{sub}} \subset \mathbb{M}$, the expectation $E(\mathbf{m})$ defined as

$$E(\mathbf{m}) = \int_{\mathbb{M}} \mathbf{m}\,\rho_m(\mathbf{m}|\mathbf{d})\,d\mathbf{m}\,, \tag{2}$$

or higher moments and marginal pdfs. All of these and other quantities may generically be written as $\int_{\mathbb{M}} \Phi(\mathbf{m})\,\rho_m(\mathbf{m}|\mathbf{d})\,d\mathbf{m}$, with a suitably chosen function $\Phi(\mathbf{m})$ (e.g. Chen et al., 2012). In most cases, and particularly in high-dimensional nonlinear problems, an analytical treatment of the integrals is not possible, thus motivating alternative strategies. One solution is to employ Monte Carlo methods that draw random samples such that the sampling density is proportional to $\rho_m(\mathbf{m}|\mathbf{d})$. The samples may then be used to approximate the integrals. Taking the example of the expectation from Eq. (2), we can compute the empirical expectation

$$E_N(\mathbf{m}) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{m}^i \rho_m(\mathbf{m}^i|\mathbf{d})\,, \tag{3}$$

from $N > 0$ random samples $\mathbf{m}^i$. As N approaches infinity, the empirical expectation approaches the true expectation, that is $E(\mathbf{m}) = \lim_{N\to\infty} E_N(\mathbf{m})$. In practice, the number of samples N is finite, and E_N approximates E only to within some error. When the samples $\mathbf{m}^i$ are statistically independent, the standard deviation of the approximation error is given by

$$\sigma = \frac{\sigma_N}{\sqrt{N}}\,, \tag{4}$$

where σ_N is the variance of the samples (MacKay, 2003). Interestingly, σ is independent of the number of dimensions, therefore a small number of samples may be sufficient to provide a reliable estimation of the expectation E. While (4) is valid for statistically independent samples, frequently employed Markov-chain Monte Carlo (MCMC) methods may produce rather correlated samples, which typically leads to slower convergence. One of the main goals of HMC is thus to improve the independence of samples.

2.2 Random walks

MCMC methods produce samples via random walks, that we summarise briefly. To reduce the notational burden, we introduce a lighter terminology. For this, we partition the model space $\mathbb{M}$ into subsets $\mathbb{M}^i$ with equal volumes M around models $\mathbf{m}^i$. Assuming the $\mathbb{M}^i$ are sufficiently small, the model prior is $P_0^i = M\rho_0(\mathbf{m}^i)$. Similarly, we define $P_d^i = M\rho_d(\mathbf{d}|\mathbf{m}^i)$ for the discretised likelihood function, and $P_m^i = M\rho_m(\mathbf{m}^i|\mathbf{d})$ for the discretised posterior. Furthermore, we denote by $P^i = M\rho(\mathbf{m}^i)$

a generic probability with density ρ.

During a random walk, a sequence of models $\{\mathbf{m}^i\}_{i=1}^N$ is visited randomly according to some rules. When the rules are designed such that the number of visits to $\mathbb{M}^i$ is proportional to P^i, the walk is said to sample the distribution ρ. Assume the walker is currently at $\mathbf{m}^j$. The conditional probability of going from $\mathbf{m}^j$ to $\mathbf{m}^i$ is called the transition probability $T^{i \leftarrow j} = T(\mathbb{M}^i | \mathbb{M}^j)$.

5   Thus, if the probability for currently being at $\mathbf{m}^j$ is P^j, the unconditional probability of walking to $\mathbf{m}^i$ is $T^{i \leftarrow j} P^j$. Since the walker must go somewhere, including the option of staying at $\mathbf{m}^j$, the transition probabilities satisfy

$$\sum_i T^{i \leftarrow j} = 1 \,. \tag{5}$$

Our particular interest is in random walks that equilibrate asymptotically as $N \to \infty$. Equilibrium means that the combined probability to arrive at $\mathbf{m}^i$ from any $\mathbf{m}^j$ is equal to the probability of actually being at $\mathbf{m}^i$, that is

10   $$\sum_j T^{i \leftarrow j} P^j = P^i \,. \tag{6}$$

Eq. (6) can be understood in terms of flows. For this, we interpret P^i as the number of some imaginary particles located at $\mathbf{m}^i$. In this sense, $T^{i \leftarrow j} P^j$ is the number of particles flowing from a specific $\mathbf{m}^j$ to $\mathbf{m}^i$, and Eq. (6) states that the total number of particles flowing from all $\mathbf{m}^j$ into $\mathbf{m}^i$ is equal to the number of particles at $\mathbf{m}^i$. A sufficient condition for equilibration at the distribution P^i is detailed balance of the flow, meaning that as many particles flow from $\mathbf{m}^j$ to $\mathbf{m}^i$ as from $\mathbf{m}^i$ to $\mathbf{m}^j$, that is

15   $$T^{i \leftarrow j} P^j = T^{j \leftarrow i} P^i \,. \tag{7}$$

Indeed, summing Eq. (7) over j, and using Eq. (5), we retrieve the equilibration condition (6).

2.3   The Metropolis rule and the Metropolis-Hastings algorithm

Our interest is to sample the posterior ρ_m. For this, we assume that we already have an algorithm that samples the prior ρ_0 such that detailed balance, $T^{i \leftarrow j} P_0^j = T^{j \leftarrow i} P_0^i$, is satisfied. Given that the walker is currently at $\mathbf{m}^j$, we use this algorithm to propose

20   a transition to $\mathbf{m}^i$. In this context, the prior ρ_0 serves as proposal density, and we accept or reject the proposed transition based on an acceptance criterion designed such that the walker samples ρ_m. More formally, the proposal to move from $\mathbf{m}^j$ to $\mathbf{m}^i$ is given by the transition probability $T^{i \leftarrow j}$. Whether the proposal is actually accepted, depends on the acceptance probability $A^{i \leftarrow j}$ that modifies the prior transition probability $T^{i \leftarrow j}$ to yield a new transition probability $\hat{T}^{i \leftarrow j}$,

$$\hat{T}^{i \leftarrow j} = A^{i \leftarrow j} T^{i \leftarrow j} \,. \tag{8}$$

25   We aim to choose $A^{i \leftarrow j}$ such that $\hat{T}^{i \leftarrow j}$ produces an equilibrium distribution equal to ρ_m^i. For this, we require the detailed balance

$$\hat{T}^{i \leftarrow j} P_m^j = \hat{T}^{j \leftarrow i} P_m^i \,. \tag{9}$$

Inserting Eq. (8) into (9), and taking into account Bayes' theorem, $MP_m^i = kP_d^i P_0^i$, we obtain the condition,

$$\frac{A^{i \leftarrow j}}{A^{j \leftarrow i}} = \frac{P_d^i}{P_d^j}. \tag{10}$$

While there are infinitely many acceptance probabilities that satisfy (10), the Metropolis rule (Metropolis et al., 1953)

$$A^{i \leftarrow j} = \min\left(1, P_d^i / P_d^j\right), \tag{11}$$

5   is employed most frequently. According to the Metropolis rule, the transition from $\mathbf{m}^j$ to $\mathbf{m}^i$ is always accepted when the likelihood of the new model, P_d^i, is higher than the likelihood of the current model, P_d^j ($A^{i \leftarrow j} = 1$). When $P_d^i < P_d^j$, however, there is still a chance to make the transition, namely with probability $A^{i \leftarrow j} = P_d^i / P_d^j$. The Metropolis rule is optimal in the sense of being as exploratory as possible (Peskun, 1973). Indeed, requesting that a transition from $\mathbf{m}^i$ to $\mathbf{m}^j$ is always accepted when the likelihood increases, that is $A^{j \leftarrow i} = 1$ for $P_d^j > P_d^i$, Eq. (10) transforms to $A^{i \leftarrow j} = P_d^i / P_d^j < 1$, which is equivalent
10   to the Metropolis rule (11).

The extended Metropolis-Hastings algorithm (EMH) (Mosegaard and Tarantola, 1995) implements a particular case of the general Metropolis-Hastings random walk (Hastings, 1970), where transitions are accepted based on the Metropolis rule, and where the prior $\rho_0(\mathbf{m})$ is used as proposal density, that is, $T^{i \leftarrow j} = P^i$. This ensures that models that are a priori more likely are visited more frequently.

15   In the following, we will use well-documented deficiencies of EMH to motivate the development of HMC. Furthermore, since EMH is widely kown and easy to implement, it will serve for illustration and as a baseline algorithm for comparison. In this context we note that HMC could of course be compared to other algorithms that may perform better than EMH in specific scenarios.

2.4   A toy example

20   For illustration, we consider a simple example that we will reuse later as reference for HMC. We assume a ten-dimensional model vector $\mathbf{m}$ and a Gaussian model space prior

$$\rho_0(\mathbf{m}) = \text{const. } \exp\left(-\frac{1}{2}\mathbf{m}^T \mathbf{C}_M^{-1} \mathbf{m}\right). \tag{12}$$

Furthermore, we let the forward problem be linear, $\mathbf{d} = \mathbf{Gm}$, where we define $\mathbf{G}$ to be a diagonal matrix with entries $G_{ii} = i/10$. The data vector components are set to $d_i^{\text{obs}} = i/5$. Assuming Gaussian measurement errors, we obtain the likelihood
25   function

$$\rho_d(\mathbf{d}|\mathbf{m}) = \text{const. } \exp\left(-\frac{1}{2}(\mathbf{d}^{\text{obs}} - \mathbf{Gm})^T \mathbf{C}_D^{-1}(\mathbf{d}^{\text{obs}} - \mathbf{Gm})\right). \tag{13}$$

For simplicity, we (arbitrarily) set the prior model and data covariance matrices to $\mathbf{C}_M = \sigma_M^2 \mathbf{I}$ and $\mathbf{C}_D = \mathbf{I}$, respectively. While the posterior mean and covariance could easily be found analytically for this example, we employ EMH for the purpose of illustration. Running EMH for $\sigma_M = 2$ with 500000 samples produces the histograms of the posterior in Fig. 1. As expected,
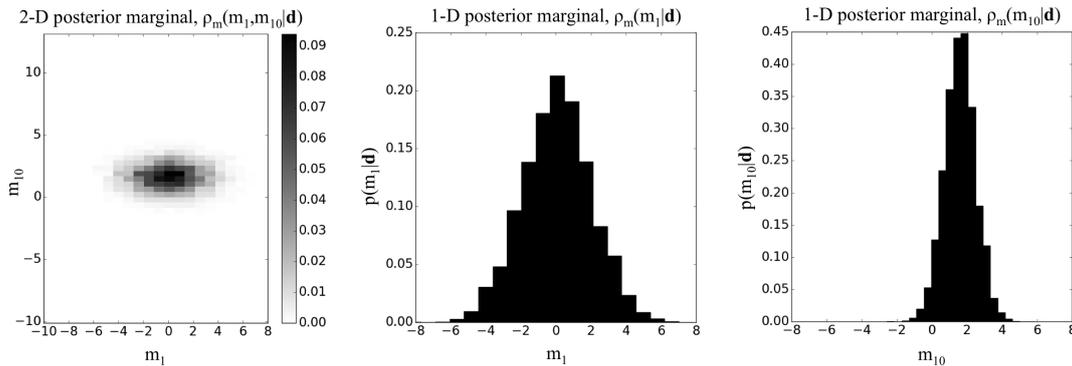
Figure 1. Two- and one-dimensional marginals of the posterior defined through the priors in equations (12) and (13), with $\sigma_M = 2$. Since d_1 depends only weakly on m_1 through the small value of G_{11}, we find that m_1 is poorly constrained. Its marginal posterior is nearly identical to the prior. In contrast, m_{10} is well constrained to be approximately 2. Better approximations of the posterior may be obtained using more random samples.

m_1 is less constrained than m_{10} because the small G_{11} allows for large variations in m_1 that still fall within the measurement error.

3   Hamiltonian Monte Carlo

3.1   Motivating example

5   The advantage of EMH lies in its ability to account for prior knowledge, the quality of which critically affects the efficiency of the algorithm. When the prior is already a good approximation of the posterior, sampling is very efficient, and a good approximation of the posterior can be obtained with a small number of samples, and vice versa. The magnitude of this effect is illustrated in Fig. 2, which uses the toy problem of section 2.4. The left panel shows the trajectory of the walker projected onto the (m_1, m_{10}) plane when the standard deviation of the prior is $\sigma_M = 1$. Each red dot represents an accepted test model. There

10   are 60176 of these dots, meaning that around 12 % of the 500000 proposed samples have been accepted. Weakening the prior by increasing its standard deviation to $\sigma_M = 3$, causes the proposed samples to scatter more widely. Therefore, more samples fail the Metropolis rule, and only ~ 7121 samples are accepted; ~ 1.4 %. Further weakening the prior to $\sigma_M = 10$, results in only 7 accepted models (on average over many runs), that is ~ 0.0014 %. The walker thus remains at a certain position for a long time before a proposed model passes the Metropolis test and the model space is further explored.

15   This undesirable phenomenon is a consequence of the curse of dimensionality (e.g. Tarantola, 2005). In a high-dimensional model space, the volume of relevant models defined by the posterior becomes excessively small compared to the volume of the whole model space. If the proposal density, that is the prior in the case of EMH, does not limit the search space very
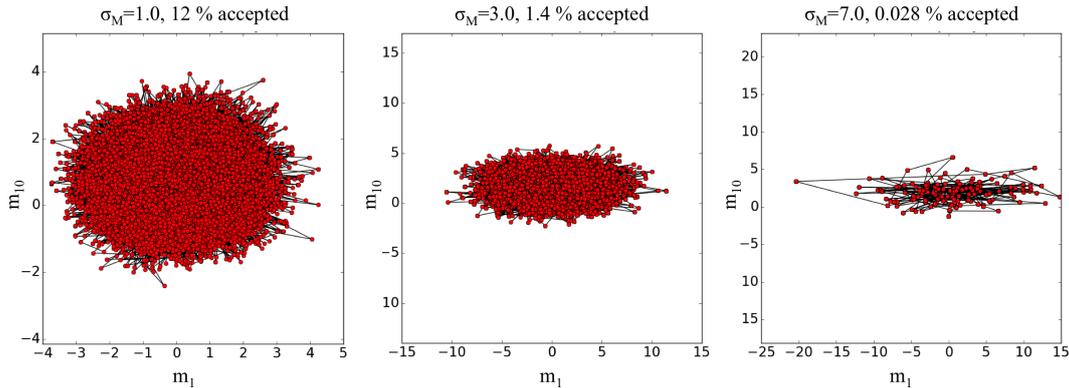
Figure 2. Trajectory of the walker for the toy example from section 2.4 for three different standard deviations of the Gaussian prior, $\sigma_M = 1$, $\sigma_M = 3$, and $\sigma_M = 7$. Each red dot represents an accepted model projected onto the (m_1, m_{10}) plane in the ten-dimensional model space. The total number of proposed samples is 500000. As the prior standard deviation increases, the number of accepted samples decreases dramatically, leading to a very poor exploration of the model space and an enormous waste of computing time.

significantly, the probability of finding a new model that is better than an already relatively good one becomes negligibly small. Therefore, the walker will mostly stay at its current position

3.2 The Hamiltonian Monte Carlo concept

HMC is designed to reduce problems related to dependent samples and low acceptance rates. In contrast to being drawn ran-
5 domly, a new proposal is already influenced by the data. Starting from the current model $\mathbf{m}$, the walker follows a Hamiltonian trajectory towards a new test model $\mathbf{m}_\tau$. The trajectory of the model as a function of an artificially introduced time variable t is controlled by Hamilton's equations (e.g. Symon, 1971; Landau and Lifshitz, 1976)

$$\frac{dm_i}{dt} = \frac{\partial H}{\partial p_i}, \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial m_i}, \tag{14}$$

with $i = 1, ..., N$. In Hamiltonian dynamics, m_i plays the role of a generalised position, and p_i is a generalised momentum
10 defined through

$$p_i = \sum_{j=1}^{n} M_{ij} \frac{dm_j}{dt}, \tag{15}$$

with a symmetric and positive definite mass matrix M_{ij}. The Hamiltonian $H(\mathbf{m}, \mathbf{p})$ is the total energy of the particle, that is $H = K + U$ with the kinetic energy $K(\mathbf{m}, \mathbf{p})$ and the potential energy $U(\mathbf{m}, \mathbf{p})$. The momenta $\mathbf{p}$ are elements of momentum space $\mathbb{P}$. Together, $\mathbf{m}$ and $\mathbf{p}$ are in phase space $\mathbb{X} = \mathbb{M} \times \mathbb{P}$. In the Appendix we review several properties of Hamiltonian
15 systems that are essential for HMC, including time reversability, volume preservation, and the preservation of phase space partitionings.

The basic concept of HMC is to sample an auxiliary distribution in the $2n-$dimensional phase space, the canonical distribution, from which samples of the posterior can be obtained by ignoring (or marginalising over) the momentum space component. Defining kinetic and potential energies as

$$K(\mathbf{p}) = \frac{1}{2} \sum_{i,j=1}^{n} p_i M_{ij}^{-1} p_j \,, \tag{16}$$

and

$$U(\mathbf{m}) = -\log \rho_m(\mathbf{m}|\mathbf{d}) \,, \tag{17}$$

the canonical distribution is defined as

$$\rho_c(\mathbf{m},\mathbf{p}) = \exp\left[-H(\mathbf{m},\mathbf{p})\right] = \exp\left[-U(\mathbf{m}) - K(\mathbf{p})\right] = \rho_m(\mathbf{m}|\mathbf{d}) \exp\left(-\frac{1}{2} \sum_{i,j=1}^{n} p_i M_{ij}^{-1} p_j\right) . \tag{18}$$

In other words, the canonical distribution is the product of the posterior $\rho_m(\mathbf{m}|\mathbf{d})$ and a Gaussian distribution of the momenta $\mathbf{p}$. If we manage to sample $\rho_c(\mathbf{m},\mathbf{p})$, we may simply take the model space component of the samples to obtain samples of the posterior.

3.3 The algorithm

The step of a random walker in HMC consists of two phases: First, momenta p_i are drawn from their Gaussian distribution

$$\rho_p(\mathbf{p}) = \text{const.} \exp\left(-\frac{1}{2} \sum_{i,j=1}^{n} p_i M_{ij}^{-1} p_j\right) \,, \tag{19}$$

as previously defined in Eq. (18). Second, the current state $(\mathbf{m},\mathbf{p})$ is propagated for some time τ along a Hamiltonian trajectory by numerically solving Hamilton's equations (A1). The new state $(\mathbf{m}_\tau, \mathbf{p}_\tau)$ is then accepted with probability

$$\min\left[1, \frac{\rho_c(\mathbf{m}_\tau,\mathbf{p}_\tau)}{\rho_c(\mathbf{m},\mathbf{p})}\right] = \min\left[1, \frac{\rho_m(\mathbf{m}_\tau|\mathbf{d})}{\rho_m(\mathbf{m}|\mathbf{d})} \exp\left(-K(\mathbf{p}_\tau) + K(\mathbf{p})\right)\right] = \min\left[1, \exp\left(-H(\mathbf{p}_\tau) + H(\mathbf{p})\right)\right] \,, \tag{20}$$

which can be seen as a modified version of the Metropolis rule (11). If the transition is accepted, the new model $(\mathbf{m}_\tau, \mathbf{p}_\tau)$ is counted; otherwise, the original model $(\mathbf{m},\mathbf{p})$. Then, the walker's next step starts again with the first phase, and the procedure is repeated. To obtain samples of the posterior, the momentum component of the phase space samples is simply ignored.

Equation (20) has the interesting implication that all proposals would be accepted if we were able to solve Hamilton's equations analytically. In practice, this is not generally the case, and so the equations need to be solved numerically. In choosing a specific numerical integrator, care must be taken to avoid breaking those properties of a Hamiltonian system upon which the proof of the HMC algorithm relies critically. As we will see in section 3.4, these are time reversability and volume preservation. Numerical integrators with these properties are called symplectic integrators. The most commonly used one, at least in the context of HMC, is the leap-frog method, briefly explained in section 3.6.1. As many other symplectic integrators, the leap-frog method does not preserve energy exactly. As a consequence, the acceptance rate of HMC typically differs from the theoretical ideal, which is 1.

3.4 Detailed balance

To demonstrate that ρ_c is indeed the equilibrium distribution of HMC, we need to assert reversibility of the transition in phase space in the sense of the detailed balance equation (7). Thus, for two phase space points $(\mathbf{m}^i, \mathbf{p}^i) \in \mathbb{X}^i$ and $(\mathbf{m}^j, \mathbf{p}^j) \in \mathbb{X}^j$, we must show

$$T(\mathbb{X}^j | \mathbb{X}^i) P_c(\mathbb{X}^i) = T(\mathbb{X}^i | \mathbb{X}^j) P_c(\mathbb{X}^j), \tag{21}$$

for any $\mathbb{X}^i$ and $\mathbb{X}^j$ that partition the phase space. As shown in the Appendix, all images of $\mathbb{X}^i$ under movement along a Hamiltonian trajectory partition the phase space as well. This implies that $(\mathbf{m}^j, \mathbf{p}^j)$ must be inside some $\mathbb{X}^j_\tau$, and is therefore equal to some particular image, that is $(\mathbf{m}^j, \mathbf{p}^j) = (\mathbf{m}^j_\tau, \mathbf{p}^j_\tau)$. Since Hamiltonian dynamics uniquely translates $\mathbb{X}^i$ into its image $\mathbb{X}^i_\tau$, the transition probability $T(\mathbb{X}^j_\tau | \mathbb{X}^i)$ must be zero unless $i = j$. Furthermore, since Hamiltonian dynamics is time reversible (see Appendix), $T(\mathbb{X}^i | \mathbb{X}^j_\tau)$ is also zero unless $i = j$. In summary,

$$T(\mathbb{X}^j_\tau | \mathbb{X}^i) = T(\mathbb{X}^i | \mathbb{X}^j_\tau) = 0, \quad \text{if } i \neq j. \tag{22}$$

Thus, Eq. (21) is trivially satisfied for $i \neq j$, and it remains to consider the more special case

$$T(\mathbb{X}^i_\tau | \mathbb{X}^i) P_c(\mathbb{X}^i) = T(\mathbb{X}^i | \mathbb{X}^i_\tau) P_c(\mathbb{X}^i_\tau). \tag{23}$$

Under the assumption that our partitioning is so fine that we can consider the canonical distribution constant within each $\mathbb{X}^i$, we have

$$P_c(\mathbb{X}^i) = X \rho_c(\mathbf{m}^i, \mathbf{p}^i), \quad \text{and} \quad P_c(\mathbb{X}^i_\tau) = X \rho_c(\mathbf{m}^i_\tau, \mathbf{p}^i_\tau). \tag{24}$$

Note that Eq. (24) only holds because Hamiltonian mechanics preserves the volume X along the trajectory (also see Appendix). With the help of Eq. (20), which determines the transition probability, Eq. (23) now transforms to

$$\min\left[1, \frac{\rho_c(\mathbf{m}^i_\tau, \mathbf{p}^i_\tau)}{\rho_c(\mathbf{m}^i, \mathbf{p}^i)}\right] \rho_c(\mathbf{m}^i, \mathbf{p}^i) = \min\left[1, \frac{\rho_c(\mathbf{m}^i, \mathbf{p}^i)}{\rho_c(\mathbf{m}^i_\tau, \mathbf{p}^i_\tau)}\right] \rho_c(\mathbf{m}^i_\tau, \mathbf{p}^i_\tau). \tag{25}$$

That Eq. (25), and therefore detailed balance, indeed holds, can now be shown by considering the two cases $\rho_c(\mathbf{m}^i_\tau, \mathbf{p}^i_\tau) > \rho_c(\mathbf{m}^i, \mathbf{p}^i)$ and $\rho_c(\mathbf{m}^i_\tau, \mathbf{p}^i_\tau) \leq \rho_c(\mathbf{m}^i, \mathbf{p}^i)$.

3.5 Illustration: Linear problems with Gaussian uncertainties

To gain intuition for HMC and to provide a relation to well-studied linear inverse problems (e.g. Parker, 1994; Tarantola, 2005; Nolet, 2008), we continue the example from section 2.3, where we assumed Gaussian uncertainties and a linear forward model. Again, we note that this problem may be solved analytically, thus not being a case where Monte Carlo sampling is truly needed. Following Eq. (17), the resulting potential energy is

$$U(\mathbf{m}) = -\log \rho_m(\mathbf{m}|\mathbf{d}) = \frac{1}{2} \left(\mathbf{d}^{\text{obs}} - \mathbf{Gm}\right)^T \mathbf{C}_D^{-1} \left(\mathbf{d}^{\text{obs}} - \mathbf{Gm}\right) + \frac{1}{2} \mathbf{m}^T \mathbf{C}_M^{-1} \mathbf{m} + \text{const.}. \tag{26}$$

Hamilton's equations (A1) now take the specific form

$$\frac{d\mathbf{m}}{dt} = \nabla_p H = \nabla_p K = \mathbf{M}^{-1}\mathbf{p}, \tag{27}$$

$$\frac{d\mathbf{p}}{dt} = -\nabla_m H = -\nabla_m U = \mathbf{G}^T \mathbf{C}_D^{-1} \left(\mathbf{d}^{\text{obs}} - \mathbf{Gm}\right) - \mathbf{C}_M^{-1}\mathbf{m}. \tag{28}$$

5   Equations (27) and (28) constitute a coupled system of ordinary differential equations. Differentiating Eq. (27) with respect to time, and substituting $\frac{d\mathbf{p}}{dt}$ from Eq. (28), we obtain an independent, second-order differential equation for the model parameters $\mathbf{m}$,

$$\frac{d^2\mathbf{m}}{dt^2} + \mathbf{M}^{-1}\left(\mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{G} + \mathbf{C}_M^{-1}\right)\mathbf{m} = \mathbf{M}^{-1}\mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{d}^{\text{obs}}. \tag{29}$$

Denoting by $\mathbf{W}$ the square root of the matrix $\mathbf{M}^{-1}\left(\mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{G} + \mathbf{C}_M^{-1}\right)$, that is

10   $$\mathbf{W} = \sqrt{\mathbf{M}^{-1}\left(\mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{G} + \mathbf{C}_M^{-1}\right)}, \tag{30}$$

and by $\mathbf{f}$ the vector $\mathbf{M}^{-1}\mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{d}^{\text{obs}}$, we can write Eq. (29) in more condensed vector-matrix notation,

$$\frac{d^2\mathbf{m}}{dt^2} + \mathbf{W}^2\mathbf{m} = \mathbf{f}. \tag{31}$$

Using the matrix sine and cosine, the general solution of (31) is given by

$$\mathbf{m}(t) = \sin(\mathbf{W}t)\,\mathbf{a} + \cos(\mathbf{W}t)\,\mathbf{b} + \left(\mathbf{W}^2\right)^{-1}\mathbf{f}, \tag{32}$$

15   with constants $\mathbf{a}$ and $\mathbf{b}$ determined by the initial state of the trajectory. An analogous equation can be obtained for the momentum vector $\mathbf{p}$. Eq. (32) shows that Hamiltonian trajectories for a linear problem with Gaussian uncertainties are harmonic oscillations, the frequencies of which are controlled by the matrix $\mathbf{W}$.

In the context of a linear inverse problem, the matrix-valued frequency $\mathbf{W}$ can be interpreted in terms of the posterior covariance matrix, $\tilde{\mathbf{C}}_M = \left(\mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{G} + \mathbf{C}_M^{-1}\right)^{-1}$, corresponding to the Gaussian priors defined in Eqs. (12) and (13) (Tarantola, 20   2005). Thus, we find $\mathbf{W} = (\mathbf{M}\tilde{\mathbf{C}}_M)^{-1/2}$, meaning that comparatively small posterior covariances result in comparatively fast oscillations, and vice versa. The mass matrix $\mathbf{M}$ can serve as tuning parameter to adjust the speed with which the phase space is traversed, as we will discuss in section 3.6.

To make this more concrete, we reuse the example from section 2.4 with unit covariances $\mathbf{C}_M = \mathbf{I}$ and $\mathbf{C}_D = \mathbf{I}$. For illustration, we choose a unit mass matrix, $\mathbf{M} = \mathbf{I}$. Fig. 3 shows a random realisation of a Hamiltonian trajectory through the 25   ten-dimensional model space projected onto various two-dimensional subspaces. It is clearly visible that model parameters with small derivatives G_{ii} and correspondingly large posterior covariances traverse the model space comparatively slowly, and have a large oscillation amplitude.
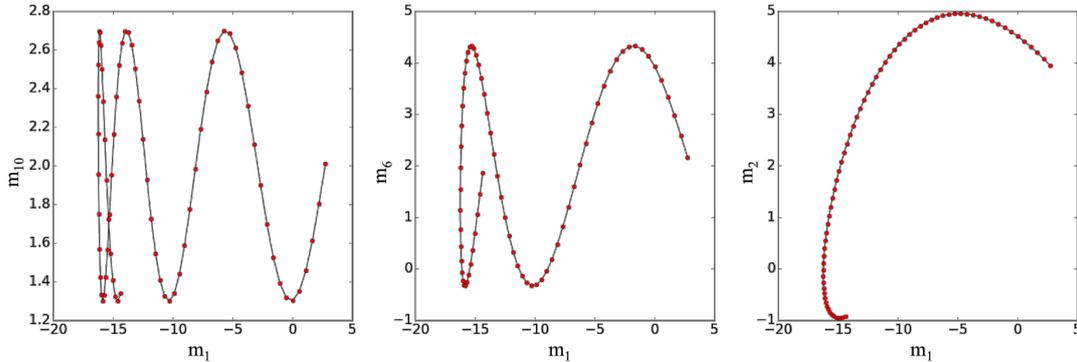
Figure 3. Realisation of a ten-dimensional Hamiltonian trajectory projected onto the two-dimensional planes (m_1, m_{10}), (m_1, m_6) and (m_1, m_2). Model parameters with small forward problem derivative G_{ii}, such as m_1, oscillate slowly with large amplitude, and vice versa.
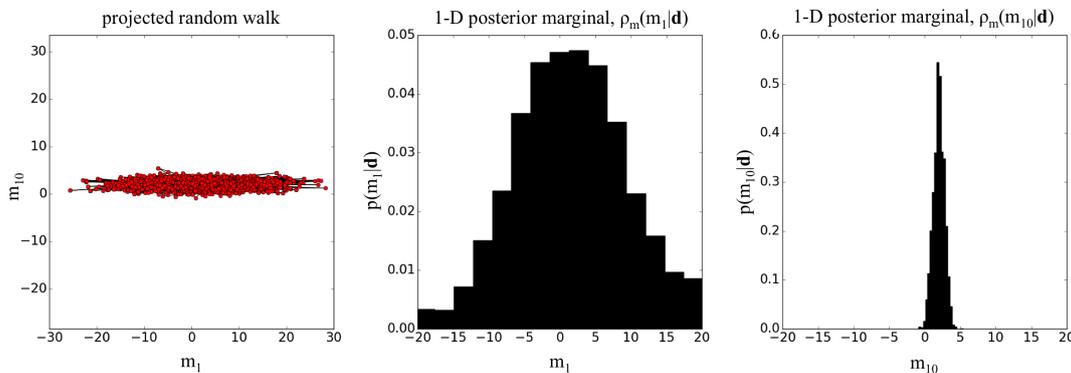

Figure 4. HMC applied to the example problem from section 2.3 and figure 2. The random walker's trajectory is shown in the left panel. Out of 10000 proposed models, 2681 were accepted, that is $\sim 27\,\%$. The samples are distributed very favourably so that a good approximation of the posterior, shown as 1-D marginals, can be obtained despite their small number.

An HMC run with 10000 test models is summarised in Fig. 4. Using a Hamiltonian trajectory length of 20 s, 2681 of the proposed models, $\sim 27\,\%$, are accepted. Their distribution is so favourable that the posterior can be approximated very well, despite the small number of samples. To put this result into perspective, we compare to the numbers from the introductory example in Fig. 2 where a more informative Gaussian prior with standard deviation $\sigma_M = 10$ led to an acceptance rate of $\sim 0.0014\,\%$. If we wanted to also have around 2700 accepted models from the EMH algorithm, we would need to propose more than 19 million test models.

3.6 Hamiltonian Monte Carlo in practice

In contrast to EMH, HMC involves tuning and subjective choices that are not a priori determined, and that may affect the performance of the algorithm. They include the numerical integrator of Hamilton's equations, the mass matrix $\mathbf{M}$, and the

12

length of the trajectories. We are thus presented with a potentially delicate tuning problem that is unlikely to have generally valid solutions. In practice, HMC unavoidably requires some trial and error (Neal, 2011). Below, we summarise some of the possible choices, not claiming general superiority of any of them.

3.6.1 Numerical integration

5 With the exception of few special cases, Hamilton's equations (A1) need to be solved numerically, using methods that maintain key properties of the Hamiltonian system needed for detailed balance, namely volume preservation and time reversability. Most frequently, the leapfrog method with time step Δt is used:

$$p_i(t + \Delta t/2) = p_i(t) - \frac{1}{2}\Delta t \left.\frac{\partial U}{\partial m_i}\right|_t, \tag{33}$$

$$m_i(t + \Delta t) = m_i(t) + \Delta t \left.\frac{\partial K}{\partial p_i}\right|_{t+\Delta t/2}, \tag{34}$$

10 $$p_i(t + \Delta t) = p_i(t + \Delta t/2) - \frac{1}{2}\Delta t \left.\frac{\partial U}{\partial m_i}\right|_{t+\Delta t}. \tag{35}$$

However, alternative (symplectic) integrators are well-documented in the literature (e.g. de Vogelaere, 1956; Ruth, 1983; Sanz-Serna and Calvo, 1994; Leimkuhler and Reich, 1994; Blanes et al., 2014). In our numerical examples we consistently use the leapfrog scheme, mostly for simplicity and ease of implementation.

3.6.2 The mass matrix

15 An obvious potential improvement by tuning is most apparent in Fig. 3. While m_1 traverses the energy level set roughly once during the Hamiltonian trajectory, m_{10} traverses it more than 6 times, thus visiting similar positions repeatedly without exploring new regions of model space. For linear forward problems and Gaussian priors, this deficiency may be avoided by choosing the mass matrix as the inverse posterior covariance

$$\mathbf{M} = \tilde{\mathbf{C}}_M^{-1} = \mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{G} + \mathbf{C}_M^{-1}. \tag{36}$$

20 Following Eqs. (30) and (32), the matrix-valued oscillation frequency $\mathbf{W}$ is then equal to the unit matrix $\mathbf{I}$. Therefore, all components oscillate through model space with an identical circular frequency of 1. For illustration, we again return to the toy example from section 2.4. Choosing the mass matrix according to Eq. (36) results in the trajectories shown in Fig. 5. In contrast to Fig. 3, where we chose $\mathbf{M} = \mathbf{I}$, all components of the model space vector now oscillate with equal frequency, thus traversing the phase space with equal speed.

25 While being attractive through its apparent simplicity, the choice of Eq. (36) deserves some additional remarks. (1) Already having to know the inverse posterior covariance from the outset clearly defeats the purpose of sampling. Nevertheless, this special case of linear forward modelling and Gaussian priors establishes an optimistic baseline that may be used for comparison, as for example in section 4.1. It furthermore serves as a useful means to gain intuition. (2) When the prior model covariance $\mathbf{C}_M^{-1}$ vanishes or when the prior distribution is uniform, Eq. (36) transforms to $\mathbf{M} = \tilde{\mathbf{C}}_M^{-1} = \mathbf{G}^T\mathbf{C}_D^{-1}\mathbf{G}$, which may not be
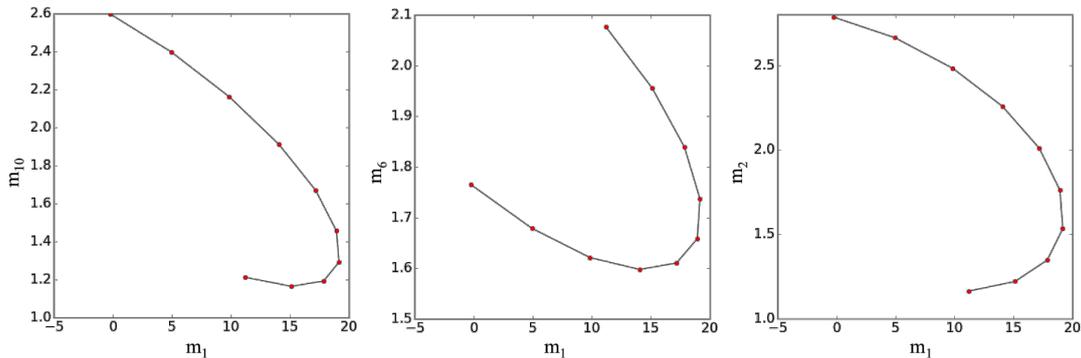
Figure 5. The same as Fig. 3, but with the mass matrix $\mathbf{M}$ defined according to Eq. (36). All components of the model space vector traverse the model space with identical frequency.

a proper covariance matrix with strictly non-zero eigenvalues. In this case, $\mathbf{M}$ may require regularisation, for instance, by adding a scaled unit matrix, $\varepsilon \mathbf{I}$, with sufficiently large $\varepsilon > 0$. (3) Most importantly, for nonlinear problems, Eq. (36) does not apply directly. However, when nonlinearity is sufficiently weak, the Hessian $\mathbf{H}$ of the potential energy U, evaluated at the prior maximum-likelihood model, can be a useful mass matrix. This variant, is referred to as Hessian Hamiltonian Monte Carlo

5 (H-HMC). Differentiating Eq. (26) shows that the Hessian of U is equal to the inverse posterior covariance $\tilde{\mathbf{C}}_M^{-1}$ of the linear Gaussian problem, thus establishing it as a special case of nonlinear H-HMC.

The possibility to tune HMC sufficiently well for arbitrary priors and potentially nonlinear forward equations is strongly problem-dependent, as we will also see in section 4 and further discuss in section 5.4. The possibility to approximate an optimal mass matrix for weakly nonlinear problems is the main reason for our focus on these cases. Any universal statements

10 on more general problems seem difficult to make.

3.6.3 Trajectory length

Fig. 3 also indicates that trajectories can be too short or too long. If too short, the next proposal is very close to the current position, leading to slow exploration of model space. In contrast, a trajectory that is too long, revisits points in model space that it has seen before, thereby wasting valuable computing time. Furthermore, the time needed to traverse a specific energy

15 level set depends on the energy itself. Low-energy level sets require less, and high-energy level sets require more time for a traversal. An additional complication arises in the presence of periodic trajectories, such as those shown for the linear problem in Fig. 3. Accidentally fixing the trajectory length to an integer multiple of the oscillation frequency will cause the random walker to always revisit the same positions, thus preventing convergence towards the posterior distribution.

At least for (nearly) linear problems with Gaussian priors, we can estimate suitable trajectory lengths. In the case of the optimal

20 mass matrix (36), we have $\mathbf{W} = \mathbf{I}$, and a trajectory length of π will achieve the maximum distance from the starting point of the trajectory. When the forward problem is weakly nonlinear or when a diagonal approximation of Eq. (36), a trajectory length

of approximately π may still serve as a useful guiding value.

To avoid locking in periodic trajectories, the trajectory length may be chosen randomly, as suggested by Mackenzie (1989) and Neal (2011). Alternatively, the length of the trajectory may be chosen adaptively, e.g., using the *No-U-Turn criterion* that terminates the integration when the trajectory begins to return towards its starting point (Gelman et al., 2014; Betancourt, 2017).

3.7 Variants of Hamiltonian Monte Carlo

Similar to other Monte Carlo methods, variants of HMC may be devised to meet application-specific requirements. With a comprehensive coverage of all variations being beyond the scope of this article, we limit ourselves to tempered, constrained and sloppy sampling, which will partly reappear in the examples section 4. A transdimensional version of HMC may be found in Biswas and Sen (2017).

3.7.1 Tempering

Tempered HMC promotes a more exploratory model space search by replacing the original canonical distribution $\rho_c(\mathbf{m}, \mathbf{p})$ by

$$\rho_c(\mathbf{m}, \mathbf{p}; T) = \rho_c^{T_0/T}(\mathbf{m}, \mathbf{p}), \tag{37}$$

where T is a positive number interpreted as temperature. The reference temperature $T_0 > 0$ is typically set to 1. Applying HMC to the tempered canonical distribution $\rho_c(\mathbf{m}, \mathbf{p}; T)$ provides samples of the tempered posterior

$$\rho_m(\mathbf{m}|\mathbf{d}; T) = \rho_m^{T_0/T}(\mathbf{m}|\mathbf{d}). \tag{38}$$

Choosing $T > T_0$ tends to flatten the posterior so that peaks of ρ_m become less pronounced relative to troughs. As a consequence, the random walker is more likely to escape from a model with a large posterior and to continue with the exploration of other model space regions. This may help to transition more easily between different modes of the posterior, to sample regions of lower likelihood more densly, and to accelerate convergence (Graham and Storkey, 2017). After sampling the tempered canonical distribution, the sampling density of the tempered posterior must be raised to the power of T/T_0 in order to retrieve the original posterior $\rho_m(\mathbf{m}|\mathbf{d})$ that we actually wish to sample. In section 4.1 we will provide an example of tempering in a linear tomography application.

3.7.2 Constrained sampling

Sometimes the model parameters we intend to investigate are subject to constraints, such as being positive or being bounded within a minimum and a maximum value. However, the standard version of the HMC algorithm assumes the variables $\mathbf{m}$, which represent the position, to be continuous and defined as real numbers from minus to plus infinity. In this section we show how it is possible to modify the standard algorithm to obtain a version where the model parameters $\mathbf{m}$ are bounded within lower

15

and upper constraints $l_i \leq m_i \leq u_i$. We follow here the strategy proposed in Neal (2011) and we refer the reader to that work for details. The core idea is to set the potential energy $U(\mathbf{m})$ to be infinite for values of m_i which do not satisfy the constraints, inducing a zero probability for those values. The latter can be obtained by modifying the potential energy by adding a term that goes to infinity outside the constraints and consequently changing the dynamics of the system.

5   To actually run the modified dynamics one can use the technique of "splitting" the Hamiltonian (Neal, 2011, and references therein): the Hamiltonian is expressed as a sum of terms $H(\mathbf{m}, \mathbf{p}) = H_1(\mathbf{m}, \mathbf{p}) + H_2(\mathbf{m}, \mathbf{p}) + H_3(\mathbf{m}, \mathbf{p}) + \ldots$ and then the dynamics (Eq. (A1)) are simulated for each term sequentially for a time Δt and this procedure is repeated until the desired total time of simulation is reached. As an illustration of this approach, consider the symmetric splitting:

$$H(\mathbf{m}, \mathbf{p}) = H_1(\mathbf{m}) + H_2(\mathbf{p}) + H_3(\mathbf{m}) = \frac{U(\mathbf{m})}{2} + K(\mathbf{p}) + \frac{U(\mathbf{m})}{2}. \tag{39}$$

10   If we apply the dynamics to each of the three terms separately for a time Δt, we retrieve the leapfrog method described in section 3.6.1 (Eqs. (33) to (35)). To constrain the model parameters with upper and lower bounds $l_i \leq m_i \leq u_i$, one possible approach is then to use the following splitting of the Hamiltonian

$$H(\mathbf{m}, \mathbf{p}) = H_1(\mathbf{m}) + H_2(\mathbf{m}, \mathbf{p}) + H_3(\mathbf{m}) = \frac{U_0(\mathbf{m})}{2} + \left(K(\mathbf{p}) + U_1(\mathbf{m}) \right) + \frac{U_0(\mathbf{m})}{2}, \tag{40}$$

where we introduce a new term for the potential energy $U_1(m_i)$ that goes to infinity if any of the components m_i are outside
15   the bounds and is equal to zero if all m_i are within the bounds. Applying the splitting method to the new Hamiltonian, we obtain an algorithm identical to the leapfrog method for H_1 and H_3 (i.e. steps (33) and (35)), while we have a modification for $H_2 = K(\mathbf{p}) + U_1(\mathbf{m})$. After running the first step (Eq. (33)) and calculating the new m_i by means of the second step (Eq. (34)) we check, one at the time, whether the new components of $\mathbf{m}$ are within the bounds. If that is the case, we have $U_1(m_i) = 0$ by definition. Therefore, the dynamics are not influenced by the new term U_1, and we can continue with the last step of the leapfrog
20   algorithm (Eq. (35)). Otherwise, if a component m_i is found to be outside the bounds, $U_1(m_i)$ plays a role and the dynamics change. The steep potential wall produced by $U_1(m_i)$ causes the trajectory to reflect with opposite momentum $p'_i = -p_i$ in order to respect the constraints. In practice, the following procedure is applied after the fist step (Eq. (33)):

1. For each i, set $p'_i = p_i(t + \Delta/2)$ and $m'_i = m_i(t) + \Delta t \left. \frac{\partial K}{\partial p'_i} \right|_{t + \Delta t/2}$ (Eq. (34))

2. For each i, while m_i is outside the bounds l_i and u_i

25   – if $m_i > u_i$ set $m'_i = u_i - (m'_i - u_i)$ and $p'_i = -p'_i$

     – if $m_i < l_i$ set $m'_i = l_i + (l_i - m'_i)$ and $p'_i = -p'_i$

3. Set $m_i(t + \Delta t) = m'_i$ and $p_i(t + \Delta t/2) = p'_i$

Finally, the third step of the leapfrog algorithm is applied (Eq. (35)), and the whole procedure is repeated until the total time of simulation is reached. This ensures that each of the model parameters is within bounds while running the dynamics.

3.7.3 Sloppy HMC

To limit the computational cost of HMC, sloppy variants may be devised to reduce the number of derivative evaluations during the numerical integration of Hamilton's equations (see section 3.6.1). The end-member case only uses one derivative of the forward problem at the starting point of a trajectory in order to approximate the derivative of the potential energy U along the trajectory. As is the case for any HMC tuning, also the efficiency of this approach must be assessed on a case-by-case basis. An example for sloppy HMC for seismic source inversion can be found in Fichtner and Simutė (2018).

4 Applications and numerical experiments

Following the more theoretical treatment of HMC, we continue with a series of tomographic/imaging applications with variable dimensionality and degree of nonlinearity. This is intended to illustrate some practical aspects related to the numerical setup and tuning of the algorithm.

4.1 Linear traveltime tomography

As first example, we consider a synthetic linear traveltime tomography based on ray theory, using Gaussian priors. Though linearised tomography ignores the dependence of the ray path on slowness perturbations relative to the background, it remains widely used (e.g. Fishwick et al., 2005; Ritsema et al., 2011; Koelemeijer et al., 2015). For simplicity, we assume Gaussian priors, noting that this may have limited meaning in real-data applications. This allows us to work within the framework established in section 3.5. As noted before, the linear Gaussian case has an analytical solution and therefore does not require Monte Carlo sampling per se. It is, nevertheless, ideal to demonstrate the critical role of the mass matrix $\mathbf{M}$.

Our numerical setup, summarised in Fig. 6, mimics a cross-hole tomography where sources with 1 m spacing emit waves recorded by receivers with 1 m spacing on the other side of the domain. The 2-D target model consists of a 1 m $\times$ 1 m chequerboard pattern, with slowness in each block being a model parameter. The total number of model parameters is $101 \times 101 = 10201$. We sample a tempered version of the posterior, as introduced in Eq. (37), with reference temperature $T_0 = 1$ and temperature $T = 450$.

We consider two choices for the mass matrix: (i) the inverse posterior covariance $\mathbf{M}_1 = \tilde{\mathbf{C}}_M^{-1} = \mathbf{G}^T \mathbf{C}_D^{-1} \mathbf{G} + \mathbf{C}_M^{-1}$, as suggested in equation (36), and (ii) the diagonal of $\tilde{\mathbf{C}}_M^{-1}$, that is, $\mathbf{M}_2 = \mathrm{diag}\, \tilde{\mathbf{C}}_M^{-1}$. For $\mathbf{M}_1 = \tilde{\mathbf{C}}_M^{-1}$ stable posterior means and covariances, shown in Fig. 7a, can be obtained with only 10^3 samples. In contrast, for $\mathbf{M}_2 = \mathrm{diag}\, \tilde{\mathbf{C}}_M^{-1}$ on the order of 10^6 samples are needed to achieve similar results, as illustrated in Figs. 7b-d.

While requiring the posterior covariance from the outset obviously defeats the purpose of sampling, the case of $\mathbf{M}_1 = \tilde{\mathbf{C}}_M^{-1}$ nevertheless establishes an optimistic baseline that may be used for comparison. The example shows that the choice of the mass matrix is not a detail, but an important factor that may change the number of required samples by orders of magnitude.
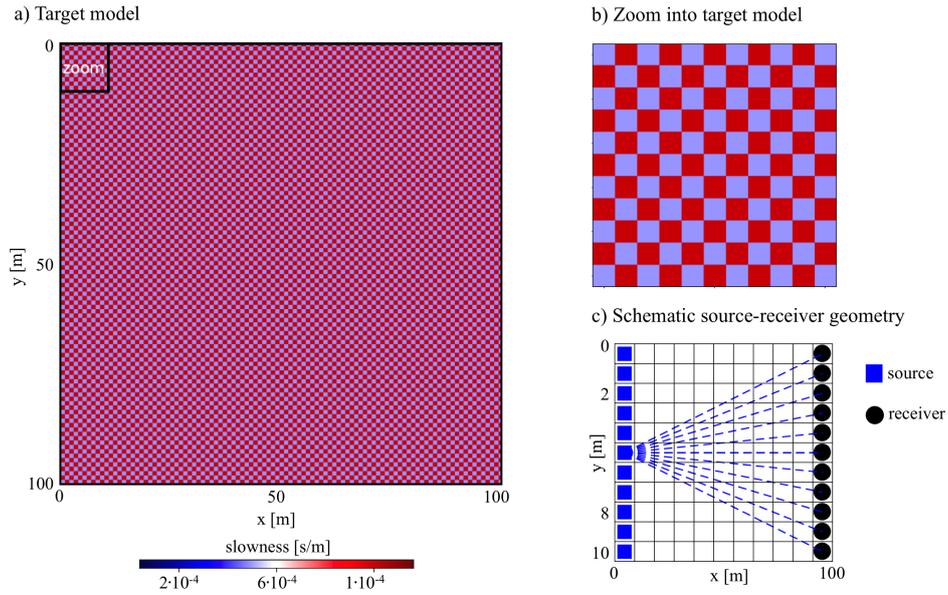
Figure 6. Setup of the synthetic linear traveltime tomography. a) The target model consists of a chequerboard pattern within a 2-D domain of 100 m × 100 m. b) Zoom into the target model. c) The source-receiver geometry mimics a cross-hole setup where sources with 1 m spacing along the left boundary emit waves recorded by receivers with 1 m spacing along the right boundary.
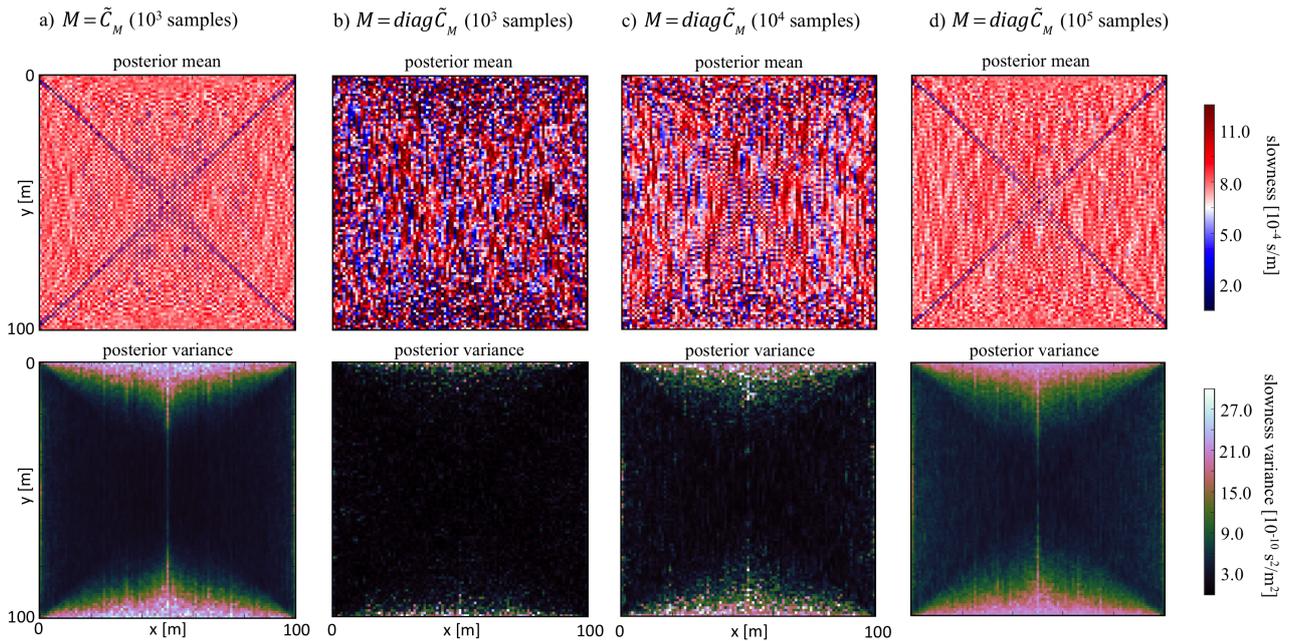

Figure 7. Posterior means and covariances for the linear tomography example for different choices of the mass matrix and the number of HMC samples. a) $\mathbf{M} = \tilde{\mathbf{C}}_M^{-1}$ with 10^3 samples. b)-d) $\mathbf{M} = \text{diag} \, \tilde{\mathbf{C}}_M^{-1}$ with 10^3, 10^4 and 10^5 samples.
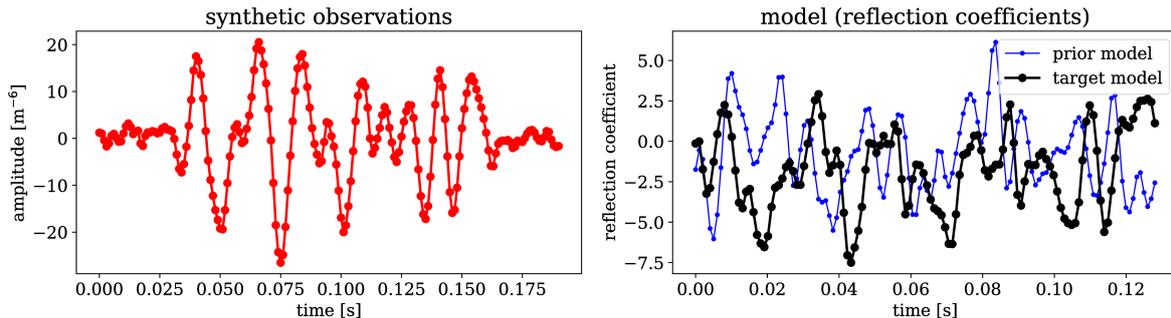
Figure 8. Synthetic observed data (left), target and prior model (right) for the seismic reflection numerical experiment.

4.2 Reflection seismology

We continue with a conceptual example where seismic reflection data are inverted for a 1-D subsurface model. This is intended to illustrate the generation of independent samples by HMC, compared to EMH. For simplicity, we consider only a single trace. The linear forward equations relate reflection coefficients $\mathbf{m}$ to a seismogram $\mathbf{d}$ via the matrix $\mathbf{G}$ that contains the source

5 wavelet and performs the convolution operation, according to the standard convolutional model of reflection seismology (e.g. Yilmaz, 2001). In this example we use 128 reflection coefficients and a seismogram consisting of 128 wave amplitude values. Fig. 8 shows the synthetic observed data with the addition of Gaussian noise (left) generated from the target model, and the prior reflectivity model (right). With a linear forward problem and Gaussian errors, this example falls into the framework discussed previously in section 3.5. As mass matrix we use an approximation of the inverse posterior covariance, $\mathbf{M} = \mathbf{G}^T \mathbf{C}_D^{-1} \mathbf{G}$.

10 To assess the efficiency of HMC relative to EMH, we compute the standard error σ of the Monte Carlo integral, given by Eq. (4) for the case of uncorrelated samples. When the samples are correlated, Eq. (4) must be corrected, by replacing the sample size N by the smaller effective sample size N_{eff}, which increases the error (Neal, 2011; Gelman et al., 2014). Following Neal (2011), we may approximate N_{eff} by

$$N_{\text{eff}} = \frac{N}{1 + 2 \sum_{n=1}^{\infty} a_n}, \tag{41}$$

15 where a_n is the auto-correlation of the sample chain at lag n. More sophisticated definitions of N_{eff} exist (e.g. Gelman et al., 2014), however, we comply with simplicity in this example.

Figs. 9a,b show correlation coefficients of 100 successive reflectivity models, starting from iteration 10000 to skip the burn-in phase. While correlation coefficients of the HMC samples change rapidly, groups of EMH samples tend to be strongly correlated. This suggests that EMH explores model space only slowly compared to HMC. This is also due to the fact that in

20 HMC all model parameters are updated in each iteration, while in EMH only subsets of parameters are updated in each iteration to maintain a relatively high level of acceptance. This last strategy is common practice, particularly when sophisticated forms of prior information are used (e.g. Zunino et al., 2015).
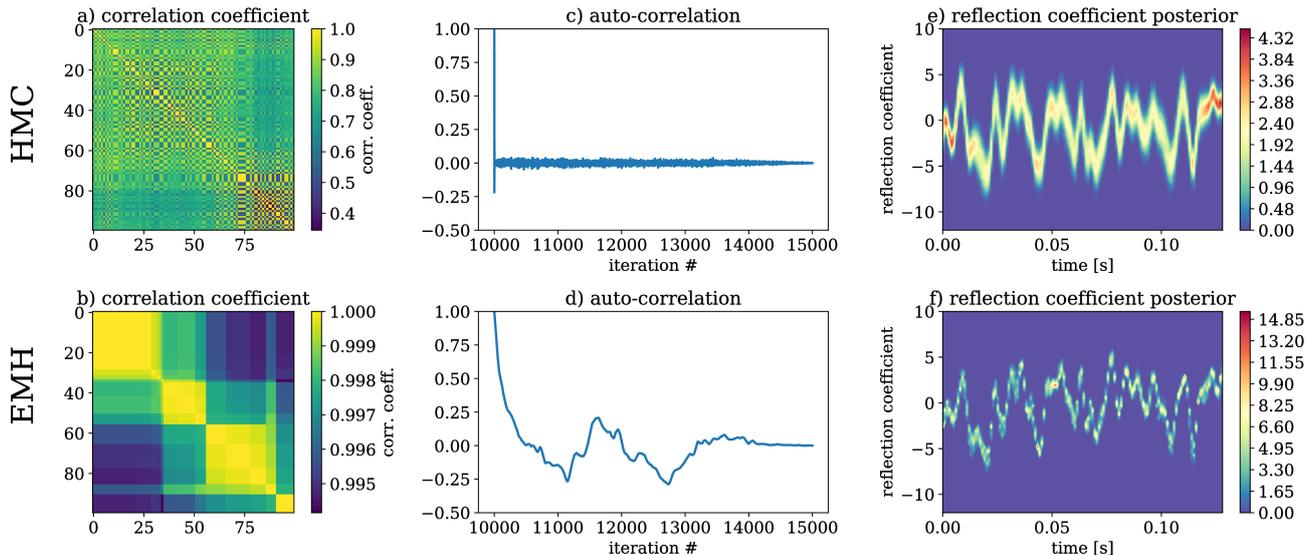
19

Figure 9. Comparision of sample independence for HMC (top row) and EMH (bottom row). a,b) Correlation coefficients of 100 successive reflectivity models, starting from iteration 10000. c,d) Auto-correlation a_n of 5000 successive samples. e,f) Posterior histograms of the reflection coefficients.

As another measure of sample independence, Figs. 9c,d show the auto-correlation a_n of the sample chain, which enters the effective sample size in Eq. (41). While the auto-correlation for HMC immediately drops close to 0, EMH requires around 500 samples to achieve the same level of independence. This is again an indication of the fact that HMC has produced more uncorrelated samples, and that the long-distance moves of HMC promote mixing of the chain. The improved model space 5 exploration of HMC, relative to EMH, also improves the representation of the posterior, which remains patchy for EMH, as shown in Figs. 9e,f. Finally, Fig. 10 depicts the effective sample size, N_{eff} for each of the 128 parameters of the problem calculated from iteration 10000 to 15000. The difference between HMC and EMH is striking: the percentage of independent samples for HMC is mostly close to 100 %, while it is less than 1 % for EMH. This means, again, that less iterations are needed for HMC compared to EMH to achieve the same reliability of posterior estimations.

10 **4.3 Non-linear traveltime tomography**

In this section, we present a nonlinear traveltime tomography that accounts for the deflection of ray paths in response to lateral heterogeneities (e.g. Nolet, 2008) an example of solution to the non-linear problem of traveltime tomography (Cerveny, 2001; Nolet, 2008). Nonlinear arrival time tomography has a long history of application from global to local scales (e.g. Zhang and Toksöz, 1998; Bijwaard and Spakman, 2000; Widiyantoro et al., 2000; Rawlinson et al., 2006). This example is an extension of 15 the linear case presented in section 4.1, with posterior distributions being non-Gaussian due to the nonlinear relation between
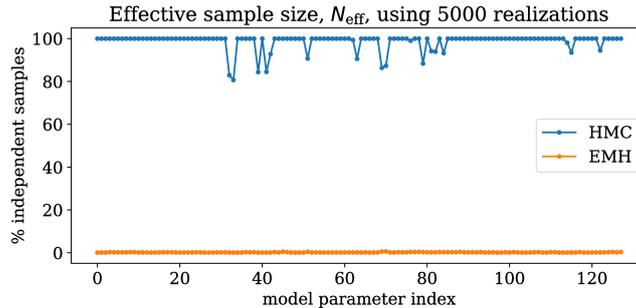
20

Figure 10. Effective sample size, N_{eff} for each model parameter in the HMC and EMH runs, considering 5000 iterations after burn-in. The HMC chain shows complete independence of realisations for most of the parameters, while the EMH run shows an average percentage of independent samples below 1 %. Much fewer iterations are thus needed for HMC algorithm compared to EMH to achieve the same error on the posterior expectation.

seismic velocity (the model parameters) and traveltime (the observable data). We consider a 2-D setup representing a vertical cross-section of the Earth. As shown in Fig. 11a, earthquakes with known hypocenters are located near the bottom of the model, and a set of receivers is aligned at the surface. The model is parameterized using 70×40 grid points, thus leading to a 2800-dimensional inverse problem in a 70 km $\times$ 40 km wide domain. The forward relationship is provided by the eikonal

5   equation (e.g. Cerveny, 2001), describing the traveltime τ of first arrivals as a function of velocity v at position (x, y). In 2-D it reads

$$|\nabla \tau|^2 = \left(\frac{\partial \tau}{\partial x}\right)^2 + \left(\frac{\partial \tau}{\partial y}\right)^2 = \frac{1}{v^2(x, y)}\,. \tag{42}$$

The misfit functional, measuring the degree of fit between calculated, and observed data and corresponding to the potential energy $U(\mathbf{m})$, with $\mathbf{m} \equiv \mathbf{v}$ in the HMC framework, is given by

10    $$S(\mathbf{v}) = \frac{1}{2} \sum_{i=1}^{N_r} \frac{\left(\tau_i^{\text{calc}}(\mathbf{v}) - \tau_i^{\text{obs}}\right)^2}{\sigma_i^2}\,, \tag{43}$$

where i denotes the receiver index ranging from 1 to N_r. The synthetic traveltime data were calculated from the target model in Fig. 11a, with the addition of uncorrelated Gaussian noise, thus producing 780 (26 sources and 30 receivers) traveltime measurements.

The traditional approach to solve the inverse problem, that is, to determine the velocity structure from the traveltime data,

15   is to linearize the forward relationship around a given model $\mathbf{v}_0$ and then to compute all individual rays between source-receiver pairs. From this, a matrix $\mathbf{G}$ relating a velocity model and the traveltime vector $\boldsymbol{\tau}$ can be constructed: $\boldsymbol{\tau} = \boldsymbol{\tau}^{\text{calc}}(\mathbf{v}_0) + \mathbf{G}(\mathbf{v}_0)(\mathbf{v}_n - \mathbf{v}_0)$. Subsequently, the gradient of the misfit function can be computed, and linear inversion methods can be used iteratively to solve the inverse problem. However, computing the gradient of the misfit functional in the abovementioned way may be detrimental for the HMC method, as it requires a large amount of computations and memory to construct the Fréchét

21

derivative $\mathbf{G}$.

We thus rely on a different approach based on the adjoint method and compute directly the gradient of the misfit as proposed in Leung and Qian (2006) and Taillandier et al. (2009). In the adjoint formulation the gradient of the misfit is calculated directly using numerical methods, which roughly have the computational cost of two forward simulations and require a limited

5 amount of memory. This provides a significant improvement in the computation of the gradient of the misfit function, making the HMC implementation significantly more efficient. It furthermore improves accuracy by avoiding the linearization of the forward problem, so that "fat-rays" are automatically taken into account.

To compute the traveltime given a velocity model we use the code FDTIMES from Podvin and Lecomte (1991), while for the adjoint calculations we use an in-house code based on the fast sweeping method (Zhao, 2005; Leung and Qian, 2006;

10 Taillandier et al., 2009; Bretaudeau et al., 2014) . The HMC has been implemented and parallelized using the Julia language (Bezanson et al., 2017).

Regarding the setup of the HMC algorithm, we use a mass matrix equivalent to an estimated inverse of the covariance matrix on the model parameters $\mathbf{M} = \mathbf{C}_M^{-1}$. Moreover, we manually tuned the integration time step Δt and the trajectory length of the leapfrog algorithm to achieve an acceptance rate of ~ 65 %. Since velocity can only be positive, we employed the constrained

15 HMC algorithm described in section 3.7.2. In total, we produced 10^6 HMC samples, of which we discarded the first 150000 as burn-in. To reduce storage requirements, we saved every 25th sample to produce an final collection of of 34000 posterior velocity models.

Fig. 11 shows a selection of results, including the model with the highest probability (Fig. 11b), a randomly chosen model (Fig. 11 c), and the mean model (Fig. 11 d). The posterior ensemble allows us to compute model statistics and to estimate

20 uncertainties. One may, for instance, calculate the histogram of plausible velocity values at given locations, which approximates the posterior marginal pdf of one velocity value. This is shown for locations, $(x, y) = (10, 26)$ km and $(x, y) = (50, 20)$ km in Figs. 11d,f. Since the problem is nonlinear, the distribution of velocity is not Gaussian and, particularly in the first case, appears multimodal with two nearly equiprobable peaks. Moreover, the spread of the histograms provides an estimate of uncertainty on those parameters. This example illustrates that the combination of two powerful techniques, the adjoint method and HMC,

25 may provide comprehensive solutions to high-dimensional nonlinear tomographic problems.

5 Discussion

In the following we discuss philosophical and technical issues of HMC in the context of tomographic inversions.
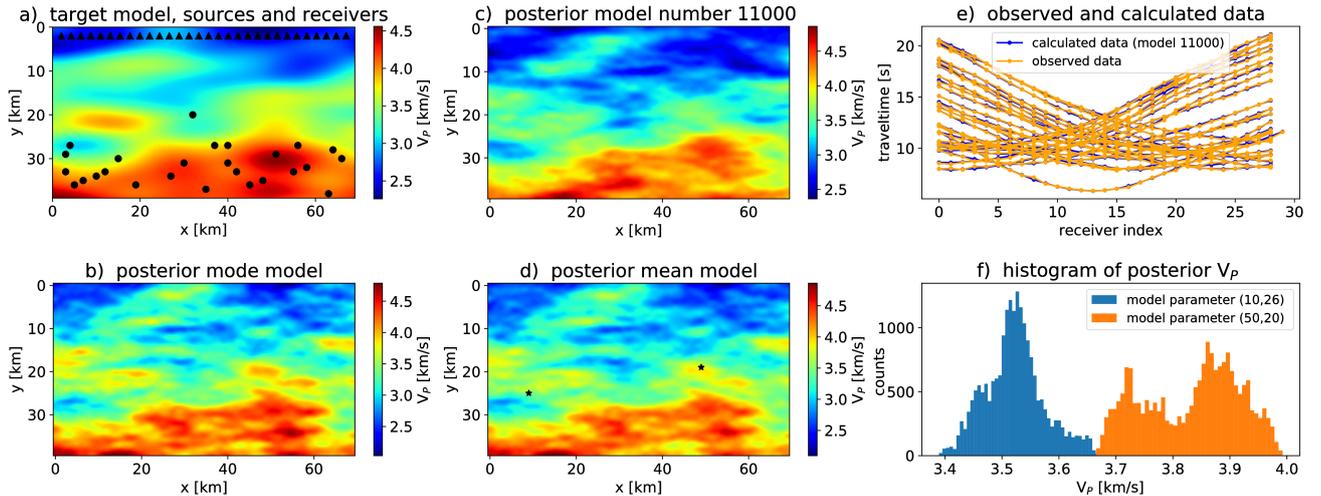
Figure 11. Setup and results of nonlinear tomography using HMC. a) The target velocity model used to compute, with the addition of Gaussian noise, the synthetic observations depicted in e). Dots represent source positions, and triangles the positions of receivers where the traveltime is recorded. b),c),d) Selection of results from the posterior ensemble: b) is the highest likelihood model, c) a randomly selected model (number 11000) and d) the mean model. e) The synthetic observations (orange) and the traveltimes calculated for model number 11000. f) Posterior histogram of velocity at locations $(x, y) = (10, 26)$ km and $(x, y) = (50, 20)$ km, whose position is indicated in plot d) by stars.

5.1 HMC for geophysical inversion

While HMC was developed more than 3 decades ago (Duane et al., 1987) and applied in various domains for at least 2 decades (Neal, 1996), its potential for the solution of geophysical inverse problems is only starting to be recognised (Muir and Tkalčić, 2015; Biswas and Sen, 2017; Sen and Biswas, 2017). Though the precise reasons for this delayed migration of HMC into geophysics are difficult to quantify, the fact that most existing HMC literature is highly domain-specific is likely to play a role.

One of the main motivations of this work was therefore a description of HMC within the context, and using the language, of geophysical inversion, and of tomography in particular. As discussed below, universally valid statements concerning the performance of any algorithm are impossible to make. Neverthelss, many geophysical problems seem to fall inside the range of efficiency of HMC because their nonlinearities are weak, and their derivatives are easy to compute.

5.2 Nonlinear traveltime tomography

The second focus of this work is on the applications, especially the nonlinear traveltime tomography of section 4.3. Most importantly, it illustrates that a nonlinear tomographic problem with nearly 3000 free parameters can be solved fully probabilistically by combining adjoint techniques and HMC.

23

We emphasise that the inversion was run on a standard laptop computer (2.5 days on a 6-core Intel Xeon), thus not making use of any supercomputing resources. Though extrapolations are inherently difficult, this suggests that the probabilistic solution of 3-D nonlinear tomographic problems is within reach.

5.3 Potentials: Sampling high-dimensional model spaces

5 The potential of HMC lies in its ability to efficiently sample high-dimensional model spaces, as illustrated for instance in the examples of sections 4.1 and 4.3. Fundamentally, this ability rests on the exploitation of derivative information that other Monte Carlo methods – such as EMH, genetic algorithms, or the Neighbourhood Algorithm – do not consider.

Though the use of derivatives is also the main limitation of HMC, many geophysical inverse problems with easily computable derivatives are likely to fall within its range of applicability. These include the large classes of tomographic and potential 10 field inversions. For emerging applications where the forward problem is solved purely numerically, adjoint techniques may be used for the efficient computation of derivatives, as proposed in seismology (e.g. Tromp et al., 2005; Fichtner et al., 2006; Plessix, 2006), geodynamics (e.g. Bunge et al., 2003; Liu and Gurnis, 2008; Colli et al., 2018), geomagnetism (Li et al., 2014), helioseismology (Hanasoge et al., 2011), and other fields. Furthermore, the extension to second-order adjoints (Santosa and Symes, 1988; Fichtner and Trampert, 2011) offers the possibility to use an approximation of the Hessian as nearly optimal 15 mass matrix.

5.4 Limitations: No free lunch

While HMC is widely recognised as an efficient alternative to other – and at least in geophysics more established – sampling methods, it is important to note that efficiency is unavoidably context-dependent. As formalised by various No-Free-Lunch Theorems (e.g. Wolpert and Macready, 1997; Mosegaard, 2012), efficiency is not a universal property of a method. Instead, it 20 arises from the combination of the method with prior knowledge about the nature of a specific problem.

In the case of HMC, the No-Free-Lunch Theorems act on two levels. First, concerning the method as a whole, HMC can only be efficient when the derivative of the forward problem can be computed quickly and with sufficient accuracy. Otherwise, the solution of Hamilton's equations will be too time-consuming. Second, within HMC, there is no universally optimal tuning, which includes the choice of the mass matrix, the numerical integration scheme, and the length of the trajectories. Again, all 25 tuning is problem-specific.

Consequently, it is inherently impossible to make general statements concerning the efficiency of HMC, however quantified. As other inverse methods, HMC requires some trial and error, for both tuning and comparison with other methods. In this context we note that HMC is not the only method designed to limit the curse of dimensionality at least to some extent. Alternatives include the Neighborhood Algorithm (Sambridge, 1999a, b), parallel tempering (e.g. Marinari and Parisi, 1992; Geyer and 30 Thompson, 1995; Sambridge, 2014), and trans-dimensional sampling (e.g. Green, 1995; Sambridge et al., 2006; Bodin and Sambridge, 2009; Sambridge et al., 2013).

Since universally valid statements on the efficiency of HMC are difficult to make, we limited this work to forward problems that are nearly linear. In these cases, an optimal mass matrix can be approximated using the Jacobian $\mathbf{G}$ of the forward problem at some point in model space. Again, as a consequence of the No-Free-Lunch Theorems, the term nearly linear is not sharply defined, but application specific.

6 Conclusions

We presented an introduction to HMC in the context of tomographic inverse problems. The relation to well-studied linear inverse problems with Gaussian uncertainties provides valuable insight into the mechanics of HMC, suitable choices of tuning parameters, and properties of the posterior ensemble. The advantage of HMC is its ability to make long-distance moves in model space to produce uncorrelated samples, while maintaining a high acceptance rate. These benefits are balanced by the need to compute derivatives, so that the overall efficiency of HMC is unavoidably problem-specific.

Examples of linear and nonlinear traveltime tomography illustrate the applicability of HMC to large classes of geophysical inverse problems. Combining HMC and adjoint techniques, the 2-D traveltime inversion provides the posterior pdf of nearly 3000 model parameters, running on a single processor without resorting to supercomputing resources.

Appendix A: Properties of Hamiltonian dynamics

In the following paragraphs, we provide proofs of three properties of Hamiltonian dynamics that are essential for HMC: time reversability, the preservation of a phase space partition, and volume preservation. More complete treatments of Hamiltonian systems can be found in the classic textbooks of Landau and Lifshitz (1976) or Symon (1971).

For completeness, we repeat Hamilton's equations

$$\frac{dm_i}{dt} = \frac{\partial H}{\partial p_i}, \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial m_i}. \tag{A1}$$

A key property of Hamiltonian systems with a time-independent Hamiltonian H is *time reversability*. Indeed, replacing t by $-t$ transforms the generalised momenta according to

$$p_i \rightarrow \sum_j M_{ij} \frac{dm_j}{d(-t)} = -p_i. \tag{A2}$$

For Hamilton's equations (A1) we then find

$$\frac{dm_i}{dt} = \frac{\partial H}{\partial p_i} \rightarrow \frac{dm_i}{d(-t)} = \frac{\partial H}{\partial(-p_i)} \quad \Rightarrow \quad \frac{dm_i}{dt} = \frac{\partial H}{\partial p_i}, \tag{A3}$$

and

$$\frac{dp_i}{dt} = -\frac{\partial H}{\partial m_i} \rightarrow \frac{d(-p_i)}{d(-t)} = -\frac{\partial H}{\partial m_i} \quad \Rightarrow \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial m_i}. \tag{A4}$$

25

Thus, Hamilton's equations are invariant under time reversal, provided that H has no explicit time dependence. Time reversability means that we cannot infer the direction of the time arrow from observing the trajectory of a particle.

An important consequence of time reversability is the preserved partitioning of the phase space under movement along a Hamiltonian trajectory: When the $\mathbb{X}^i$ partition that phase space $\mathbb{X}$, meaning that they completely fill $\mathbb{X}$ without overlapping, then the images $\mathbb{X}_\tau^i$ partition the phase space as well. To see this, we first assume that the image space has holes, so that there is an element $(\mathbf{m}_\tau, \mathbf{p}_\tau)$ that is not an image. Using time reversal, $(\mathbf{m}_\tau, \mathbf{p}_\tau)$ translates to some $(\mathbf{m}, \mathbf{p})$ of the phase space. Then, again by time reversal, $(\mathbf{m}, \mathbf{p})$ maps to $(\mathbf{m}_\tau, \mathbf{p}_\tau)$, thus contradicting the assumption that it is not an image. Conversely, assume that images $\mathbb{X}_\tau^i$ and $\mathbb{X}_\tau^j$ overlap, meaning that there is an $(\mathbf{m}_\tau, \mathbf{p}_\tau)$ located within both $\mathbb{X}_\tau^i$ and $\mathbb{X}_\tau^j$. It follows that there must have been distinct $(\mathbf{m}^i, \mathbf{p}^i) \in \mathbb{X}^i$ and $(\mathbf{m}^j, \mathbf{p}^j) \in \mathbb{X}^j$ mapping to the same point. This, however, contradicts time reversability which maps $(\mathbf{m}_\tau, \mathbf{p}_\tau)$ back to one unique $(\mathbf{m}, \mathbf{p})$.

In addition to being time-reversible, motion under Hamiltonian mechanics preserves volume in phase space. To see this, we consider a group of particles with position vectors $\mathbf{x} = (\mathbf{m}, \mathbf{p})$ within a volume X in phase space $\mathbb{X}$. As illustrated in figure A1, a small movement of the particles induces a small volume change given by

$$\mathbf{n} \cdot d\mathbf{x} \, dA = \mathbf{n} \cdot \frac{d\mathbf{x}}{dt} \, dt \, dA \,, \tag{A5}$$

where dA is a small surface element with normal $\mathbf{n}$. The total volume change per time, integrated over the surface ∂X of the volume X is then

$$\frac{dX}{dt} = \int_{\partial X} \mathbf{n} \cdot \frac{d\mathbf{x}}{dt} \, dA = \int_X \nabla \cdot \frac{d\mathbf{x}}{dt} \, d\mathbf{x} \,, \tag{A6}$$

where we applied Gauss' theorem to transform the surface into a volume integral. Using Hamilton's equations, we now find

$$\frac{dX}{dt} = \int_X \nabla \cdot \frac{d\mathbf{x}}{dt} \, d\mathbf{x} = \int_X \sum_{i=1}^n \left(\frac{\partial}{\partial m_i} \frac{dm_i}{dt} + \frac{\partial}{\partial p_i} \frac{dp_i}{dt} \right) d\mathbf{x} = \int_X \sum_{i=1}^n \left(\frac{\partial}{\partial m_i} \frac{\partial H}{\partial p_i} - \frac{\partial}{\partial p_i} \frac{\partial H}{\partial m_i} \right) d\mathbf{x} = 0 \,. \tag{A7}$$

This proves that the phase space volume is preserved under Hamiltonian motion. In other words, the volume occupied by the group of particles remains constant.
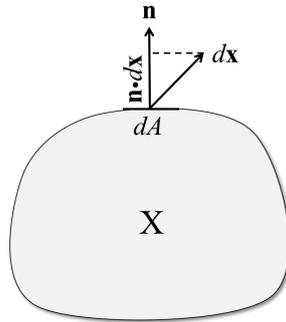
Figure A1. Volume preservation in phase space under Hamiltonian mechanics. A volume X in phase space. On the boundary, a particle is displaced by $d\mathbf{x}$, which induces a small volume change $\mathbf{n} \cdot d\mathbf{x}\, dA$, where $\mathbf{n}$ is the surface normal, and dA is a small surface element.

References

Aki, K., Christoffersson, A., and Husebye, E. S.: Determination of three-dimensional seismic structure of the lithosphere, J. Geophys. Res., 81, 277–296, 1976.

Backus, G. E. and Gilbert, F.: The resolving power of gross Earth data, Geophys. J. Roy. Astr. Soc., 16, 169–205, 1968.

5 Backus, G. E. and Gilbert, F.: Uniqueness in the inversion of inaccurate gross Earth data, Phil. Trans. R. Soc. London, A, 266, 123–192, 1970.

Bayes, T. and Price, R.: An Essay towards solving a Problem in the Doctrine of Chance, Phil. Trans. Roy. Soc. London, 53, 370–418, 1763.

Betancourt, M.: A conceptual introduction to Hamiltonian Monte Carlo, arXiv:1701.02434 [stat.ME], 2017.

Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B.: Julia: A fresh approach to numerical computing, SIAM Review, 59, 65–98, 2017.

10 Bijwaard, H. and Spakman, W.: Non-linear global P-wave tomography by iterated linearized inversion, J. Geophys. Res., 141, 71–82, 2000.

Bishop, C. M.: Pattern recognition and machine learning, Springer, New York, 2006.

Biswas, R. and Sen, M.: 2D full-waveform inversion and uncertainty estimation using the reversible jump Hamiltonian Monte Carlo, SEG Expanded Abstracts, pp. 1280–1285, 2017.

Blakely, R.: Potential theory in gravity and magnetic applications, Cambridge University Press, Cambridge, U.K., 2002.

15 Blanes, S., Casas, F., and Sanz-Serna, J. M.: Numerical integrators for the Hybrid Monte Carlo method, SIAM J. Sci. Comp., 36, A1556–A1580, 2014.

Bodin, T. and Sambridge, M.: Seismic tomography with the reversible jump algorithm, Geophys. J. Int., 178, 1411–1436, 2009.

Bretaudeau, F., Brossier, R., Virieux, J., and Métivier, L.: First-arrival delayed tomography using 1st and 2nd order adjoint-state method, in: SEG Technical Program Expanded Abstracts 2014, pp. 4757–4762, Society of Exploration Geophysicists, 2014.

20 Bunge, H.-P., Hagelberg, C. R., and Travis, B. J.: Mantle circulation models with variational data assimilation: Inferring past mantle flow and structure from plate motion histories and seismic tomography, Geophys. J. Int., 152, 280–301, 2003.

Cerveny, V.: Seismic ray theory, Cambridge University Press, 2001.

Chen, M.-H., Shao, Q.-M., and Ibrahim, J. G.: Monte Carlo methods in Bayesian computation, Springer Science & Business Media, 2012.

Colli, L., Ghelichkhan, S., Bunge, H.-P., and Oeser, J.: Retrodictions of Mid-Paleogene mantle flow and dynamic topography in the Atlantic region from compressible high-resolution adjoint mantle convection models: Sensitivity to deep mantle viscosity and tomographic input model, Gondwana Research, 53, 252–272, 2018.

de Vogelaere, R.: Methods of integration which preserve the contact transformation property of Hamiltonian equations, Tech. Rep. Report No. 4, Department of Mathematics, University of Notre Dame, Notre Dame, Ind., USA, 1956.

Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D.: Hybrid Monte Carlo, Phys. Lett. B, 195, 216–222, 1987.

Dubbledam, D., Calero, S., Ellis, D. E., and Snurr, R. Q.: RASPA: molecular simulation software for adsorption and diffusion in flexible nanoporous materials, Mol. Sim., 42, 81–101, 2016.

Dziewoński, A. M., Chou, T.-A., and Woodhouse, J. H.: Determination of earthquake source parameters from waveform data for studies of global and regional seismicity, J. Geophys. Res., 10, 2825–2852, 1981.

Ekström, G., Nettles, M., and Dziewonski, A. M.: The global CMT project 2004-2010: centroid moment tensors for 13,017 earthquakes, Phys. Earth Planet. Inter., 200-201, 1–9, 2012.

Elhatisari, S., Lee, D., Rupak, G., Epelbaum, E., Krebs, H., Lahde, T. A., Luu, T., and ner, U.-G. M.: Ab initio alpha-alpha scattering, Nature, 258, 111–113, 2015.

Fichtner, A. and Simutė, S.: Hamiltonian Monte Carlo inversion of seismic sources in complex media, J. Geophys. Res., p. under review, 2018.

Fichtner, A. and Trampert, J.: Hessian kernels of seismic data functionals based upon adjoint techniques, Geophys. J. Int., 185, 775–798, 2011.

Fichtner, A., Bunge, H.-P., and Igel, H.: The adjoint method in seismology - I. Theory, Phys. Earth Planet. Inter., 157, 86–104, 2006.

Fishwick, S., Kennett, B. L. N., and Reading, A. M.: Contrasts in lithospheric structure within the Australian Craton, Earth Planet. Sci. Lett., 231, 163–176, 2005.

Gallagher, K., Sambridge, M. S., and Drijkoningen, G. G.: Genetic algorithms: an evolution of Monte Carlo methods for strongly non-linear geophysical optimization problems, Geophys. Res. Lett., 18, 2177–2180, 1991.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B.: Bayesian data analysis, vol. 2, CRC press Boca Raton, FL, 2014.

Geyer, C. J.: Markov Chain Monte Carlo maximum likelihood, in: Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface, pp. 156–163, 1991.

Geyer, C. J. and Thompson, E. A.: Annealing Markov Chain Monte Carlo with applications to ancestral inference, J. Am. Stat. Assoc., 90, 909–920, 1995.

Graham, M. M. and Storkey, A. J.: Continuously tempered Hamiltonian Monte Carlo, arXiv:1704.03338 [stat.CO], 2017.

Green, P. J.: Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination, Biometrika, 82, 711–732, 1995.

Gutmann, M. U. and Hyvärinen, A.: Noise-contrastive estimation of unnormalized statistical models with applications to natural image statistics, J. Mach. Learning. Res., 13, 307–361, 2012.

Hanasoge, S. M., Birch, A., Gizon, L., and Tromp, J.: The adjoint method applied to time-distance helioseismology, Astrophys. J., 738, doi:10.1088/0004–637X/738/1/100, 2011.

Hastings, W. K.: Monte Carlo sampling methods using Markov Chains and their applications, Biometrika, 57, 97–109, 1970.

Holland, J. H.: Adaptation in natural and artificial systems, The University of Michigan Press, Ann Arbor, 1975.

Honkela, A., Peltonen, J., Topa, H., Charapitsa, I., Matarese, F., Grote, K., Stunnenberg, H. G., Reid, G., Lawrence, N. D., and Rattray, M.: Genome-wide modeling of transcription kinetics reveals patterns of RNA production delays, Proc. Nat. Acad. Sci., 112, 13 115–13 120, 2015.

Iyer, H. M. and Hirahara, K.: Seiemic tomography, Theory and Practice, Chapman & Hall, 1993.

5   Jackson, A., Jonkers, A. R. T., and Walker, M. R.: Four centuries of geomagnetic secular variation from historical records, Phil. Trans. R. Soc. Lond. A, 358, 957–990, 2000.

Jaynes, E. T.: Probability Theory - The Logic of Science, Cambridge University Press, Cambridge, UK, 2003.

Keilis-Borok, V. J. and Yanovskaya, T. B.: Inverse problems of seismology (structural review), Geophys. J. Roy. Astr. Soc., 13, 223–234, 1967.

10   Koelemeijer, P., Ritsema, J., Deuss, A., and van Heijst, H.-J.: SP12RTS: a degree-12 model of shear- and compressional-wave velocity for Earth's mantle, Geophys. J. Int., pp. 1024–1039, 2015.

Landau, L. D. and Lifshitz, E. M.: Course of Theoretical Physics, Volume 1, Mechanics, 3rd edition, Elsevier Butterworth Heinemann, Amsterdam, 1976.

Lee, E.-J., Chen, P., Jordan, T. H., Maechling, P. B., Denolle, M., and Beroza, G. C.: Full-3D tomography (F3DT) for crustal struc-
15   ture in Southern California based on the scattering-integral (SI) and the adjoint-wavefield (AW) methods, J. Geophys. Res., 119, doi: 10.1002/2014JB011 346, 2014.

Leimkuhler, B. and Reich, S.: Simulating Hamiltonian systems, Cambirdge University Press, Cambridge, UK, 1994.

Leung, S. and Qian, J.: An adjoint state method for three-dimensional transmission traveltime tomography using first-arrivals, Communications in Mathematical Sciences, 4, 249–266, 2006.

20   Li, K., Jackson, A., and Livermore, P. W.: Variational data assimilation for a forced, inertia-free magnetohydrodynamic dynamo model, Geophys. J. Int., 199, 1662–1676, 2014.

Li, Y. and Oldenburg, D.: 3-D inversion of gravity data, Geophysics, 63, 109–119, 1998.

Liu, L. and Gurnis, M.: Simultaneous inversion of mantle properties and initial conditions using an adjoint of mantle convection, J. Geophys. Res., 113, doi:10.1029/2007JB005 594, 2008.

25   MacKay, D. J.: Information theory, inference and learning algorithms, Cambridge university press, 2003.

Mackenzie, P. B.: An improved hybrid Monte Carlo method, Phys. Lett. B., 226, 369–371, 1989.

Mallet, R.: Account of Experiments Made at Holyhead (North Wales) to Ascertain the Transit-Velocity of Waves, Analogous to Earthquake Waves, through the Local Rock Formations, Phil. Trans. R. Soc. London, 151, 655–679, 1861.

Marinari, E. and Parisi, G.: Simulated tempering: a new Monte Carlo scheme, Europhys. Lett., 19, 451–458, 1992.

30   Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E.: Equations of state calculations by fast computing machines, J. Chem. Phys., 21, 1087–1092, 1953.

Mosegaard, K.: Limits to nonlinear inversion, pp. 11–21, Springer, Berlin, Heidelberg, 2012.

Mosegaard, K. and Tarantola, A.: Monte Carlo sampling of solutions to inverse problems, J. Geophys. Res., 100, 12 431–12 447, 1995.

Muir, J. B. and Tkalčić, H.: Probabilistic joint inversion of lowermost mantle P-wave velocities and core mantle boundary topography using
35   differential travel times and hierarchical Hamiltonian Monte-Carlo sampling, in: AGU 2015 Fall meeting, pp. S14A–03, AGU, 2015.

Neal, R. M.: Bayesian learning for neural networks, Springer, New York, 1996.

Neal, R. M.: MCMC using Hamiltonian dynamics, in: Handbook of Markov Chain Monte Carlo, p. Chapter 5, 2011.

Nolet, G.: A breviary of seismic tomography, Cambridge University Press, Cambridge, UK, 2008.

Oldham, R. D.: The constitution of the interior of the Earth as revealed by earthquakes, Quart. J. Geol. Soc. London, 62, 456–475, 1906.

Parker, R. L.: Geophysical Inverse Theory, Princeton University Press, 1994.

Peskun, P. H.: Optimum monte-carlo sampling using markov chains, Biometrika, 60, 607–612, 1973.

Plessix, R.-E.: A review of the adjoint-state method for computing the gradient of a functional with geophysical applications, Geophys. J. Int., 167, 495–503, 2006.

Podvin, P. and Lecomte, I.: Finite difference computation of traveltimes in very contrasted velocity models: a massively parallel approach and its associated tools, Geophysical Journal International, 105, 271–284, https://doi.org/10.1111/j.1365-246X.1991.tb03461.x, 1991.

Press, F.: Earth models obtained by Monte-Carlo inversion, J. Geophys. Res., 73, 5223–5234, 1968.

Rawlinson, N., Reading, A. M., and Kennett, B. L. N.: Lithospheric structure of Tasmania from a novel form of teleseismic tomography, J. Geophys. Res., 111, doi: 10.1029/2005JB003 803, 2006.

Ritsema, J., Deuss, A., van Heijst, H. J., and Woodhouse, J. H.: S40RTS: a degree-40 shear-velocity model for the mantle from new Rayleigh wave dispersion, teleseismic traveltime and normal-mode splitting function measurements, Geophys. J. Int., 184, 1223–1236, 2011.

Ruth, R.: A canonical integration technique, IEEE Trans. Nucl. Sci., 30, 2669–2671, 1983.

Sambridge, M. S.: Geophysical inversion with the Neighbourhood Algorithm - I. Searching a parameter space, Geophys. J. Int., 138, 479–494, 1999a.

Sambridge, M. S.: Geophysical inversion with the Neighbourhood Algorithm - II. Appraising the ensemble, Geophys. J. Int., 138, 727–746, 1999b.

Sambridge, M. S.: A parallel tempering algorithm for probabilistic sampling and multi-modal optimization, Geophys. J. Int., 196, 357–374, 2014.

Sambridge, M. S. and Drijkoningen, G. G.: Genetic algorithms in seismic waveform inversion, Geophys. J. Int., 109, 323–342, 1992.

Sambridge, M. S., Gallagher, K., Jackson, A., and Rickwood, P.: Trans-dimensional inverse problems, model comparison, and the evidence, Geophys. J. Int., 167, 528–542, 2006.

Sambridge, M. S., Bodin, T., Gallagher, K., and Tkalcic, H.: Transdimensional inference in the geosciences, Phil. Trans. R. Soc. A, 371, doi:10.1098/rsta.2011.0547, 2013.

Santosa, F. and Symes, W. W.: Computation of the Hessian for least-squares solutions of inverse problems of reflection seismology., Inverse Problems, 4, 211–233, 1988.

Sanz-Serna, J. M. and Calvo, M. P.: Numerical Hamiltonian problems, Chapman and Hall, London, UK, 1994.

Seah, Y.-L., Shang, J., Ng, H. K., Nott, D. J., and Englert, B.-G.: Monte Carlo sampling from the quantum state space. II, New J. Phys., 17, doi:10.1088/1367–2630/17/4/043 018, 2015.

Sen, M. K. and Biswas, R.: Tansdimensional seismic inversion using the reversible jump Hamiltonian Monte Carlo algorithm, Geophysics, 82, R119–R134, 2017.

Sen, M. K. and Stoffa, P. L.: Rapid sampling of model space using genetic algorithms: Examples from seismic waveform inversion, Geophys. J. Int., 108, 281–292, 1992.

Sen, M. K. and Stoffa, P. L.: Global optimization methods in geophysical inversion, Cambridge University Press, Cambridge, UK, 2013.

Symon, K. R.: Mechanics, Addison-Wesley, Reading, Massachusetts, 1971.

Taillandier, C., Noble, M., Chauris, H., and Calandra, H.: First arrival travel time tomography based on the adjoint state methods, Geophysics, 74, WCB57–WCB66, 2009.

Tarantola, A.: Inverse problem theory and methods for model parameter estimation, 2nd edition., Society for Industrial and Applied Mathematics, Philadephia, 2005.

Tarantola, A. and Valette, B.: Generalized Nonlinear Inverse Problems Solved Using the Least Squares Criterion, Rev. Geophys., 20, 219–232, 1982.

5   Tromp, J., Tape, C., and Liu, Q.: Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels, Geophys. J. Int., 160, 195–216, 2005.

von Rebeur-Paschwitz, E.: The earthquake of Tokyo, April 18, 1889, Nature, 40, 294–295, 1889.

Wei, Q., Dobigeon, N., and Tourneret, J.-Y.: Bayesian fusion of multi-band images, IEEE J. Select. Topics Signal Proc., 9, doi:10.1109/JSTSP.2015.2407 855, 2015.

10   Widiyantoro, S., Gorbatov, A., Kennett, B. L. N., and Fukao, Y.: Improving global shear wave traveltime tomography using three-dimensional ray tracing and iterative inversion, Geophys. J. Int., 141, 747–758, 2000.

Wolpert, D. H. and Macready, W. G.: No Free Lunch Theorems for optimization, IEEE Trans. Evolutionary Comp., 1, 67–82, 1997.

Yilmaz, O.: Seismic Data Analysis, Society of Exploration Geophysics, 2001.

Zhang, J. and Toksöz, M. N.: Nonlinear refraction traveltime tomography, Geophysics, 63, 1726–1737, 1998.

15   Zhao, H.: A fast sweeping method for eikonal equations, Mathematics of computation, 74, 603–627, 2005.

Zunino, A., Mosegaard, K., Lange, K., Melnikova, Y., and Mejer Hansen, T.: Monte Carlo reservoir analysis combining seismic reflection data and informed priors, Geophysics, 80, R31–R41, 2015.