

Hamiltonian Monte Carlo for the Purpose of Induced Source Characterization: Application to an M_L 3.4 Event in the Groningen Gas Field

La Ode Marzujriban Masfara, Cornelis Weemstra

¹Delft University of Technology, Stevinweg 1, 2628 CN, Delft, The Netherlands

²Royal Netherlands Meteorological Institute, Utrechtseweg 297, 3730 AE, De Bilt, The Netherlands

Key Points:

- A modified Hamiltonian Monte Carlo (HMC) algorithm is applied to recordings of a magnitude 3.4 induced event in the Groningen gas field
- Meaningful prior information is required to ensure the modified HMC algorithm samples the correct mode of the posterior probability
- Existing knowledge of the subsurface geology can be incorporated into the prior information of the algorithm

Corresponding author: L.O.M Masfara, l.o.m.masfara@tudelft.nl

Abstract

Hamiltonian Monte Carlo (HMC) is known to be highly efficient when sampling high-dimensional parameter spaces. This high efficiency can be attributed to Hamilton's equations, which guide the sampling of the model space. In the case of weakly non-linear problems, this efficiency can be increased even further by linearizing the forward problem. In this study, we exploit this for the purpose of estimating source parameters of a 3.4 magnitude induced event that originated in the Groningen gas field in 2019. In total, we estimate ten earthquake parameters: centroid (three coordinate components), moment tensor (six elements), and origin time. We demonstrate that, in the absence of a sufficiently accurate centroid prior, the linearization of the forward model necessitates multiple initial centroid priors. Here, we consider two sets of initial centroid priors. The first set is based on the known fault geometry in the Groningen reservoir, whereas the second set is obtained by placing initial centroid priors on a uniform horizontal grid at a depth of 3 km (the approximate depth of the gas reservoir). In general, the results from both sets are in good agreement with each other. Most important, however, is their agreement with the geological knowledge of the area: the posterior peaks for model vectors containing a centroid near a major fault and a moment tensor that corresponds to normal faulting along a plane which has a strike almost coinciding with the strike of that major fault.

Plain Language Summary

Earthquake source parameters, such as depth, time, and type of faulting, can be estimated using the recordings (or seismograms) of this (induced) earthquake. Being able to do this such that the uncertainty of the estimated parameters is also quantified is particularly valuable. This, however, requires the use of a probabilistic algorithm. A disadvantage of probabilistic algorithms is their computational cost. In this study, we simplify the relationship between the earthquake source parameters and the seismograms to significantly reduce computational costs. Specifically, we demonstrate that the simplified relation between the earthquake source parameters and the earthquake recordings requires the probabilistic algorithm to be provided with a sufficiently accurate (prior) estimate of these very earthquake source parameters. By means of a magnitude 3.4 induced event that originated in the Groningen gas field in 2019, we show that a geologically inspired prior can be helpful to partly overcome this: we use (known) existing faults in the reservoir to kick start the probabilistic algorithm. As such, we recover earthquake source parameters that are in line with subsurface geological information.

1 Introduction

Characterizing an earthquake is essential for a number of reasons. First, its source parameters (centroid, magnitude, slip direction, etc.) determine, to a large extent, the damage it may cause (Lui et al., 2016). This is because the depth, size, and type of rupture all affect the amount of shaking produced (Trippetta et al., 2019). Secondly, source characterization may help to improve our understanding of an event's nucleation, which is essential for developing reliable earthquake hazard models (Ellsworth et al., 2015). In addition, an increased understanding of source characteristics can potentially be used to improve earthquake early warning systems by providing (additional) information that can be used to generate alerts before strong shaking takes place (Peng et al., 2021).

Seismologists distinguish between ‘natural’ and ‘induced’ earthquakes. Induced earthquakes usually emit shorter period signals compared to tectonic earthquakes (Dais et al., 2018). This is because, on average, induced events have relatively low magnitudes compared to (stronger) tectonic earthquakes, although some induced events are reported to be as high as 5.8 (Foulger et al., 2018). In addition, induced events usually occur at rel-

actively shallow depths. Combined, shallower depths and higher frequencies imply that induced events may still cause significant damage to buildings and infrastructure (Vlek, 2018). In addition, ground motions are exacerbated by high amplification factors in some areas (Bommer et al., 2017).

A notable example of induced earthquakes is the events occurring in the Groningen gas field, the Netherlands (Sarhosis et al., 2019). The Groningen gas field, located in the northern part of the Netherlands, is the largest gas field in Europe. Since the first reported induced earthquake in 1986, there has been a gradual increase in seismic activity in the field (van Thienen-Visser & Breunese, 2015). Because of the societal unrest associated with the earthquakes (Nepveu et al., 2016), the Dutch government has recently taken steps to reduce the extraction of natural gas from the Groningen gas field. The field will close down permanently on October 1, 2024, with production expected to be halted on October 1, 2023. Parallel to the production reduction, an extensive array of seismometers was installed by the Dutch meteorological institute (KNMI, which stands for Koninklijk Nederlands Meteorologisch Instituut), funded by NAM (Nederlandse Aardolie Maatschappij), which is the major operator in the Groningen gas field (Ntinalexis et al., 2019). The array also includes borehole seismometers, enabling improved source characterization in the area (Smith et al., 2020), i.e., due to a significant increase of the signal-to-noise ratio (SNR) at depth (Ruigrok & Dost, 2019).

An earthquake source can be parameterized in several ways (Aki & Richards, 2002). In this study, we consider a moment tensor (MT) representation (Jost & Herrmann, 1989). This implies that the seismic event is collapsed to a single position (point-source representation), which is usually referred to as ‘the centroid’. Such a representation is justified in case the waveform data is analyzed at periods for which the seismic source is effectively a point source (Aki & Richards, 2002). Additionally, assuming instantaneous rupturing, we end up with ten source parameters. The first six are the moment tensor components, where the MT’s magnitude is a measure of the amount of energy released. This MT can be decomposed into isotropic (ISO), double-couple (DC), and compensated linear vector dipole (CLVD) components (Jost & Herrmann, 1989). The other four parameters are the event’s east, north, and depth coordinates and the origin time.

Various datasets and techniques have been utilized to estimate the source characteristics of Groningen earthquakes. Willacy et al. (2018) adopt a deterministic approach to estimate moment tensors and centroids. These authors employed a detailed 3D subsurface model of Groningen but restricted the search space to DC sources. In contrast, Dost et al. (2020) used a probabilistic approach to estimate the centroid and full moment tensor (implying that they allowed for the ISO and CLVD components as well) but employed (locally) 1D models. Deterministic approaches often provide faster computations compared to probabilistic approaches. However, probabilistic approaches quantify the uncertainty of the different parameters; in this case, these are the uncertainties of the ten earthquake source parameters. Also, the use of 3D subsurface models has a clear advantage over 1D subsurface models. This is because 3D models take into account the subsurface lateral heterogeneity that will affect the shape (amplitude and phase) of the seismogram generated from simulating an earthquake event using those 3D models.

In this study, we investigate the combination of a probabilistic approach with 3D subsurface models to estimate the source parameters of a real event in Groningen. To mitigate the aforementioned “inefficiency” of probabilistic approaches, we modify the workflow described in Masfara et al. (2022). This workflow relies on a variant of the Hamiltonian Monte Carlo (HMC) algorithm and has previously been tested using synthetic recordings generated using the 3D Groningen subsurface velocity model. For this study, we consider the 2019 3.4 local magnitude earthquake below the village of Westerwijtwerd (Figure 1). Since we estimate the full moment tensor, our estimation does not limit the search space to just DC components but includes the ISO and CLVD components. Also, the inclusion of origin time in the estimation quantifies the trade-off between origin time

and estimated depth. In what follows, we first describe the theory underlying the workflow. We subsequently introduce and discuss the (retrieval of the) recordings used to estimate the parameters, including the prior information that is used to increase the computational efficiency of the workflow. Finally, we compare our results to results obtained in other studies and draw conclusions, including the outlook of applying the same approach to a larger set of events in the Groningen area.

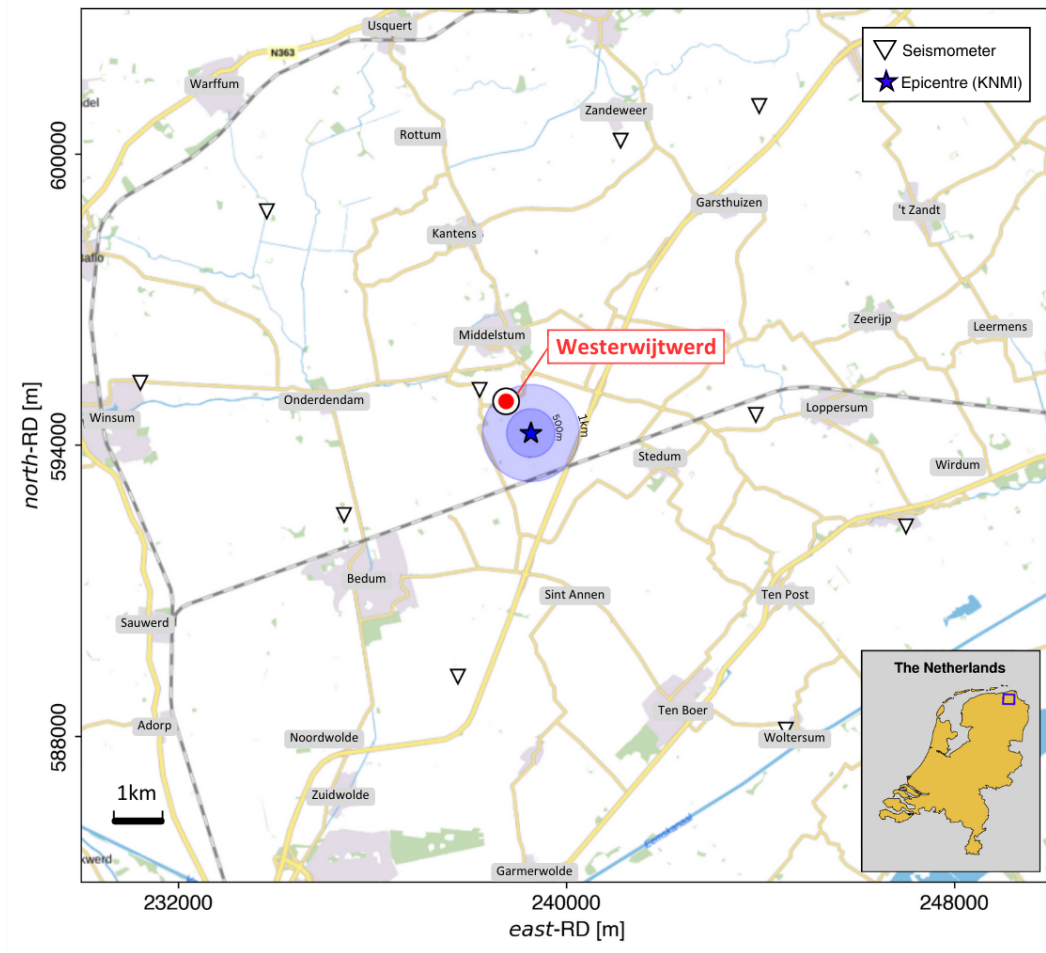


Figure 1. Map of the research area. The inverted triangles indicate the location of the KNMI seismometers, and the blue star is the epicenter of the 2019 Westerwijrtwerd earthquake, as estimated by the KNMI. Axes indicate location using the Dutch RD coordinate system. This specific coordinate system gives the geodetic coordinates for European Netherlands and is used in official national maps. The inset at the bottom right shows the location of the study area.

2 Methodology

To enable source characterization, the formal relationship between the observed data and the source (model) parameters is introduced and detailed in the first subsection. Subsequently, we introduce Bayes' theorem and, assuming Gaussianity, cast it in a form al-

lowing us to utilize it. In Subsection 2.3, we then introduce the HMC algorithm. Finally, in Subsections 2.4 and 2.5, we then describe how the algorithm’s efficiency can be enhanced via linearization of the forward problem and by choosing meaningful prior information, respectively.

2.1 The Forward Problem

In this study, the posterior probability of the model parameters is estimated by means of a Markov process. The generation of such a Markov chain is detailed further below (Subsection 2.3), but, at this point, it should be understood that for each sample in the chain, forward-modeled data is compared against measured data. In the context of our problem, a specific model \mathbf{m} (or sample) implies assigning a specific value to each of the ten aforementioned source parameters (MT, centroid, and origin time). The measured data \mathbf{d}^{obs} consists of the induced event’s waveform data, which, in our case, are recordings of particle displacement recorded by KNMI instruments. Computation of the likelihood $\rho(\mathbf{d}^{\text{obs}}|\mathbf{m})$ yields the probability of these recordings given a model \mathbf{m} and involves quantification of the misfit between the recorded particle displacements and numerically modeled particle displacements. The latter is computed by numerically solving the wave equation, i.e., they are the result of solving (what is usually referred to as) ‘the forward problem’. Mathematically, the forward problem can be written as

$$u_i(\mathbf{x}^{(r)}, t) = \sum_{j=1}^3 \sum_{k=1}^3 M_{jk}(t, T_0) * G_{ij,k}(\mathbf{x}^{(r)}, t; \mathbf{x}^{(a)}), \quad (1)$$

where u_i is the i^{th} component of the particle displacement vector ($\mathbf{u} = (u_1, u_2, u_3)$ where 1, 2, 3, correspond to the east, north, and down direction, respectively), M_{jk} represents an element of the 3×3 moment tensor \mathbf{M} at position $\mathbf{x}^{(a)}$, i.e., the centroid. Note that j and k indicate the axis along which the force is acting and the direction in which the arm is pointing, respectively (Aki & Richards, 2002). Furthermore, $\mathbf{x}^{(r)}$ denotes the position where the displacement is recorded, \mathbf{G} is the Green’s tensor, $*$ represents temporal convolution, and T_0 denotes the origin time. The comma after the second subscript of an individual element of the 3×3 Green’s tensor implies a spatial derivative in the k direction with respect to $\mathbf{x}^{(a)}$. To make the computation of $\mathbf{u}(\mathbf{x}^{(r)}, t)$ for a large number of potential centroids (i.e., a large number of $\mathbf{x}^{(a)}$) more efficient, we invoke reciprocity (Aki & Richards, 2002). In this study, the numerically modeled particle displacements are generated using SPEC3D (Komatitsch & Tromp, 2002). For this purpose, we use the 3D subsurface models of the Groningen gas field by Romijn (2017).

2.2 Bayes’ Theorem

The probabilistic workflow used in this study relies on Bayes’ theorem (or rule). In general, Bayes’ theorem describes how, in the presence of prior knowledge, the probability of a hypothesis (or model) \mathbf{m} depends on the available data \mathbf{d}^{obs} . The prior knowledge is accounted for by the prior probability distribution (often simply referred to as ‘the prior’). Ignoring the marginal probability (or ‘evidence’), Bayes’ theorem can be written as

$$\rho(\mathbf{m}|\mathbf{d}^{\text{obs}}) \propto \rho(\mathbf{d}^{\text{obs}}|\mathbf{m})\rho(\mathbf{m}), \quad (2)$$

where $\rho(\mathbf{m}|\mathbf{d}^{\text{obs}})$ is the posterior probability distribution (or simply ‘the posterior’), $\rho(\mathbf{d}^{\text{obs}}|\mathbf{m})$ the likelihood, and $\rho(\mathbf{m})$ the prior probability distribution. The model vector \mathbf{m} is a ten-component vector containing the centroid $\mathbf{x}^{(a)}$ (where a Cartesian east-north-down coordinate system implies that $\mathbf{x}^{(a)} = (x_1^{(a)}, x_2^{(a)}, x_3^{(a)})$; hence three model parameters), the moment tensor \mathbf{M} (six independent elements and hence six model parameters), and

the origin time T_0 (one model parameter). This implies that $\rho(\mathbf{m})$ represents the prior probability of these ten parameters.

Assuming Gaussian observational errors and a Gaussian distributed prior probability, the posterior in equation 2 can be written as (Fichtner & Simut , 2018; Masfara et al., 2022):

$$\rho(\mathbf{m} | \mathbf{d}^{\text{obs}}) \propto \exp \left(-\frac{1}{2} (\mathbf{d}(\mathbf{m}) - \mathbf{d}^{\text{obs}})^\top \mathbf{C}_d^{-1} (\mathbf{d}(\mathbf{m}) - \mathbf{d}^{\text{obs}}) - \frac{1}{2} (\mathbf{m} - \mathbf{m}^{(0)})^\top \mathbf{C}_m^{-1} (\mathbf{m} - \mathbf{m}^{(0)}) \right). \quad (3)$$

Here, $\mathbf{d}(\mathbf{m})$ contains the numerically modeled displacement recordings (solution of equation 1) and \mathbf{d}^{obs} the observed ones. Explicitly, for a total of N_r three-component instruments, $\mathbf{d}(\mathbf{m})$ is a concatenation of all $3 \times N_r$ modeled seismograms and \mathbf{d}^{obs} a concatenation of all $3 \times N_r$ recorded seismograms. \mathbf{C}_d , \mathbf{C}_m , and $\mathbf{m}^{(0)}$ are the data covariance matrix, prior covariance matrix, and prior mean, respectively. Evaluating equation 3 results in the (a posteriori) probability of the model parameters, i.e., their probability given observations and prior knowledge of the system (Tarantola, 2006).

2.3 Hamiltonian Monte Carlo

Although Bayes’ theorem describes how the posterior probability distribution depends on the available data \mathbf{d}^{obs} (through the likelihood) and prior knowledge $\rho(\mathbf{m})$, that posterior can usually not be estimated directly (Tarantola & Valette, 1981). In particular, a large number of model parameters and non-linearity prohibit this. To overcome this, we generate a sequence of specific models (often called ‘samples’) in what is referred to as a ‘Markov chain’. The density of these samples reflects the density of the posterior distribution we seek to find.

Numerous sampling algorithms are available to estimate $\rho(\mathbf{m} | \mathbf{d}^{\text{obs}})$, all with their own advantages and disadvantages. In this study, we implement a workflow that relies on the Hamiltonian Monte Carlo (HMC) algorithm. HMC was derived from classical mechanics, applied to statistical mechanics (Betancourt, 2017), and considered one of the most efficient probabilistic algorithms for exploring high-dimensional model spaces. HMC relies on the sequential calculation of two quantities. These are the ‘potential energy’ U , which is a function of the model vector \mathbf{m} , and the ‘kinetic energy’ K , which, in our framework, is solely a function of the momentum vector \mathbf{p} . This momentum vector is an auxiliary vector that has the same dimension as \mathbf{m} (ten in our case). Together, \mathbf{m} and \mathbf{p} make up what is often referred to as the ‘phase space,’ and their joint probability is described by the ‘canonical distribution’ $\rho(\mathbf{p}, \mathbf{m})$.

The canonical distribution can be written in terms of an invariant function $H(\mathbf{p}, \mathbf{m})$, i.e.,

$$\rho(\mathbf{p}, \mathbf{m}) = e^{-H(\mathbf{p}, \mathbf{m})}. \quad (4)$$

Here, $H(\mathbf{p}, \mathbf{m})$ is referred to as ‘the Hamiltonian’, and its value in phase space is usually called ‘the energy’ at that point (Neal et al., 2011). As such, a model \mathbf{m} can be looked upon as the position of a ‘particle’ (Betancourt, 2017).

Rewriting equation 4, and substituting the posterior probability (i.e., $\rho(\mathbf{p}, \mathbf{m}) \rightarrow \rho(\mathbf{p}, \mathbf{m} | \mathbf{d}^{\text{obs}})$), we have

$$\begin{aligned} H(\mathbf{p}, \mathbf{m} | \mathbf{d}^{\text{obs}}) &\equiv -\ln(\rho(\mathbf{p}, \mathbf{m} | \mathbf{d}^{\text{obs}})) \\ &= -\ln[\rho(\mathbf{p} | \mathbf{m})] - \ln[\rho(\mathbf{m} | \mathbf{d}^{\text{obs}})] \\ &= K(\mathbf{p}, \mathbf{m}) + U(\mathbf{m}). \end{aligned} \quad (5)$$

Here, $U(\mathbf{m}) \equiv -\ln \rho(\mathbf{m} | \mathbf{d}^{\text{obs}})$.

Equation 5 describes the more general case; in our implementation, $K(\mathbf{p}, \mathbf{m})$ is merely a function of the momentum vector and hence $K(\mathbf{p}, \mathbf{m}) \rightarrow K(\mathbf{p})$. Specifically, it is given

210 by (Fichtner & Simut , 2018; Masfara et al., 2022)

$$K(\mathbf{p}) = \mathbf{p}^T \mathcal{M}^{-1} \mathbf{p} / 2, \quad (6)$$

211 where the mass matrix \mathcal{M} acts as a tuning parameter (Fichtner et al., 2019, 2021), al-
212 lowing the particle to move through the desired areas of phase space with correspond-
213 ing potential and kinetic energy (Betancourt, 2017).

214 Starting from an initial estimate of \mathbf{m} with some prescribed initial momentum, Hamil-
215 ton’s equations, which read

$$\frac{d\mathbf{m}}{d\tau} = \frac{\partial K}{\partial \mathbf{p}}, \quad \frac{d\mathbf{p}}{d\tau} = -\frac{\partial U}{\partial \mathbf{m}}, \quad (7)$$

216 will efficiently explore areas with relatively low potential energies (corresponding to the
217 a posteriori more probable areas of the model space; see equation 5). Here, the quan-
218 tity τ is the ‘artificial time’ that is used to propagate (the particle) from the initial model
219 along trajectories of constant H . This propagation occurs for some (to-be-determined)
220 time τ_{lp} , where the subscript ‘lp’ stems from ‘leap’ as we use the leapfrog algorithm to
221 evaluate 7. The model reached at τ_{lp} , i.e., $\mathbf{m}(\tau_{lp})$, is subsequently accepted with prob-
222 ability

$$\theta = \min \left[1, \frac{\rho(\mathbf{p}(\tau_{lp}), \mathbf{m}(\tau_{lp}))}{\rho(\mathbf{p}, \mathbf{m})} \right], \quad (8)$$

223 which is usually referred to as the ‘metropolis rule’ (Tarantola, 2005). If the model $\mathbf{m}(\tau_{lp})$
224 is not accepted, the process will be repeated by introducing a new (different) momen-
225 tum vector to the initial model. If accepted, the model $\mathbf{m}(\tau_{lp})$ will serve as the start-
226 ing point for a new deterministic trajectory after being endowed with momentum.

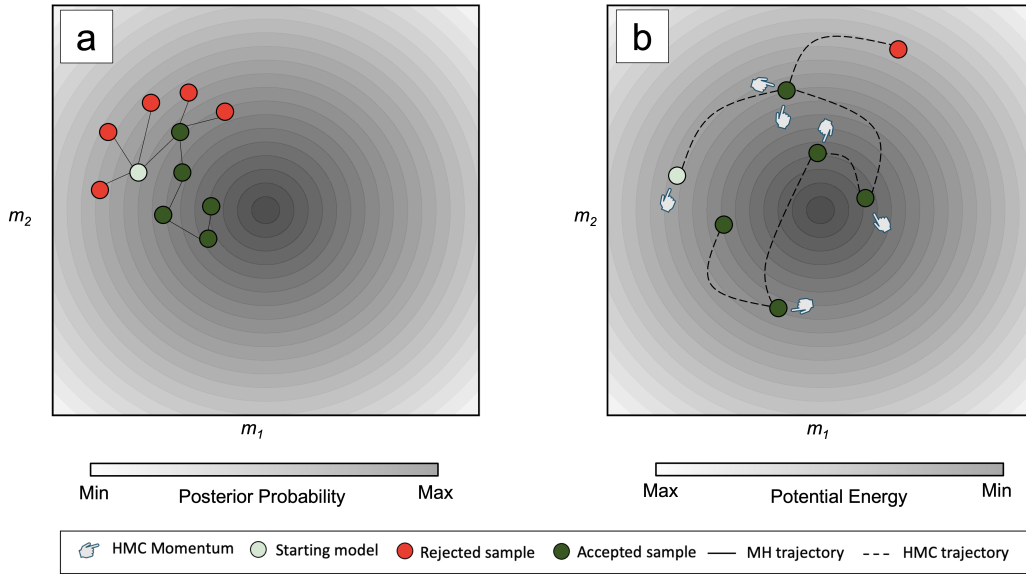


Figure 2. Illustration of model space exploration using Metropolis-Hastings (a) and Hamiltonian Monte Carlo (b) algorithms. Note that with a similar number of accepted samples, HMC explores the distribution more efficiently via a combination of iterative short and long trajectories. This is achieved by prescribing a different momentum for each trajectory and iterative computation of Hamilton’s equations. Mind that we only show the rejected samples of the first two moves/accepted samples for both algorithms.

2.4 Linearization of the Forward Problem

To ease the computation of the gradient of the potential energy in the model space, Fichtner and Simutè (2018) linearize equation 1 by means of a Taylor expansion around the prior mean $\mathbf{m}^{(0)}$ (see Appendix A). Simutè et al. (2022) use this same modification and 3D Earth models to characterize tectonic earthquakes below the Japanese peninsula. In these studies, $\mathbf{m}^{(0)}$ is obtained from an earthquake catalog, which is not always directly available for induced earthquakes. Replacing, in $\mathbf{d}(\mathbf{m})$, the numerically modeled displacements $\mathbf{u}(\mathbf{x}^{(r)}, t)$ by numerically modeled displacements resulting from a linear approximation of equation 1 implies that we assume $\mathbf{m}^{(0)}$ to be “sufficiently close” to the true model parameters. This merely applies for the centroid $\mathbf{x}^{(a)}$ and origin time T_0 . That is, since the particle displacement depends linearly on the moment tensor components, the linearization does not impose an approximation when it comes to the moment tensor components. Importantly, “sufficiently close” means that the centroid $\mathbf{x}^{(a)}$ and origin time T_0 should be at sub-wavelength and sub-period distance from the true centroid and origin time, respectively.

In our case, the assumption that $\mathbf{m}^{(0)}$ is sufficiently close to the true model parameters is usually not met. This will render the application of HMC ineffective (to state the least). In order to apply HMC (including a linear approximation of equation 1) to induced earthquakes, two main challenges, therefore, need to be addressed. First, the recorded seismograms are often dominated by high-frequency signals (>1 Hz), increasing the non-linearity of the forward problem. Second, as mentioned earlier, the prior information is often unavailable or rather inaccurate. To address these challenges, in this study, we use the multi-stage workflow introduced by Masfara et al. (2022). This means that we iteratively update $\mathbf{m}^{(0)}$, which is detailed in the remainder of this section. In addition, we run this workflow multiple times (in parallel), each starting from a different $\mathbf{m}^{(0)}$. This is explained in Section 2.5. In the remainder of this paper, we will refer to the HMC variant that involves a linearization of the forward problems as ‘linearized HMC’. It should be understood, however, that this does not involve a linearization of Hamilton’s equations itself.

Figure 3 illustrates the embedding of linearized HMC in the proposed multi-stage workflow. Iteratively updating $\mathbf{m}^{(0)}$ partly overcomes deviations of the estimated posterior from the true posterior, as such addressing the first challenge. Given a first $\mathbf{m}^{(0)}$, the three quantities in equations A4-A6 need to be computed only once in order to sample a “local posterior” around that $\mathbf{m}^{(0)}$. These quantities are used to compute the gradient of the potential energy and hence evaluate Hamilton’s equations and the Hamiltonian itself (equations 7 and 5, respectively). Importantly, in the absence of a linearization of the forward problem, the computation of equations 7 and 5 requires the forward problem to be evaluated during each deterministic trajectory. Linearization of equation 1, resulting in the three aforementioned quantities, renders this unnecessary for each individual stage (Masfara et al., 2022).

When $\mathbf{m}^{(0)}$ does not coincide with the true model parameters, the linearized HMC algorithm will explore a “local posterior” that deviates from the true posterior distribution despite being computationally efficient. This is illustrated in Figure 3(a), where the linearized HMC can only explore the area above the orange curve. To obtain a better approximation of the posterior, the workflow uses the result of exploring the local posterior in Figure 3(a) to obtain a new $\mathbf{m}^{(0)}$ (essentially taking the mean of the local posterior and using that as $\mathbf{m}^{(0)}$). Linearization of the forward problem about the updated $\mathbf{m}^{(0)}$ and re-computation of the aforementioned quantities allows for a new exploration of the model space in Figure 3(b) and (c). After five Taylor expansions about the new $\mathbf{m}^{(0)}$, six local posteriors are estimated. The associated distributions are, for each stage, depicted in Figure 3(d). Having the results from all stages in (d), the workflow then uses variance reduction (e.g., Mustač & Tkalčić, 2016; Masfara et al., 2022) as a

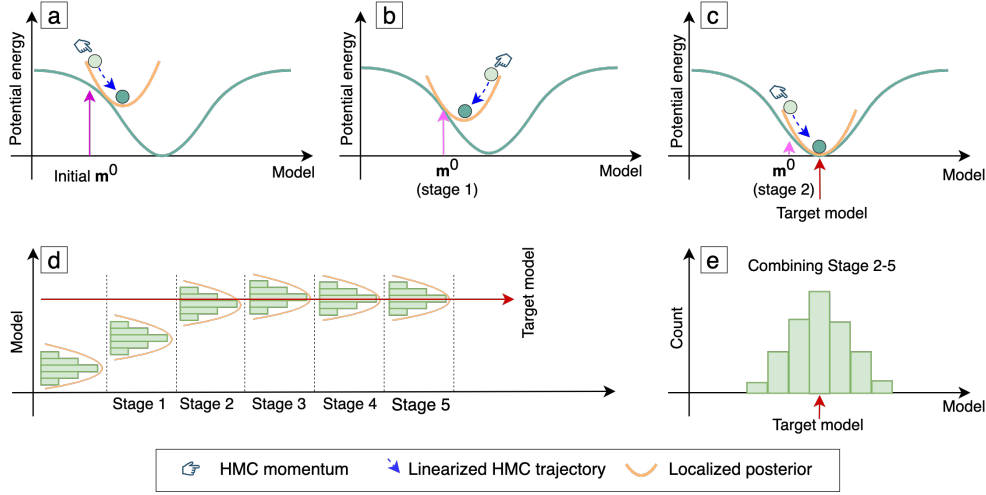


Figure 3. Illustration of linearized HMC embedded in the multi-stage workflow detailed in Masfara et al. (2022). Panel (a) to (c) depict the sampling of a local posterior associated with different $\mathbf{m}^{(0)}$. In (a), $\mathbf{m}^{(0)}$ is the initial prior mean. In the next stage (b), $\mathbf{m}^{(0)}$ is updated using the results of the exploration of the local posterior associated with this initial prior mean until $\mathbf{m}^{(0)}$ (almost) coincides with the most likely model (c). The workflow’s progression up to five stages is shown in (d). (e) is the final posterior composed using variance reduction criterion, which discriminates the first two stages from stages 2 to 5.

criterion to select stages that should be included in the estimate of the final posterior. This is depicted in Figure 3(e).

2.5 The Importance of the Prior

Having an inaccurate $\mathbf{m}^{(0)}$ can only partly be overcome by updating $\mathbf{m}^{(0)}$ in progressive stages. That is, the multi-stage workflow will still be ineffective when the initial $\mathbf{m}^{(0)}$ is located in a “local mode” of the posterior distribution (i.e., associated with a local minimum of the potential energy). The chance of this happening increases with an increase in the non-linearity between the model parameters and the observed displacement recordings (i.e., higher frequencies). In practice, this happens when the centroid $\mathbf{x}^{(a)}$ and origin time T_0 in $\mathbf{m}^{(0)}$ are separated from the true centroid and true origin time by more than (approximately) half a wavelength or half a period, respectively. To address this, we additionally use multiple initial $\mathbf{m}^{(0)}$ concatenated in a list which we denote by $\mathbf{m}_{\text{list}}^{(0)}$ (the list consists of $\mathbf{m}_1^{(0)}$, $\mathbf{m}_2^{(0)}$, ..., $\mathbf{m}_N^{(0)}$ with N being the total number of $\mathbf{m}^{(0)}$). These initial $\mathbf{m}_i^{(0)}$ differ to the extent that the centroid position is different for each of them. The use of $\mathbf{m}_{\text{list}}^{(0)}$ is to ensure some of the individual $\mathbf{m}_i^{(0)}$ are contained in the global minimum. The same criterion is used to select which (local posterior) distributions can be included in the final posterior (i.e., which stages). That is, the variance reduction is now computed for all stages associated with the individual (initial) $\mathbf{m}_i^{(0)}$ in $\mathbf{m}_{\text{list}}^{(0)}$. We illustrate the process of using multiple $\mathbf{m}^{(0)}$ in Figure 4. We depict three initial $\mathbf{m}^{(0)}$, with one located in the “correct” lobe, that is, $\mathbf{m}_2^{(0)}$. Each of the $\mathbf{m}^{(0)}$ will then be updated in a similar fashion as shown in Figure 3. While $\mathbf{m}_1^{(0)}$ and $\mathbf{m}_3^{(0)}$ ended up sampling the wrong lobe, the updated $\mathbf{m}_2^{(0)}$ enables the linearized HMC algorithm to sample the correct lobe. In Figure 4(b), we detail the last stage of the multi-stage workflow that started with $\mathbf{m}_2^{(0)}$ in the red circle.




Figure 4. (a) Illustration of using multiple initial $\mathbf{m}^{(0)}$ while sampling a complex/multimodal posterior distribution using linearized HMC. (b) zoom of the last stage of the multiple stages associated with the initial model prior $\mathbf{m}_2^{(0)}$.

We end this section by emphasizing that although being very efficient in sampling the posterior distribution (through the potential energy), the proposed multi-stage workflow (including the use of multiple initial priors $\mathbf{m}_i^{(0)}$) ultimately only results in an approximate posterior distribution. This is because the true observational errors are not necessarily Gaussian and uncorrelated (which we assume in this study) and because we linearized the relation between observed particle displacement and model parameters. In addition, the 3D velocity model used to model (numerically) displacement recordings (according to equation 1) is assumed to coincide with the true velocity model. Since this will not be the case, another “source of error” is introduced, which in practice will result in a deviation of the estimated posterior from the true posterior. Moreover, since a Markov process only approaches the true posterior asymptotically, a Markov-chain-based estimate of the posterior is, by definition, an approximation. Whereas the latter two cannot be circumvented (we don’t have the exact subsurface model and also cannot run a Markov chain for an infinite amount of time), the linearization is, in principle, not necessary, and also Gaussian observational errors do not need to be assumed. Not doing so, however, would make the computational demands prohibitively large.

3 Data

In this study, \mathbf{d}^{obs} contains the $3 \times N_r$ recordings of displacements (\mathbf{u}^{obs}) due to an induced event that occurred close to the village of Westerwijtwerd in 2019, the province of Groningen (see Figure 1). The KNMI estimates the magnitude of the earthquake to be 3.4 local magnitude. We collected \mathbf{u}^{obs} from ten G-network seismometers. These seismometers are selected based on their distance and azimuthal coverage with respect to the estimated epicenter. In Figure 5(a), we depict the ten seismometers as white inverted triangles and the location of the KNMI-estimated epicenter by a blue star. The seismometers are part of the KNMI borehole network: each borehole contains four vertically-separated seismometers. The number at the end of their ID indicates their depth, i.e., their IDs run from ..1 to ..4, with the instruments numbered ..1 being at 50 m depth and the instruments numbered ..4 being at 200 m depth. We illustrated the configuration of a string of borehole seismometers in Figure 5(b).

From the four seismometers in each borehole, we solely used the seismograms recorded by the deepest seismometers: they have a higher signal-to-noise ratio than the shallower seismometers (Dost et al., 2012). Furthermore, all seismometers experience a horizon-

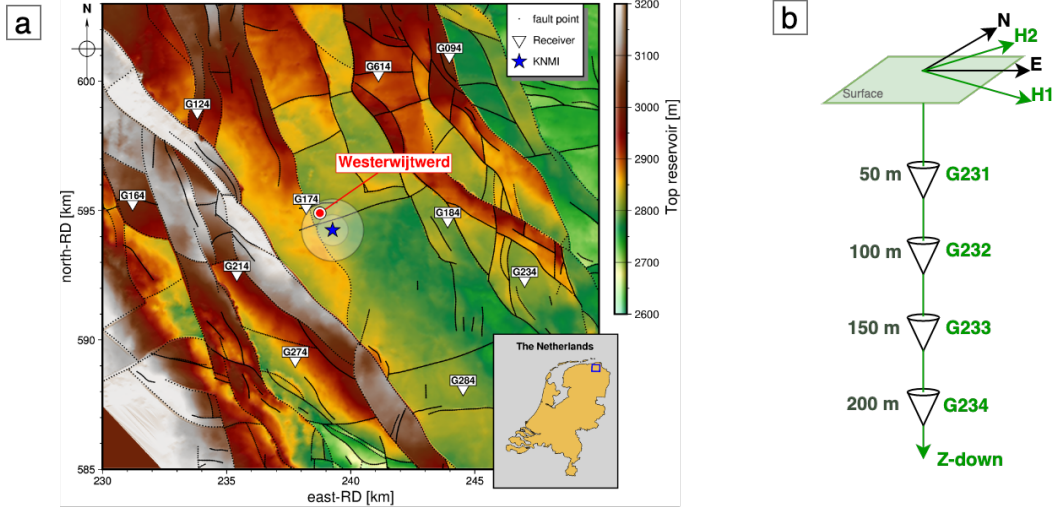


Figure 5. (a) Depth of top Rotliegend (reservoir) in area of interest. Solid black dots delineate mapped faults (Bourne & Oates, 2017). The inset at the bottom right shows the location of the study area. (b) Illustration of borehole seismometers in the G-network.

tal rotation while lowering them in the borehole. Consequently, a rotation needs to be carried out for projecting the horizontal recordings to specific preferred orientations, which in our case are to the east-west(x_1 -axis) and north-south(x_2 -axis) orientations, respectively. In Figure 5(b), we illustrate the orientation of the deepest borehole seismometer. The axis H1 and H2 are proxies of east and north. We then rotate the data to the true east and north using the angles given in Ruigrok et al. (2019). We depict the original seismograms (obtained from the KNMI) and the rotated seismograms of the selected seismometers in Figure 6.

Dost et al. (2020) have used the same recordings to characterize the Westerwijtwerd event probabilistically. These authors, however, use local 1D velocity models to solve the forward problem. Furthermore, they separately use 0.5 and 1s windows of P and S waves, respectively, where the P-wave is given more weight and evaluated at higher frequencies (i.e., 2-4 Hz for P and 1-3 Hz for S-wave). The P-wave waveform is given a higher weight because of the higher accuracy of the employed P-wave velocity models (compared to the S-wave velocity models). Also, these authors only use the vertical components of the recorded P-wave and the transverse component of the recorded S-waves. To account for inaccuracy in the velocity models, they allow individual, station-specific shifts of 0.1 s for both wave types. Another study in the area is by Smith et al. (2020), which uses a coherence method. This study focuses on determining the hypocenter. They find most Groningen earthquakes to systematically originate approximately 200 m above the reservoir layer. In this study, we exclusively use P-wave seismograms due to the significantly higher accuracy of the P-wave model. Furthermore, we use both the vertical and horizontal components and filter the recordings using a passband of 1-4 Hz, similar to the frequency range used by Dost et al. (2020). As for the length of the measurement window, we use 2.5s for all components and taper both ends with a 0.5 s cosine taper. For the data covariance, we use a diagonal matrix representing uncorrelated noise and estimate this to be 5% of each component's maximum amplitude. By taking a certain fraction of the maximum amplitude, we overestimate the 'true noise'. The reason for this is that we want to account for (part) of the waveform misfit arising from the deviation of the employed velocity model Romijn (2017) from the true (unknown) velocity model.




Figure 6. Observed seismograms before (green) and after rotation/polarity switch (black). Recordings are normalized (individually) with respect to maximum particle displacement (written in blue).

Before applying it to the field data, we perform a synthetic experiment, detailed in the next section.

4 Synthetic Experiment

In this section, we test the validity of the proposed workflow and data processing parameters (i.e., frequency band, length of the measurement window, and noise criteria) on a synthetic event. For this, we first generate synthetic data using the KNMI-estimated hypocenter as the centroid of our synthetic earthquake. We then set T_0 to 3 s, and for the MT, we use the values of 0.2E13 Nm, 2.86E13 Nm, -3.07E13 Nm, 0.76E13 Nm, -0.45E13 Nm, -1.71E13 Nm for M_{nn} , M_{ee} , M_{dd} , M_{ne} , M_{nd} , and M_{ed} respectively. These values represent pure shear normal faulting (rake of -90°) along a geological fault with a strike of 165°, a dip of 60°, and a moment magnitude of 3. We then corrupt the data in the frequency domain to simulate the presence of uncorrelated noise. This is implemented using the same approach as Mustač and Tkalčić (2016). In the time domain, the uncorrelated noise results in amplitude variations that affect the estimation of our centroid and MT, and shift the observed recordings in time (resulting in uncertainty in T_0). To effectively test the workflow, we first choose a (single) $\mathbf{m}^{(0)}$ that significantly deviates from the actual value (i.e., the synthetic earthquake parameters). For the centroid, we impose a shift of 200 m along each axis, i.e., the centroid in $\mathbf{m}^{(0)}$ deviates 200 m from $x_{east}^{(a)}$, $x_{north}^{(a)}$, and $x_{depth}^{(a)}$. For the MT, we simply assign a uniform value to each MT component, and for T_0 , we impose a shift of 0.5 s. We then run our workflow for 20 stages (i.e., the prior mean $\mathbf{m}^{(0)}$ is updated twenty times). The results are presented in Figure 7.

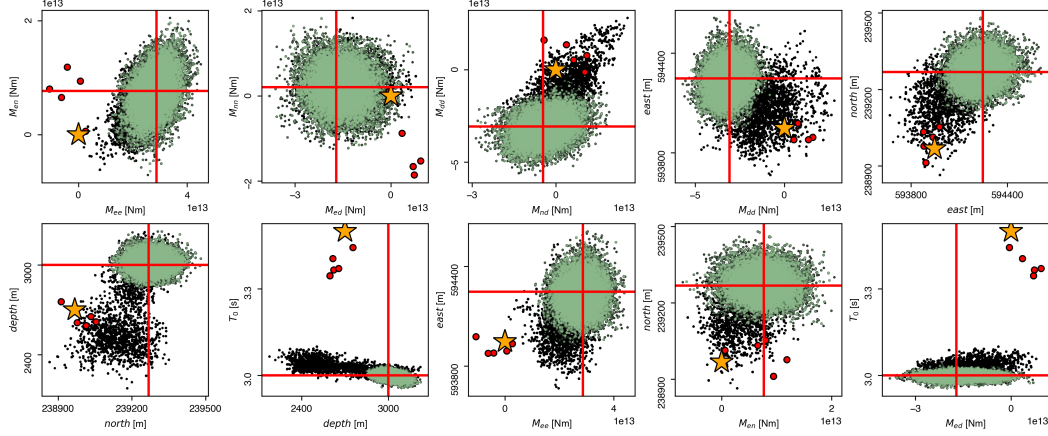


Figure 7. Marginal posterior probabilities obtained through applying the proposed linearized HMC workflow to synthetic recordings. The stars represent the initial $\mathbf{m}^{(0)}$. The black and green dots represent all accepted samples from all 20 stages and samples from selected stages (i.e., the VR-score exceeds 0.95) used to compose the final posterior, respectively. The red lines represent the true (synthetic) model parameters, and the red dots are the samples generated running the generic (non-linearized) HMC algorithm.

The yellow stars represent the initial $\mathbf{m}^{(0)}$, and the red lines represent the true synthetic earthquake parameters. The black dots are the samples generated from all 20 stages, which are equivalent to samples used to build all the histograms from exploring local posteriors in Figure 3(d). Whereas the green dots are the samples from selected stages based on a VR criterion, equivalent to the samples from the selected stages in Figure 3(e). The red dots represent samples resulting from a generic HMC run (i.e., HMC without linearizing the forward problem). This run was terminated as soon as the number of times for which the forward problem needed to be solved coincided with the number of times the forward problem was solved while running the multi-stage workflow in which the forward problem was linearized. Mind that each solution of the forward problem involves the computation of 3×10 seismograms (recall from Section 3 that $N_r = 10$).

Let us demonstrate the computational benefit of the multi-stage workflow (in conjunction with a linearization of the forward problem) over generic HMC (which does not involve this linearization). The number of times the forward problem needs to be solved in order to generate four model samples using generic HMC (represented by the star and the red dots in Figure 7) is 404. Here, each ‘solution of the forward problem’ in practice involves a separate computation of the $u_i(\mathbf{x}^{(r)}, t)$ in equation 1. We arrive at 404 as follows: it depends on the number of generated samples N_s (4 in this case), the number of leaps N_{lp} to arrive at $\mathbf{m}(\tau_{lp})$ (here we use 5), and the number of model parameters N_m (10 in our case). First, with the prescribed five leaps to arrive at a new model started from the current model, we evaluate equation 7 five times. Second, the evaluation of equation 7 requires the computation of $\frac{\partial U}{\partial \mathbf{m}}$. For that, we use a central difference approximation, which means that for each of the ten parameters in \mathbf{m} , we must evaluate U twice. Additionally, after the five leaps, we still have to compute $\rho(\mathbf{p}(\tau_{lp}), \mathbf{m}(\tau_{lp}))$ to evaluate equation 8, which requires one additional solution of the forward problem per sample. Consequently, the total number of forward problem solutions coincides with $N_s \times N_{lp} \times 2N_m + N_s = 404$. Linearization of the forward problem reduces this number dramatically. In fact, for every stage of the multi-stage workflow, the number of samples that can be generated is unlimited in the sense that it does not require additional

solutions to the forward problem. The forward problem just needs to be run $2 \times N_m = 20$ times per stage. This number stems from the (one-time) computation of the derivatives of U . These derivatives are included in the A_{pq} , b_p , and c (equations A4, A5, and A6 in appendix Appendix A, respectively). Therefore, to generate all samples for a total of 20 stages (i.e., 20 updates of $\mathbf{m}^{(0)}$), the number of times the solution to the forward problem needs to be computed is just 400.




Figure 8. Seismograms modeled using the posterior mean (gray) compared to the modeled observed recordings with noise added (brown) and the modeled observed recordings without noise (green).

We use the mean of the approximate posterior resulting from our multi-stage workflow to generate displacement recordings. In Figure 8, we compare these recordings with the observed (synthetic) recordings. The observed recordings are depicted in brown (recall that noise is added to these seismograms). The recordings associated with the mean values of our estimated posterior are depicted in grey and align well with the noise-free recordings associated with the true source parameters (depicted in green).

5 Prior Knowledge

In Subsection 2.5, we discussed the importance of using $\mathbf{m}_{\text{list}}^{(0)}$ to avoid getting trapped in a local mode. For the purpose of generating $\mathbf{m}_{\text{list}}^{(0)}$, we make use of the available fault map of Groningen's subsurface by Bourne and Oates (2017). This is inspired by research that reveals a strong correlation between hypocenters and major faults in Groningen's subsurface (Pickering, 2015; Spetzler & Dost, 2017; Willacy et al., 2018). In this context, we also evaluate the importance of the displacement along the horizontal components for the estimated posterior. The reason for this is potential errors arising from possible incorrect rotations of the horizontal displacements (see Section 3). Combined, we,

therefore, investigate three different cases: two centroid prior configurations (i.e., with different $\mathbf{m}_{\text{list}}^{(0)}$) of which one is used in conjunction with both the vertical component recordings and the three-component recordings. The configuration that uses known faults in the reservoir as a basis to generate the $\mathbf{m}_{\text{list}}^{(0)}$, in conjunction with the vertical component recordings only, is referred to as ‘1C-fault’. The same configuration, but used to estimate the posterior based on the recordings by all three components, is referred to as ‘3C-fault’. The other centroid prior configuration we consider consists of a square grid that covers not just the fault but also the surrounding area. This configuration of $\mathbf{m}_{\text{list}}^{(0)}$ is only used in conjunction with the recordings by all three components and is referred to as the ‘3C-grid’. This centroid prior configuration is considered to evaluate whether the recovered posterior might peak at a centroid position that deviates from the known fault geometry. The two different centroid prior configurations are depicted in Figure 9.

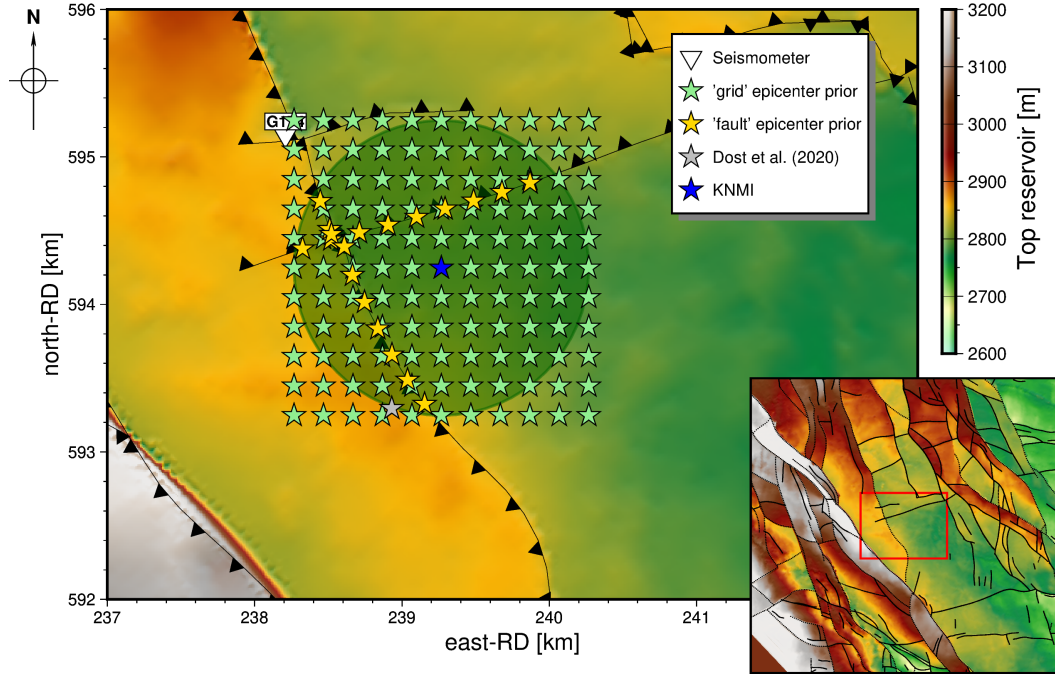


Figure 9. Horizontal positions of the different centroid priors for the two different prior configurations considered. The first configuration is guided by the known fault geometry inside the green circle and is represented by the yellow stars. This circle has a 1km radius and is centered at the epicenter estimated by the KNMI (blue star). The second centroid prior configuration uses a 2 km x 2 km grid with again the KNMI-estimated epicenter at its center. These centroid priors are depicted as green stars.

To generate the entries (individual $\mathbf{m}^{(0)}$) in $\mathbf{m}_{\text{list}}^{(0)}$ of the two considered prior configurations, we first draw a circle with a 1 km radius around the epicenter estimated by the KNMI. The enclosed area is colored dark green in Figure 9. Next, we discretize the fault inside the circle using a spatial sampling criterion based on the approximate seismic P-wave velocity within the circle and the highest frequency we use while fitting the waveforms. This criterion provides a rough estimate of the minimum “wavelength” of the posterior distribution. By discretizing the fault such that the individual centroids (associated with individual $\mathbf{m}_i^{(0)}$) in $\mathbf{m}_{\text{list}}^{(0)}$ are separated by less than half this wavelength, we therefore, ensure that at least one of the initial priors is located in the “correct” lobe, i.e., similar to what we have illustrated in Figure 4. Given the P-wave velocities at reser-

voir depth and a maximum frequency of 4 Hz (recall that we filter the recordings using a passband of 1-4 Hz), we arrive at a value of 200 m for this criterion. This is hence the separation along the fault at which individual centroid priors are placed. We depict these initial centroid priors in Figure 9 as yellow stars. At the same time, the fault orientations at these positions are used to determine the six moment tensor entries in the initial priors. As for the depth and origin time T_0 in the $\mathbf{m}_i^{(0)}$, we use the values estimated by the KNMI for both configurations (i.e., 3 km for the depth and 2019-05-22T03:49:00.075s for the origin time). In total, 19 individual $\mathbf{m}^{(0)}$ are concatenated in $\mathbf{m}_{\text{list}}^{(0)}$ for 1C-fault and 3C-fault. For the third case, we consider a centroid prior configuration consisting of a square grid of 2 km \times 2 km, with the center again being the epicenter estimated by the KNMI. We use the same criterion (200 m) to determine the horizontal spacing between the individual centroid priors. In Figure 9, we depict these as green stars. For the depth and origin time, we use identical values. Furthermore, for the MT, we assign a uniform value to each MT component for each individual $\mathbf{m}^{(0)}$. In total, we obtain 121 initial $\mathbf{m}^{(0)}$ for this configuration.

6 Application to Field Data

For all cases described above (1C-fault, 3C-fault, 3C-grid), our multi-stage workflow consists of 20 stages. For the centroid prior configuration derived from the geometry of the known faults within the reservoir $\mathbf{m}_{\text{list}}^{(0)}$ contains $N = 19$ $\mathbf{m}^{(0)}$, which implies a total of 380 stages. For the 3C-grid, a total of 121 initial priors serve as starting model of the 121 multi-stage workflows (see Figure 9), resulting in a total of 2420 stages for this configuration. For each stage, we then compute the VR score based on the recordings $\mathbf{u}(\mathbf{x}^{(r)}, t; \bar{\mathbf{m}})$ associated with the mean model $\bar{\mathbf{m}}$ of all 3000 individual models within that stage. Stages for which the VR score exceeds 0.95 are subsequently used to build our final posterior distribution. For each of the three cases considered, and for each of the initial centroid prior means, we show in Figure 10 the VR score associated with that $\bar{\mathbf{m}}_i^{(0)}$ of the 20 $\mathbf{m}^{(0)}$ in $\mathbf{m}_{\text{list}}^{(0)}$ for which the VR score attains its maximum. Note that this model's centroid is usually not at the location of the initial centroid prior mean (i.e., the centroid in $\mathbf{m}_0^{(0)}$) because the models for which the waveforms best fit the observed recordings are often found in one of the later stages; see also Figure 3. For all three cases considered here, the highest VR scores are obtained in those chains for which the initial centroid prior mean is close to a fault.

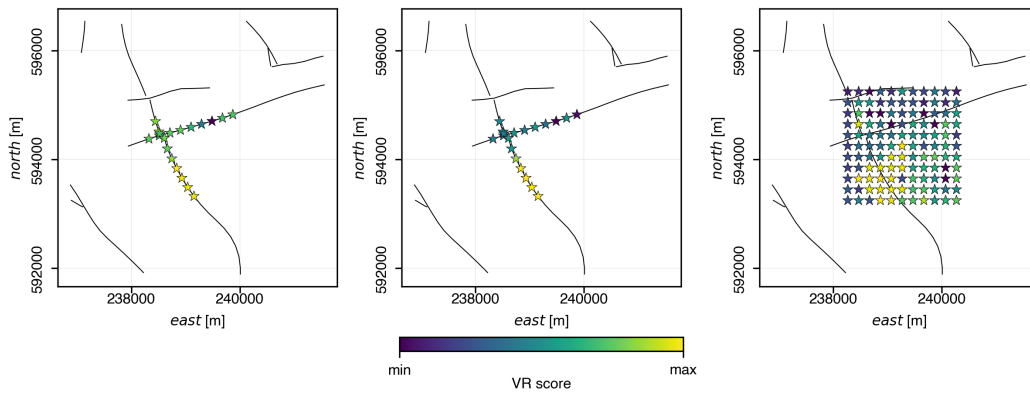


Figure 10. Maximum VR score in each of the chains associated with the different initial $\mathbf{m}^{(0)}$ for the three different cases considered (from left to right: 1C-fault, 3C-fault, and 3C-grid, respectively). Note that here we represent them by plotting the initial prior means (of the lateral positions) of the centroid.

6.1 Estimated Posterior

In Figure 11, we display the 1D marginal posterior distributions obtained from the selected stages of each configuration. In general, the mean value of these posteriors is fairly consistent across configurations, especially for 3C-fault and 3C-grid. For the 1C-fault case, the mean of the posteriors slightly deviates while at the same time having a slightly broader distribution compared to the other two cases. We attribute this to the fact that, for 3C-fault and 3C-grid, the additional data reduces the uncertainty of the estimates. In Figure 12, and for 3C-fault, we also plot the progression of the different stages associated with one of the individual centroid priors included in one of the $\mathbf{m}^{(0)}$ in $\mathbf{m}_{\text{list}}^{(0)}$. Specifically, we show the progression of that workflow (i.e., starting from that $\mathbf{m}_i^{(0)}$) that contains the stage that results in the overall maximum VR score. The vertical lines represent the start of different stages, and the red horizontal lines are the posterior means computed using the selected stages (after evaluating the VR scores for all stages). The progression follows a trend identical to the illustration in Figure 3(d), especially for the origin time T_0 with a slight variation for some others, such as for the depth and M_{nd} that shift monotonically to lower values. It is important to add that an initial estimate of T_0 was obtained using the envelope of the traces. This is described in detail in Section 6.1 of Masfara et al. (2022).

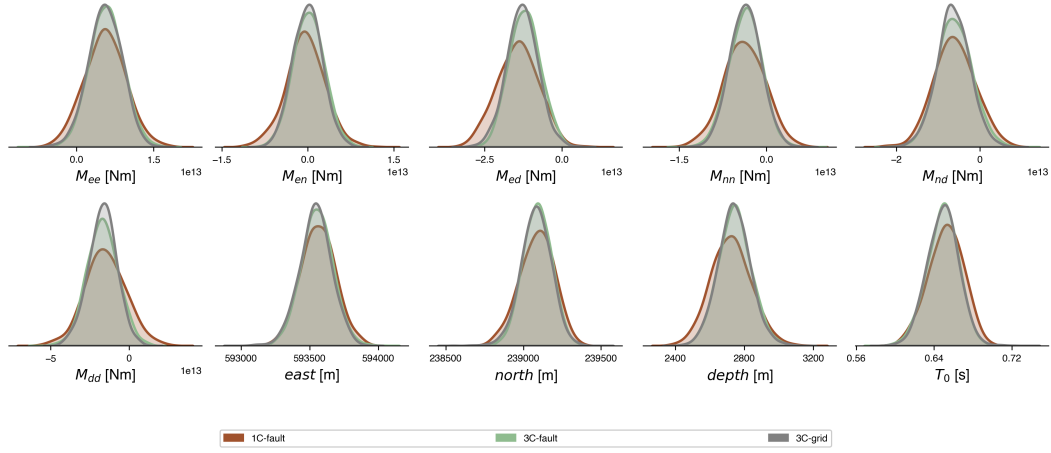


Figure 11. 1D marginal posterior distributions for the three different cases considered. '1C-fault': initial centroid prior configuration derived from the geometry of the known faults within the reservoir, and only the vertical particle displacement recordings are used. '3C-fault': initial centroid prior configuration derived from the geometry of the known faults within the reservoir, but both horizontal and vertical particle displacement recordings are used. '3C-grid': initial centroid prior located on a regular grid in a horizontal plane at the approximate (expected) depth of the event, and both horizontal and vertical particle displacement recordings are again used.

6.2 Traces Associated With the Posterior Distribution

Using the posterior mean in Figure 11, we generate synthetic data and compare these with the observed data in Figure 13. In our workflow, the misfit in equation A1 is based on 2.5 seconds of the observed particle displacement, bandpass filtered between 1 and 4 Hz. Here, for consistency, we adopted the same values for these parameters. Additionally, we show in Figure 13 the maximum and minimum bounds using synthetic data

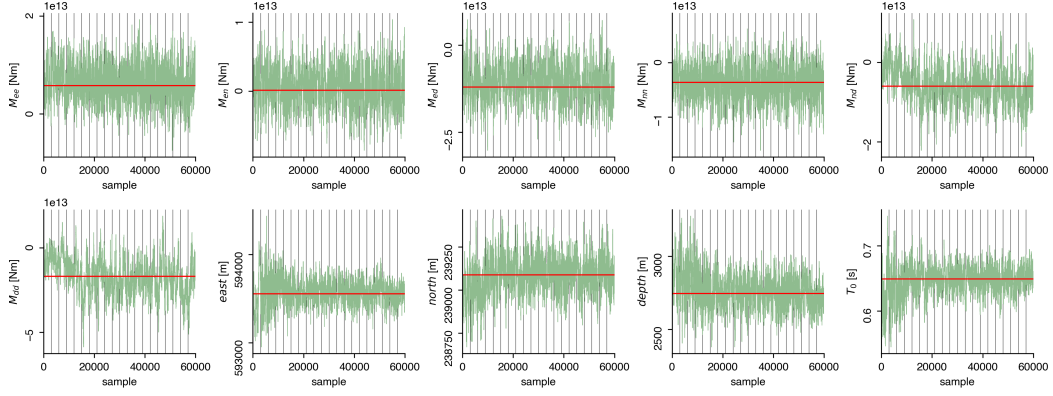


Figure 12. Progression of 20 stages from using one of the $\mathbf{m}^{(0)}$ in the 3C-fault configuration. The vertical lines represent different stages, whereas the red lines are the posterior mean (i.e., the mean of the green distributions in Figure 11) obtained from the selected stages based on the VR criteria for 3C-fault configuration.

generated from 1000 models drawn from the posterior distribution. We depict those bounds as a shaded area in Figure 13.

6.3 Source Characteristics

To investigate the source characteristics of the analyzed induced event, we first decompose the MTs of the posteriors shown in Figure 11. In this study, we do not limit our solutions to a single mechanism. We, therefore, decompose our moment tensor solutions into their ISO, DC, and CLVD components. We do this for each case (1C-fault, 3C-fault, and 3C-grid) and depict the decompositions in the Hudson plots in Figure 14. The mean MT for each case is represented by the beachball with the red outline. For all cases, the DC “region” is densely clustered (i.e., the center of the plot), with negative ISO components clearly outnumbering positive ISO components. This is often attributed to the compaction due to the gas extraction (Dost et al., 2020). We show the posterior distributions of the different MT components in Figure 15 (bottom row). Furthermore, in the top row, we depict the translation of the MT solutions in Figure 14 to distributions of strike, dip, and rake. Here, we only show solutions with strikes between 90° and 180° , which are in accordance with the orientation of the fault close by (given the centroid posterior distributions).

We visualize the centroid posterior distributions using horizontal and vertical slices of the Groningen subsurface (Figure 16). In the top row, we show the depth of the top reservoir as a contour map, including the location of faults from Bourne and Oates (2017) at that depth. On top of these contour maps, we show the samples used to generate the 2D marginal posterior distributions of the lateral position of the centroids. We also plot the result from Dost et al. (2020) and the KNMI as the black beachball and blue star, respectively. The red beachball represents the mean MT which is also depicted in Figure 14 (beachball with red outlines). Not only do the posterior means of the (lateral) centroid positions coincide with the known fault, but also does the moment tensor solution agree quite well with the strike of the nearby fault. On the vertical slices (middle and bottom rows), we depict the depth of the top reservoir as solid black lines. The location of the east-west vertical cross section and the north-south vertical cross section are shown as red and blue lines in the contour maps, respectively. For this specific earthquake, we find the posterior mean of the centroid to be slightly shallower than the centroid esti-




Figure 13. The comparison between observed and numerically modeled seismograms. The modeled seismograms are generated given the posterior mean for estimated for each of the considered cases (see Figure 11). The shaded area is within the maximum and minimum bounds of a total of 1000 waveforms generated using 1000 models drawn from the posterior distributions in Figure 11. Each seismogram is filtered and tapered using the same parameters used in the multi-stage HMC workflow. The duration of each trace plotted here is 3.25s.

546 mated by Dost et al. (2020). In fact, instead of being within the reservoir, we find the
 547 probability of having the earthquake nucleated above the reservoir is higher. The earth-
 548 quake (model) parameter that has the strongest trade-off with depth is origin time. This
 549 is because an earlier origin time can be translated to an earthquake occurring at greater
 550 deeper and vice versa. In this study, origin time uncertainty is considered, and the re-
 551 sult shows that the estimated T_0 from the KNMI is lagging by a few milliseconds. As
 552 a caveat, however, we do not consider the uncertainty in the 3D velocity models, which
 553 may not only introduce amplitude variations but also affect the origin time and/or depth.
 554 For a more detailed comparison, in Table 1, we list the mean and standard deviation of
 555 our estimated parameters (for the MTs, we convert these into strike, dip, and rake so-
 556 lutions) and compare them with the result of Dost et al. (2020) and the KNMI (hypocen-
 557 ter only).

Configuration	Strike [°]	Dip [°]	Rake [°]	North [m]	East [m]	Depth [m]	T_0 [s]
1C-fault	163.9 ± 16.9	71.5 ± 10.6	-100.7 ± 11.4	593562 ± 125	239086 ± 111	2722 ± 115	0.652 ± 0.018
3C-fault	162.5 ± 15.0	69.8 ± 10.0	-99.8 ± 11.7	593554 ± 116	239088 ± 91	2744 ± 96	0.649 ± 0.016
3C-grid	162.2 ± 12.4	70.6 ± 8.6	-100.7 ± 8.0	593543 ± 117	239081 ± 97	2742 ± 94	0.649 ± 0.016
Dost et al. (2020)	163 ± 15	70 ± 10	-100 ± 12	593295	238931	2953 ± 94	-
KNMI	-	-	-	594247	239268	3000	-

Table 1. Comparison between earthquake parameters estimated in this study with estimation from Dost et al. (2020) and the KNMI. Specifically for T_0 , the value is relative to the origin time estimated by the KNMI.




Figure 14. Hudson plot that shows the decomposition of the source mechanisms given the posterior distributions in Figure 11. The beachball with the red outlines represents the mean MT.




Figure 15. Top: The distributions of strike, dip, and rake solutions given the beachballs in Figure 14. Here we only show one part of the solutions closer to the orientation of the nearby major faults. Bottom: The marginal posterior distributions for different earthquake mechanisms given the decomposition in Figure 14.

7 Discussion and Conclusion

Using a probabilistic workflow incorporating the HMC algorithm, we estimate the source characteristics of a 3.4 M_L induced earthquake associated with gas extraction from the Groningen gas field. Specifically, we estimate the posterior probability density of ten earthquake parameters using two different sets of initial prior probabilities, of which one is used in conjunction with two sets of data: one consisting only of vertical component displacement recordings and a second one composed of the particle displacement in all three directions (east, north, down). We find that the posteriors estimated using both horizontal and vertical components of the seismograms (i.e., the latter data set) have similar shapes. At the same time, the one that only depends on the vertical component recordings yields a posterior that deviates (slightly) from the results of the other two cases while simultaneously being slightly broader. However, we find no substantial difference in the modeled seismograms associated with the different posterior means. In terms of runtime, using an 8-core MacBook Pro (2018 version), it took us a maximum of 3 minutes to run the 19 multi-stage workflows of the 1C-fault and 3C-fault case, and 12 minutes for the 121 multi-stage workflows of 3C-grid.

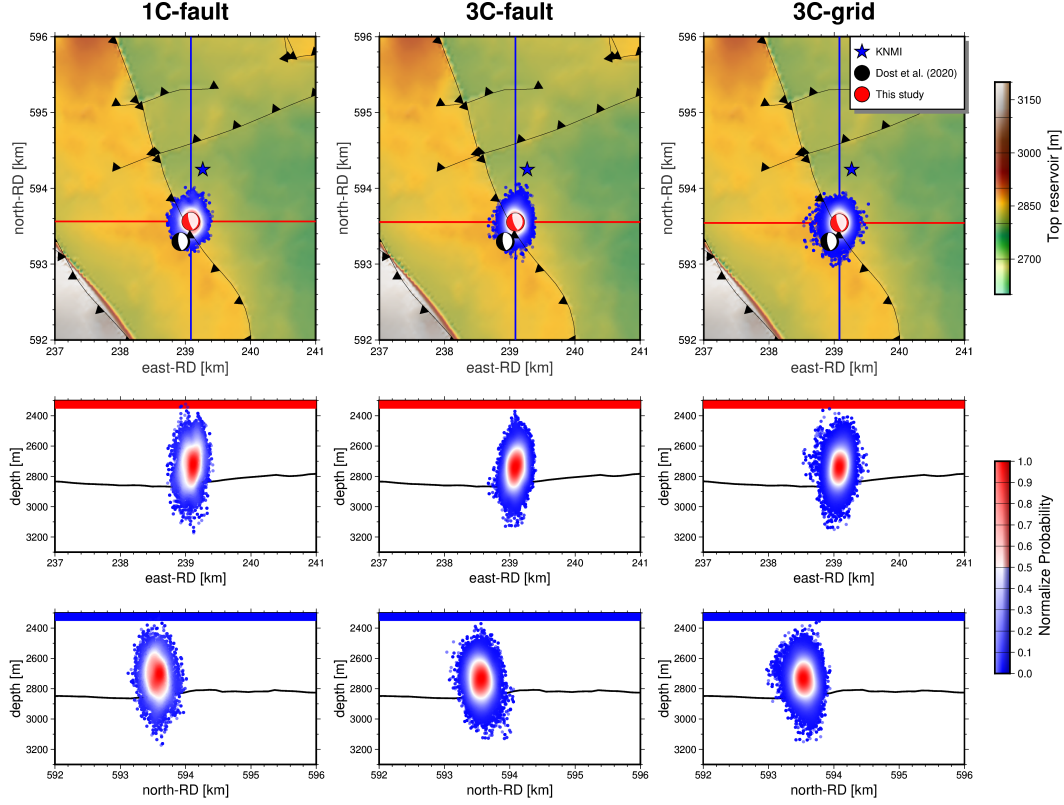


Figure 16. Comparison of samples used to generate centroid posterior distributions in Figure 11 (east, north, and depth) with the centroid estimated by Dost et al. (2020) and the KNMI. The samples are color coded with the density of centroid posteriors. The black line in the last two rows represent the top reservoir obtained from slicing the top reservoir map based on the red and blue line in the top row.

The main factor that affects the shape of the posteriors is uncertainty, which, in this case, is formulated as data and model uncertainty. In our study, we choose a uniform distribution for the model parameters to encode a state of ignorance (i.e., $\sigma_m \rightarrow \infty$). Whereas the data uncertainty is estimated individually for each component on each seismometer (and hence captured by σ_{ri} in A_{pq} , see Equation A4, where the indices r and i are associated with a specific receiver and component, respectively). It is assumed that the noise is uncorrelated. Prescribing the noise to be correlated will make the workflow more complex and computationally more costly and requires us to estimate data covariance matrices. In addition, a study by Gu et al. (2018) reveals that in the case of induced seismicity, accounting for (potentially) correlated noise has relatively little effect compared to the uncertainty arising from the inaccuracy of the velocity model. Ideally, the latter is also formally included. The relation between a specific source model (i.e., a specific set of model parameters) and the particle displacement at the surface will, in that case, be quantified by means of a probability density function (Tarantola & Valette, 1981). Due to limited computational resources, however, we disregard the uncertainty associated with the velocity model. Including it (for our 3D velocity model) will require enormous computational effort as each ‘cell’ in the model must be varied according to their variance when computing the forward problem represented by Equation 1 (effectively, the Green’s functions will become probability density functions). While using 1D velocity models, lateral heterogeneity is not considered, and therefore, the number of cells

will be exponentially reduced and hence the computational burden. In general, using 3D models has improved the characterization of earthquake sources since they better represent the subsurface compared to 1D models (Hingee et al., 2011; Hejrani et al., 2017; Wang & Zhan, 2020).

Many studies involving MT inversions limit the model space to purely double-couple sources. Often, this limitation is justified by (presumed) a priori information of the source type. For example, a DC mechanism is usually sufficient to explain faulting in tectonically active areas where volumetric components can be expected to be negligible. In the context of induced seismicity, however, numerous studies have found that non-DC components explain a substantive part of the observed recordings (Caputa et al., 2021; Cesca et al., 2013; Šílený & Milev, 2008). In the context of the gas extraction below Groningen, a study by Willacy et al. (2019) uses waveform data to obtain moment tensor solutions assuming that the earthquakes can be explained by DC mechanisms. Hence, they only estimated the best DC mechanisms of each observed earthquake. Meanwhile, another study by Kühn et al. (2020) (also focusing on the events in Groningen) reveals that ignoring non-DC components significantly affects the solution and data fit. In this study, we find the DC component to be dominant but still need the ISO and CLVD components to be non-zero in order to explain the data.

As for the centroid, we find that it is likely that the earthquake nucleated above the reservoir. In our case, the posterior mean is located a bit above 2.8 km depth. This is a small shift from the estimate by Dost et al. (2020), who estimate the earthquake to be located inside the reservoir. A recent study by Smith et al. (2020), however, finds that most of the Groningen earthquakes nucleated just above the reservoir, although this study does not include the event we are using here. Considering both the centroid and MT solution, we find that the models that best explain the recorded particle displacements correlate well with the nearby fault (see Figure 16).

For the workflow to be applied to a larger number of induced earthquakes, we believe a couple of additions would be beneficial. The first is related to the estimation of the data uncertainty. Since the workflow relies on Bayesian inference, the data uncertainty is rather critical while shaping reliable final posterior distributions. A second addition would be to allow for correlated noise. Particularly for the Groningen earthquakes, the effect of correlated noise for source characterization is not considered in any of the publications cited in this manuscript. Quantifying its effect on source parameters estimations would therefore be relevant.

Appendix A Linearization of the Forward Problem

In the context of Hamiltonian Monte Carlo, a model \mathbf{m} can be interpreted as the position of a particle in the $2N_m$ -dimensional phase space (Betancourt, 2017). Using equation 3, this particle’s potential energy U , which is defined as $U(\mathbf{m}) \equiv -\ln \rho(\mathbf{m} \mid \mathbf{d}^{\text{obs}})$, therefore reads (Fichtner & Simutè, 2018; Masfara et al., 2022)

$$U(\mathbf{m}) = \frac{1}{2} (\mathbf{d}(\mathbf{m}) - \mathbf{d}^{\text{obs}})^\top \mathbf{C}_d^{-1} (\mathbf{d}(\mathbf{m}) - \mathbf{d}^{\text{obs}}) - \frac{1}{2} (\mathbf{m} - \mathbf{m}^{(0)})^\top \mathbf{C}_m^{-1} (\mathbf{m} - \mathbf{m}^{(0)}). \quad (\text{A1})$$

where we have, for convenience, ignored the proportionality constant (this does not affect our results as it is independent of \mathbf{m} .)

In this study, the HMC variant that involves a linearization of the forward problems is referred to as ‘linearized HMC.’ Linearization of the forward model implies a Taylor expansion of equation 1 about the prior mean $\mathbf{m}^{(0)}$, and subsequently dropping higher

order terms. The linear approximation of u_i reads (Fichtner & Simut , 2018)

$$u_i(\mathbf{x}^{(r)}, t; \mathbf{m}) = u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) + \sum_{p=1}^{N_m} \frac{\partial}{\partial m_p} u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) (m_p - m_p^{(0)}). \quad (\text{A2})$$

Substituting this approximation in equation A1, $U(\mathbf{m})$ can be written as follows:

$$U(\mathbf{m}) = \frac{1}{2} \sum_{p,q=1}^{N_m} (m_p - m_p^{(0)}) A_{pq} (m_q - m_q^{(0)}) + \sum_{p=1}^{N_m} b_p (m_p - m_p^{(0)}) + \frac{1}{2} c, \quad (\text{A3})$$

where A_{pq} , b_p , and c read

$$A_{pq} \equiv \frac{1}{T} \sum_{r=1}^{N_r} \sum_{i=1}^3 \sigma_{ri}^{-2} \int_0^T \frac{\partial}{\partial m_p} u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) \frac{\partial}{\partial m_q} u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) dt + \frac{1}{N_m \sigma_m^2}, \quad (\text{A4})$$

$$b_p \equiv \frac{1}{T} \sum_{r=1}^{N_r} \sum_{i=1}^3 \sigma_{ri}^{-2} \int_0^T \left[u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) - u_i^{\text{obs}}(\mathbf{x}^{(r)}, t) \right] \frac{\partial}{\partial m_p} u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) dt, \quad (\text{A5})$$

and

$$c \equiv \frac{1}{T} \sum_{r=1}^{N_r} \sum_{i=1}^3 \sigma_{ri}^{-2} \int_0^T \left[u_i(\mathbf{x}^{(r)}, t; \mathbf{m}^{(0)}) - u_i^{\text{obs}}(\mathbf{x}^{(r)}, t) \right]^2 dt, \quad (\text{A6})$$

respectively. Here, σ_{ri}^{-2} encodes the data uncertainty for receiver number r and component i . This formulation implies that the noise is assumed to be uncorrelated. Similarly, a pairwise uncorrelated prior probability of the model parameters is considered. (i.e., \mathbf{C}_m is diagonal). Using the three quantities above, $\frac{\partial U}{\partial \mathbf{m}}$ in equation 7 can be replaced by

$$\frac{\partial U}{\partial m_p} = \sum_{q=1}^{N_m} A_{pq} (m_q - m_q^{(0)}) + b_p. \quad (\text{A7})$$

Acknowledgments

We thank NWO for providing the financial support for this study. We thank Thomas Cullison for his assistance in generating our synthetic datasets. We thank Daniela Kuhn for the fruitful discussion and suggestion regarding probabilistic inversions.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

The seismograms of the earthquake were downloaded from the KNMI seismic and acoustic network (KNMI, 1993) via Obspy (Beyreuther et al., 2010). The Forward modeled seismograms were generated using SPECFEM (Komatitsch & Tromp, 2002) available at <https://github.com/geodynamics/specfem3d>. The input for SPECFEM is Sub-surface Groningen models provided by Romijn (2017) available at <https://nam-onderzoeksrapporten>

.data-app.nl/reports/download/groningen/en/3b4f8b0d-0277-40e0-8ff5-9a385c08327d (last accessed July 2023). The models were first pre-processed using Python packages, gnam (Cullison et al., 2022), and PyAspect (Cullison & Masfara, 2022). Those packages are stored at <https://github.com/code-cullison/gnam> and <https://github.com/code-cullison/pyaspect>, respectively.

References

- Aki, K., & Richards, P. G. (2002). *Quantitative seismology* (2nd ed.). University Science Books. Hardcover. Retrieved from <http://www.worldcat.org/isbn/0935702962>
- Betancourt, M. (2017). A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*.
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010). Obspy: A python toolbox for seismology. *Seismological Research Letters*, 81(3), 530–533.
- Bommer, J. J., Stafford, P. J., Edwards, B., Dost, B., van Dedem, E., Rodriguez-Marek, A., ... Ntinalaxis, M. (2017). Framework for a ground-motion model for induced seismic hazard and risk analysis in the groningen gas field, the netherlands. *Earthquake Spectra*, 33(2), 481–498.
- Bourne, S., & Oates, S. (2017, October). Induced seismicity within the groningen gas field. Retrieved from <https://doi.org/10.5281/zenodo.1035226> doi: 10.5281/zenodo.1035226
- Caputa, A., Rudziński, L., & Cesca, S. (2021). How to assess the moment tensor inversion resolution for mining induced seismicity: A case study for the rudna mine, poland. *Frontiers in Earth Science*, 9, 671207.
- Cesca, S., Rohr, A., & Dahm, T. (2013). Discrimination of induced seismicity by full moment tensor inversion and decomposition. *Journal of seismology*, 17, 147–163.
- Cullison, T., & Masfara, L. O. M. (2022, August). *Pyaspect: First release*. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.6987368> doi: 10.5281/zenodo.6987368
- Cullison, T., Masfara, L. O. M., & Hawkins, R. (2022, August). *gnam: First release*. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.6987375> doi: 10.5281/zenodo.6987375
- Dais, D., Bal, I. E., & Smyrou, E. (2018). "differences" between induced and natural seismic events. In *16th european conference on earthquake engineering*.
- Dost, B., Goutbeek, F., van Eck, T., & Kraaijpoel, D. (2012). Monitoring induced seismicity in the north of the netherlands: status report 2010. *KNMI Scientific report*, 2003.
- Dost, B., van Stiphout, A., Kühn, D., Kortekaas, M., Ruigrok, E., & Heimann, S. (2020). Probabilistic moment tensor inversion for hydrocarbon-induced seismicity in the groningen gas field, the netherlands, part 2: Application. *Bulletin of the Seismological Society of America*, 110(5), 2112–2123.
- Ellsworth, W. L., Llenos, A. L., McGarr, A. F., Michael, A. J., Rubinstein, J. L., Mueller, C. S., ... Calais, E. (2015). Increasing seismicity in the us midcontinent: Implications for earthquake hazard. *The Leading Edge*, 34(6), 618–626.
- Fichtner, A., & Simuté, S. (2018). Hamiltonian monte carlo inversion of seismic sources in complex media. *Journal of Geophysical Research: Solid Earth*, 123(4), 2984–2999.
- Fichtner, A., Zunino, A., & Gebraad, L. (2019). Hamiltonian monte carlo solution of tomographic inverse problems. *Geophysical Journal International*, 216(2), 1344–1363.
- Fichtner, A., Zunino, A., Gebraad, L., & Boehm, C. (2021). Autotuning hamiltonian monte carlo for efficient generalized nullspace exploration. *Geophysical Journal*

- International*, 227(2), 941–968.
- Foulger, G. R., Wilson, M. P., Gluyas, J. G., Julian, B. R., & Davies, R. J. (2018). Global review of human-induced earthquakes. *Earth-Science Reviews*, 178, 438–514.
- Gu, C., Marzouk, Y. M., & Toksöz, M. N. (2018). Waveform-based bayesian full moment tensor inversion and uncertainty determination for the induced seismicity in an oil/gas field. *Geophysical Journal International*, 212(3), 1963–1985.
- Hejrani, B., Tkalčić, H., & Fichtner, A. (2017). Centroid moment tensor catalogue using a 3-d continental scale earth model: Application to earthquakes in papua new guinea and the solomon islands. *Journal of Geophysical Research: Solid Earth*, 122(7), 5517–5543.
- Hingee, M., Tkalčić, H., Fichtner, A., & Sambridge, M. (2011). Seismic moment tensor inversion using a 3-d structural model: applications for the australian region. *Geophysical Journal International*, 184(2), 949–964.
- Jost, M. u., & Herrmann, R. (1989). A student’s guide to and review of moment tensors. *Seismological Research Letters*, 60(2), 37–57.
- KNMI. (1993). *Netherlands Seismic and Acoustic Network*. Royal Netherlands Meteorological Institute (KNMI), Other/Seismic Network. (10.21944/e970fd34-23b9-3411-b366-e4f72877d2c5)
- Komatitsch, D., & Tromp, J. (2002). Spectral-element simulations of global seismic wave propagation—i. validation. *Geophysical Journal International*, 149(2), 390–412.
- Kühn, D., Heimann, S., Isken, M. P., Ruigrok, E., & Dost, B. (2020). Probabilistic moment tensor inversion for hydrocarbon-induced seismicity in the groningen gas field, the netherlands, part 1: Testing. *Bulletin of the Seismological Society of America*, 110(5), 2095–2111.
- Lui, S. K., Helmberger, D., Yu, J., & Wei, S. (2016). Rapid assessment of earthquake source characteristics. *Bulletin of the Seismological Society of America*, 106(6), 2490–2499.
- Masfara, L. O. M., Cullison, T., & Weemstra, C. (2022). An efficient probabilistic workflow for estimating induced earthquake parameters in 3d heterogeneous media. *Solid Earth*, 13(8), 1309–1325.
- Mustać, M., & Tkalčić, H. (2016). Point source moment tensor inversion through a bayesian hierarchical model. *Geophysical Journal International*, 204(1), 311–323.
- Neal, R. M., et al. (2011). Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11), 2.
- Nepveu, M., van Thienen-Visser, K., & Sijacic, D. (2016). Statistics of seismic events at the groningen field. *Bulletin of Earthquake Engineering*, 14, 3343–3362.
- Ntinalexis, M., Bommer, J. J., Ruigrok, E., Edwards, B., Pinho, R., Dost, B., . . . van Elk, J. (2019). Ground-motion networks in the groningen field: usability and consistency of surface recordings. *Journal of Seismology*, 23(6), 1233–1253.
- Peng, C., Jiang, P., Ma, Q., Wu, P., Su, J., Zheng, Y., & Yang, J. (2021). Performance evaluation of an earthquake early warning system in the 2019–2020 m 6.0 changning, sichuan, china, seismic sequence. *Frontiers in Earth Science*, 9, 699941.
- Pickering, M. (2015). An estimate of the earthquake hypocenter locations in the groningen gas field. *Nederlandse Aardolie Maatschappij. Assen*.
- Romijn, R. (2017). Groningen velocity model 2017—groningen full elastic velocity model september 2017. *Technical Rept..*
- Ruigrok, E., Domingo-Ballesta, J., van den Hazel, G.-J., Dost, B., & Evers, L. (2019). Groningen explosion database. *First Break*, 37(8), 37–41.
- Ruigrok, E., & Dost, B. (2019). Seismic monitoring and site-characterization with near-surface vertical arrays. In *Near Surface Geoscience Conference and Exhi-*

- bition (pp. 1–5). doi: 10.3997/2214-4609.201902455
- Sarhosis, V., Dais, D., Smyrou, E., & Bal, İ. E. (2019). Evaluation of modelling strategies for estimating cumulative damage on groningen masonry buildings due to recursive induced earthquakes. *Bulletin of Earthquake Engineering*, 17(8), 4689–4710.
- Šilený, J., & Milev, A. (2008). Source mechanism of mining induced seismic events—resolution of double couple and non double couple models. *Tectonophysics*, 456(1-2), 3–15.
- Simutè, S., Boehm, C., Krischer, L., Gokhberg, A., Vallée, M., & Fichtner, A. (2022). Bayesian seismic source inversion with a 3-d earth model of the japanese islands. *Journal of Geophysical Research: Solid Earth*, e2022JB024231.
- Smith, J. D., White, R. S., Avouac, J.-P., & Bourne, S. (2020). Probabilistic earthquake locations of induced seismicity in the groningen region, the netherlands. *Geophysical Journal International*, 222(1), 507–516.
- Spetzler, J., & Dost, B. (2017). Hypocentre estimation of induced earthquakes in groningen. *Geophysical Journal International*, 209(1), 453–465.
- Tarantola, A. (2005). *Inverse Problem Theory and methods for model parameter estimation*. Philadelphia, PA, USA: Society of industrial and applied mathematics.
- Tarantola, A. (2006). Popper, bayes and the inverse problem. *Nature physics*, 2(8), 492–494.
- Tarantola, A., & Valette, B. (1981, October). Inverse problems = Quest for information. *Journal of Geophysics*, 50(1), 159–170. Retrieved from <https://journal.geophysicsjournal.com/JofG/article/view/28>
- Trippetta, F., Petricca, P., Billi, A., Collettini, C., Cuffaro, M., Lombardi, A. M., ... Doglioni, C. (2019). From mapped faults to fault-length earthquake magnitude (flem): a test on italy with methodological implications. *Solid Earth*, 10(5), 1555–1579.
- van Thienen-Visser, K., & Breunese, J. (2015). Induced seismicity of the groningen gas field: History and recent developments. *The Leading Edge*, 34(6), 664–671.
- Vlek, C. (2018). Induced earthquakes from long-term gas extraction in groningen, the netherlands: statistical analysis and prognosis for acceptable-risk regulation. *Risk analysis*, 38(7), 1455–1473.
- Wang, X., & Zhan, Z. (2020). Moving from 1-d to 3-d velocity model: automated waveform-based earthquake moment tensor inversion in the los angeles region. *Geophysical Journal International*, 220(1), 218–234.
- Willacy, C., van Dedem, E., Minisini, S., Li, J., Blokland, J. W., Das, I., & Droujinine, A. (2018). Application of full-waveform event location and moment-tensor inversion for groningen induced seismicity. *The Leading Edge*, 37(2), 92–99.
- Willacy, C., van Dedem, E., Minisini, S., Li, J., Blokland, J.-W., Das, I., & Droujinine, A. (2019). Full-waveform event location and moment tensor inversion for induced seismicity. *Geophysics*, 84(2), KS39–KS57.

Figure 1.

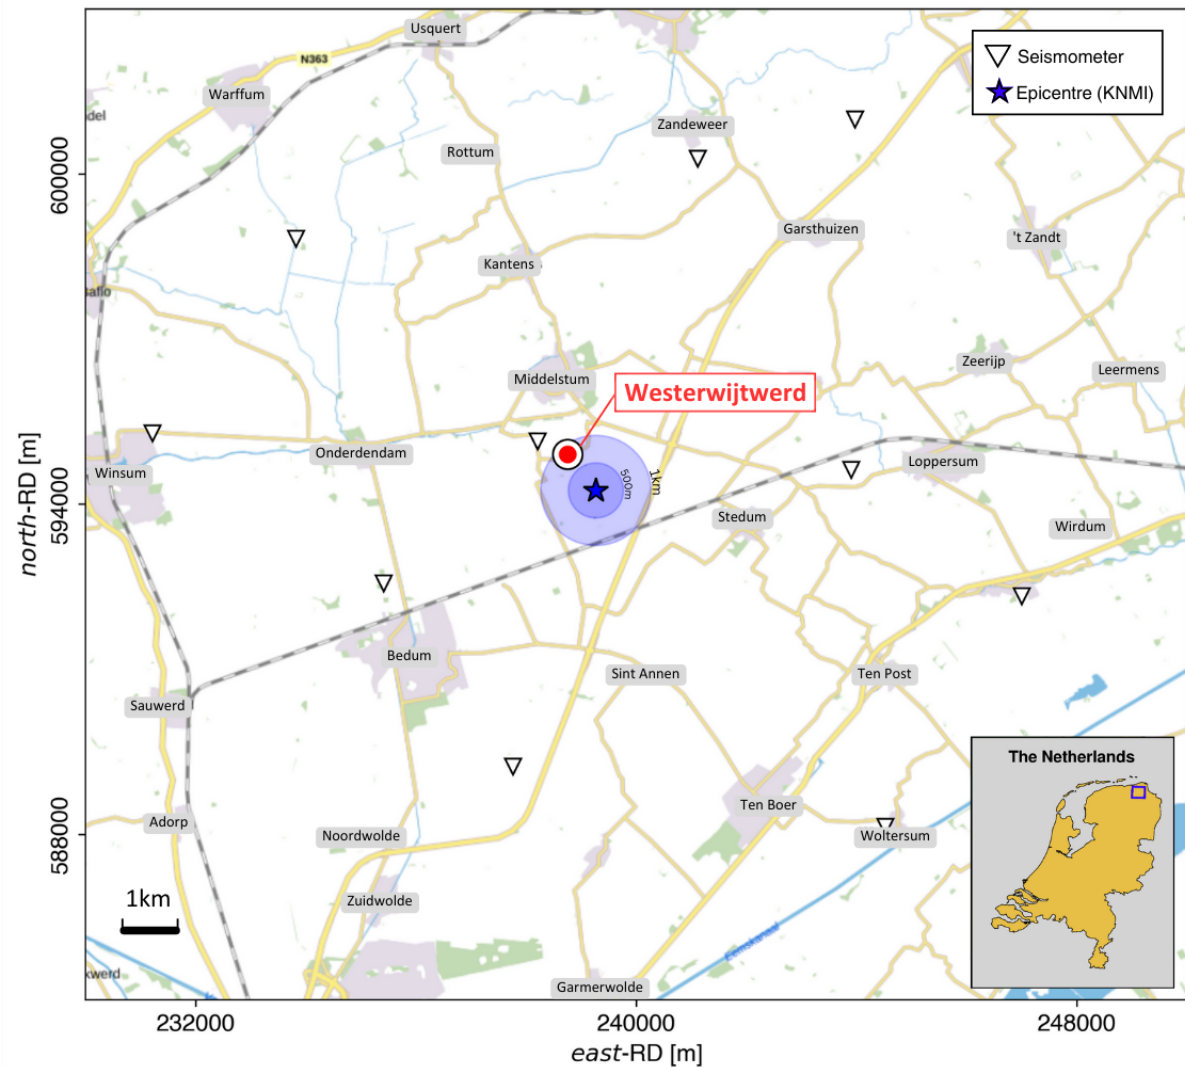

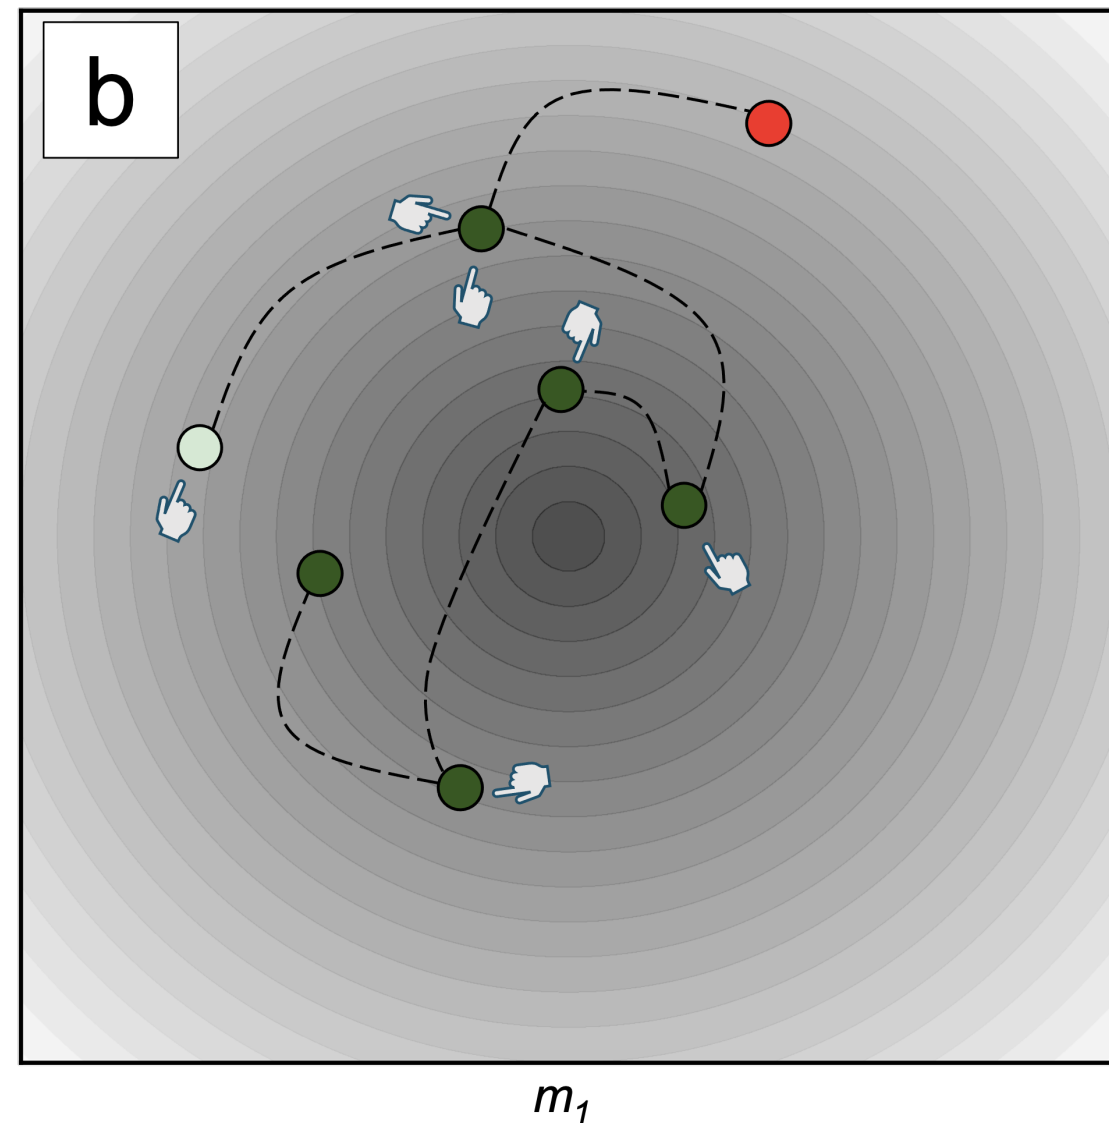


Figure 2.





Min Posterior Probability Max



Max Potential Energy Min



HMC Momentum Starting model Rejected sample Accepted sample — MH trajectory --- HMC trajectory

Figure 3.



 HMC momentum
  Linearized HMC trajectory
  Localized posterior

Figure 4.



HMC momentum



Linearized HMC trajectory



Localized posterior

Figure 5.


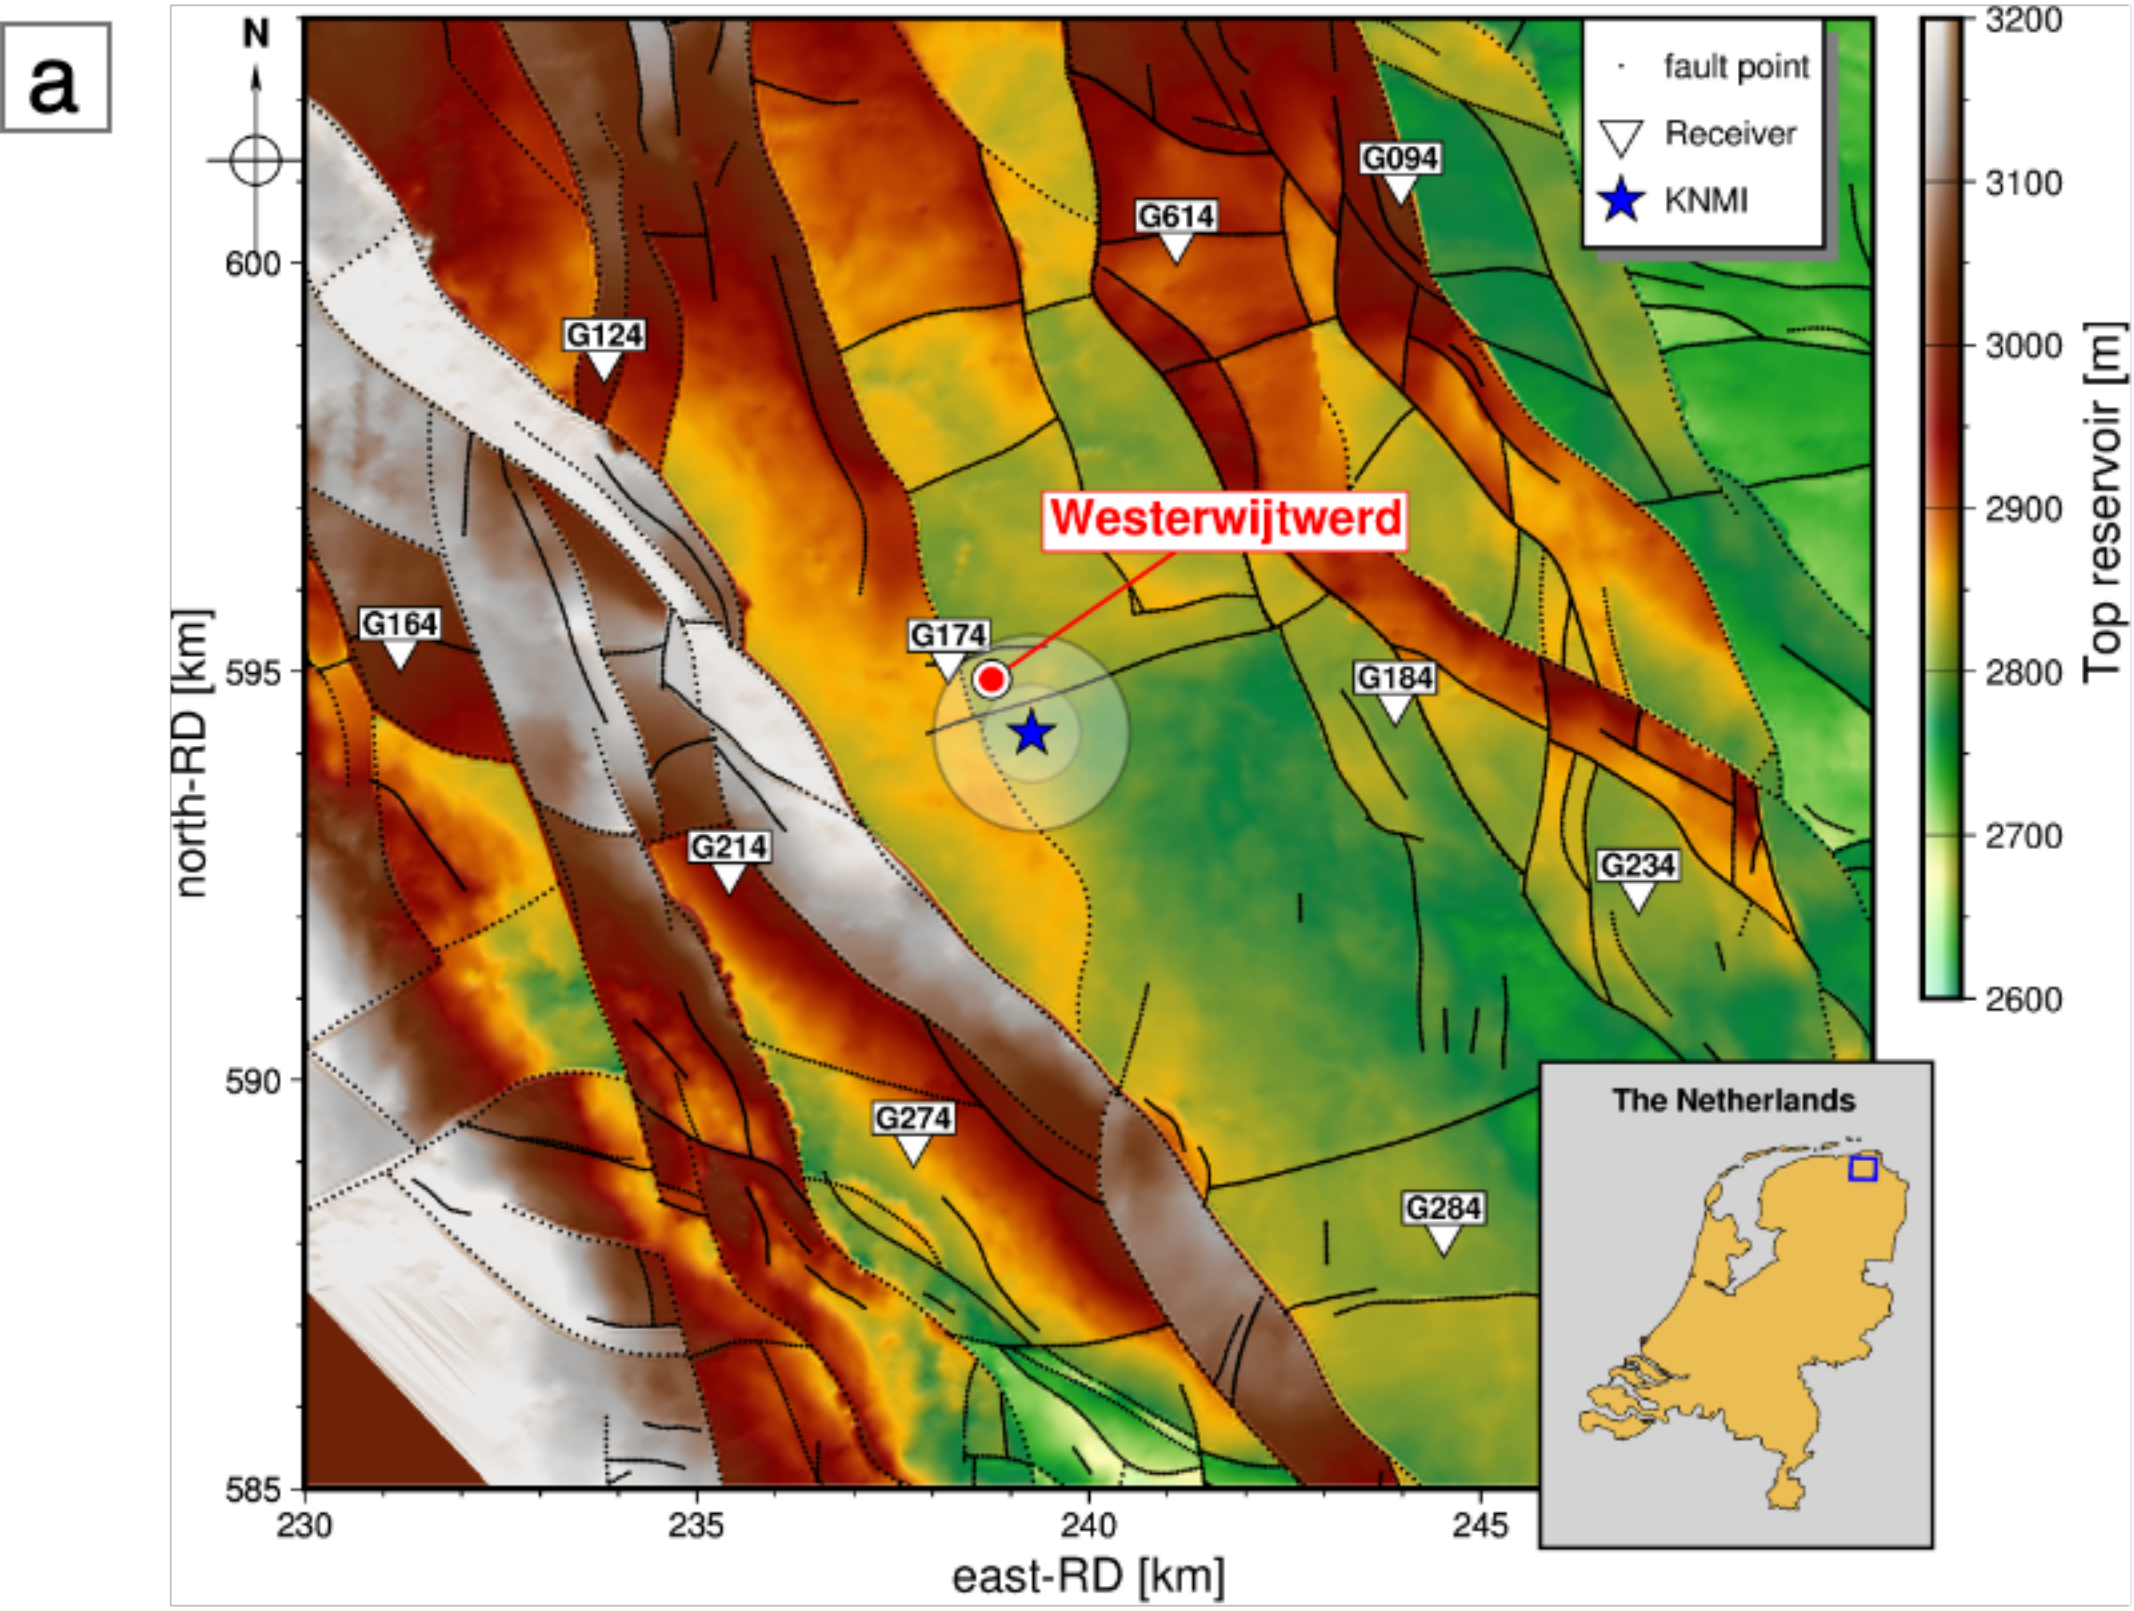


Figure 6.




Figure 7.

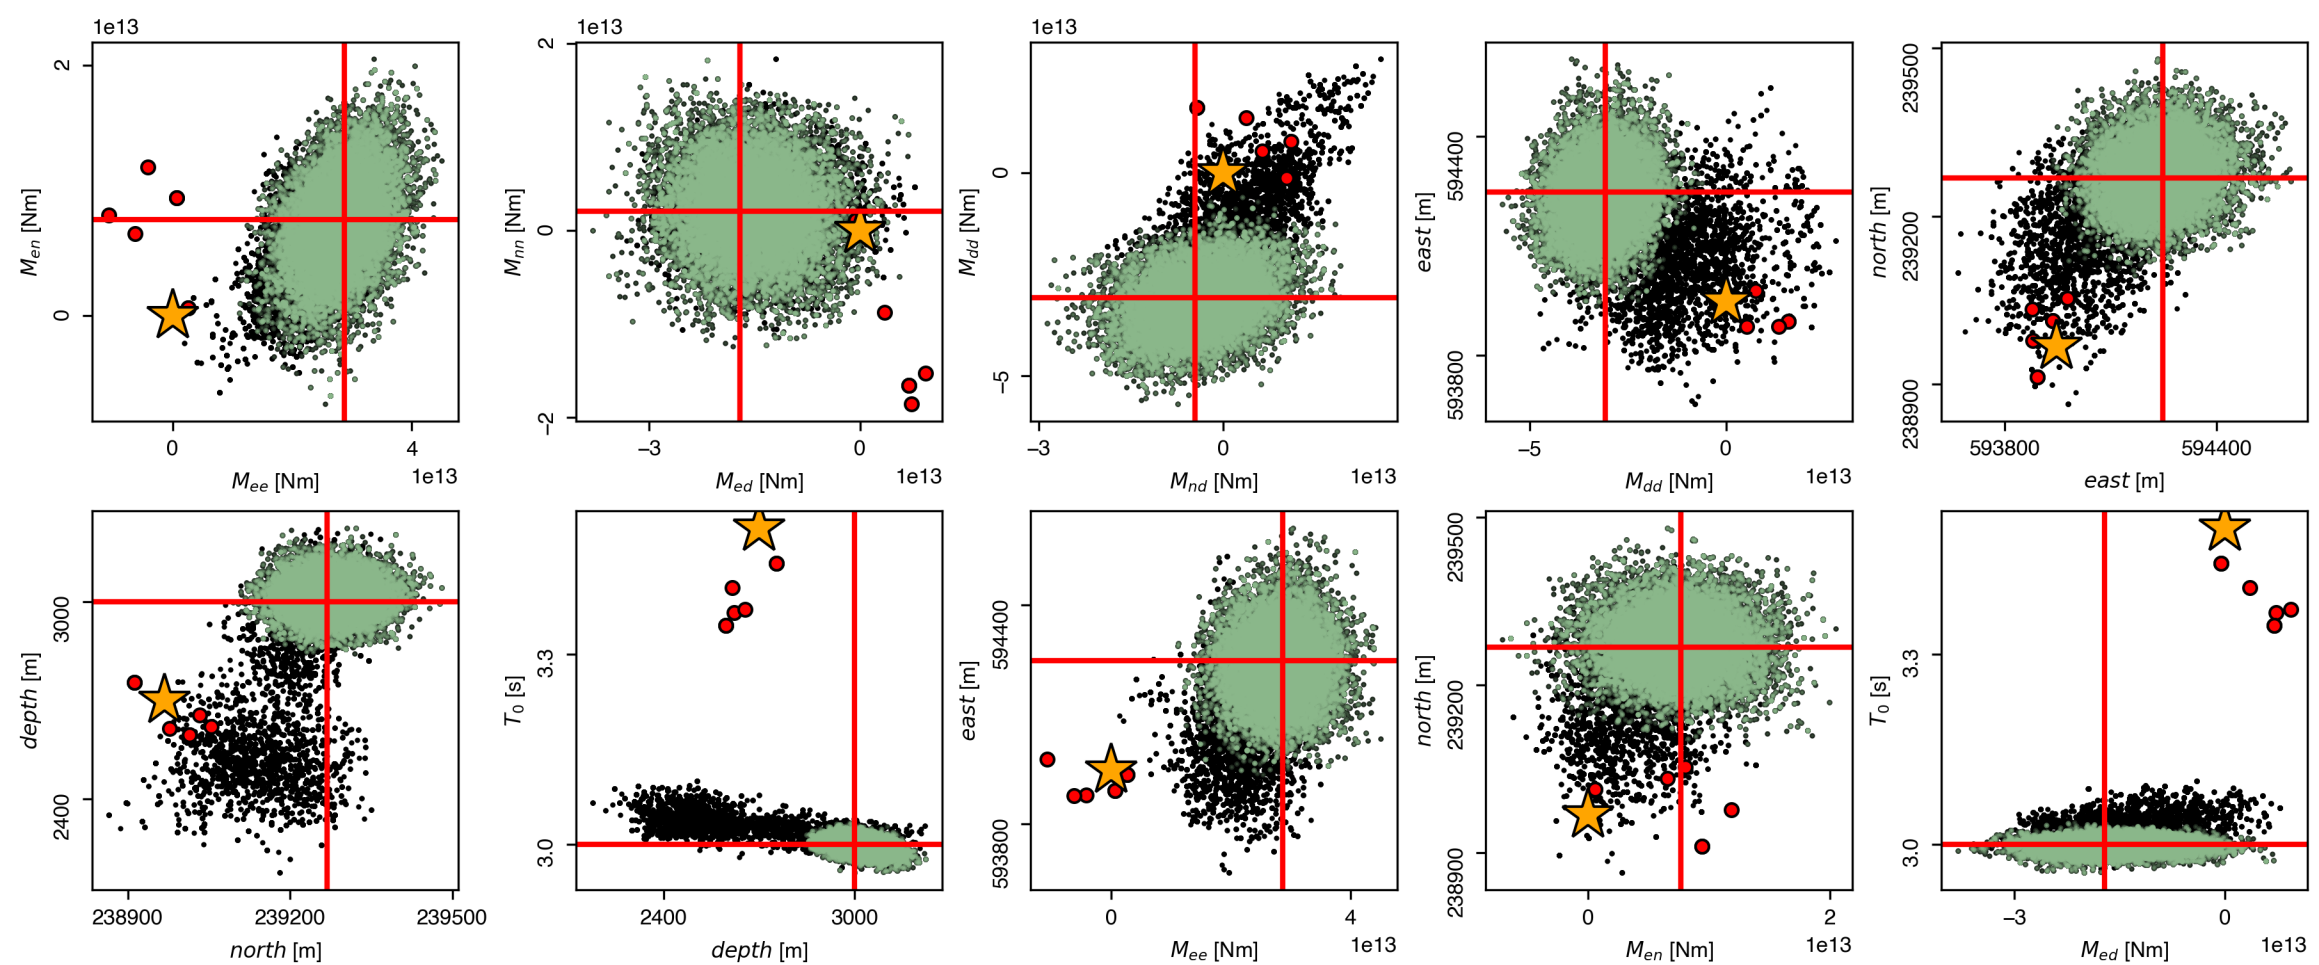



Figure 8.



modeled "observed data+noise" modeled "observed data" modeled using posterior mean

Figure 9.

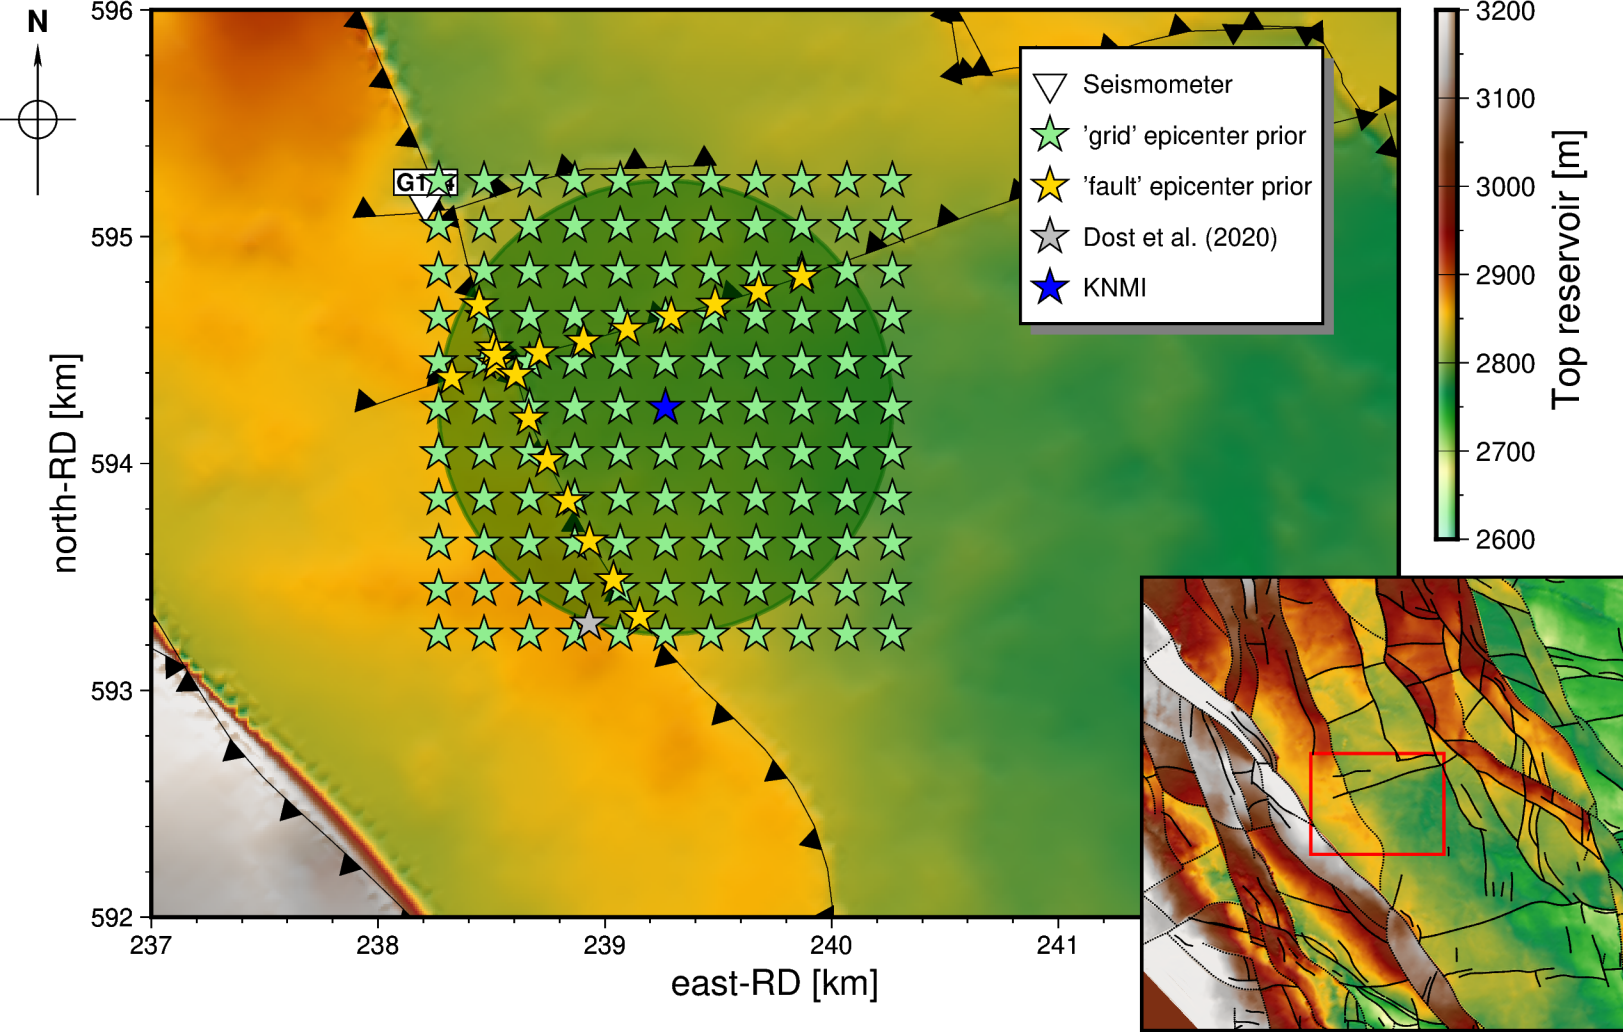


Figure 10.







Figure 11.

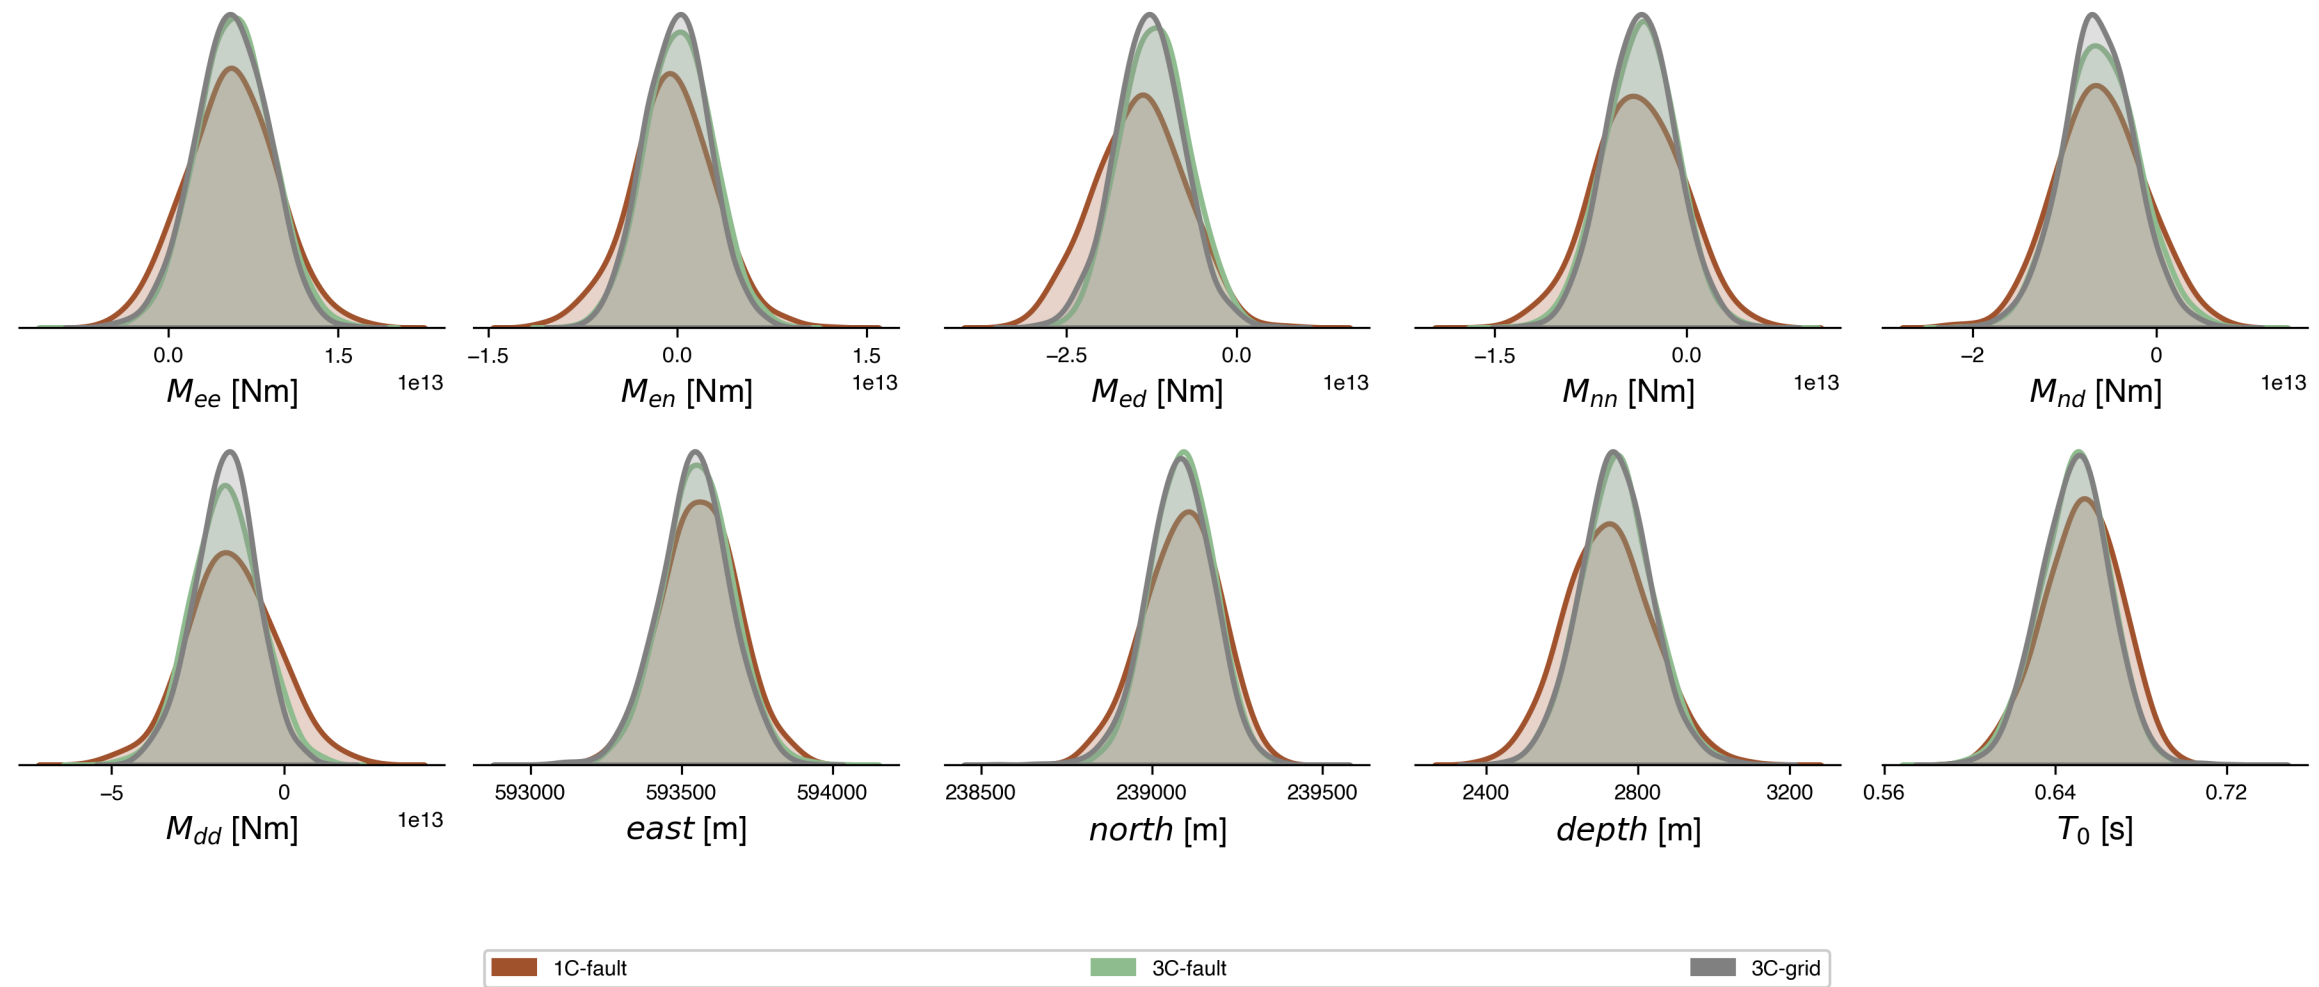


Figure 12.

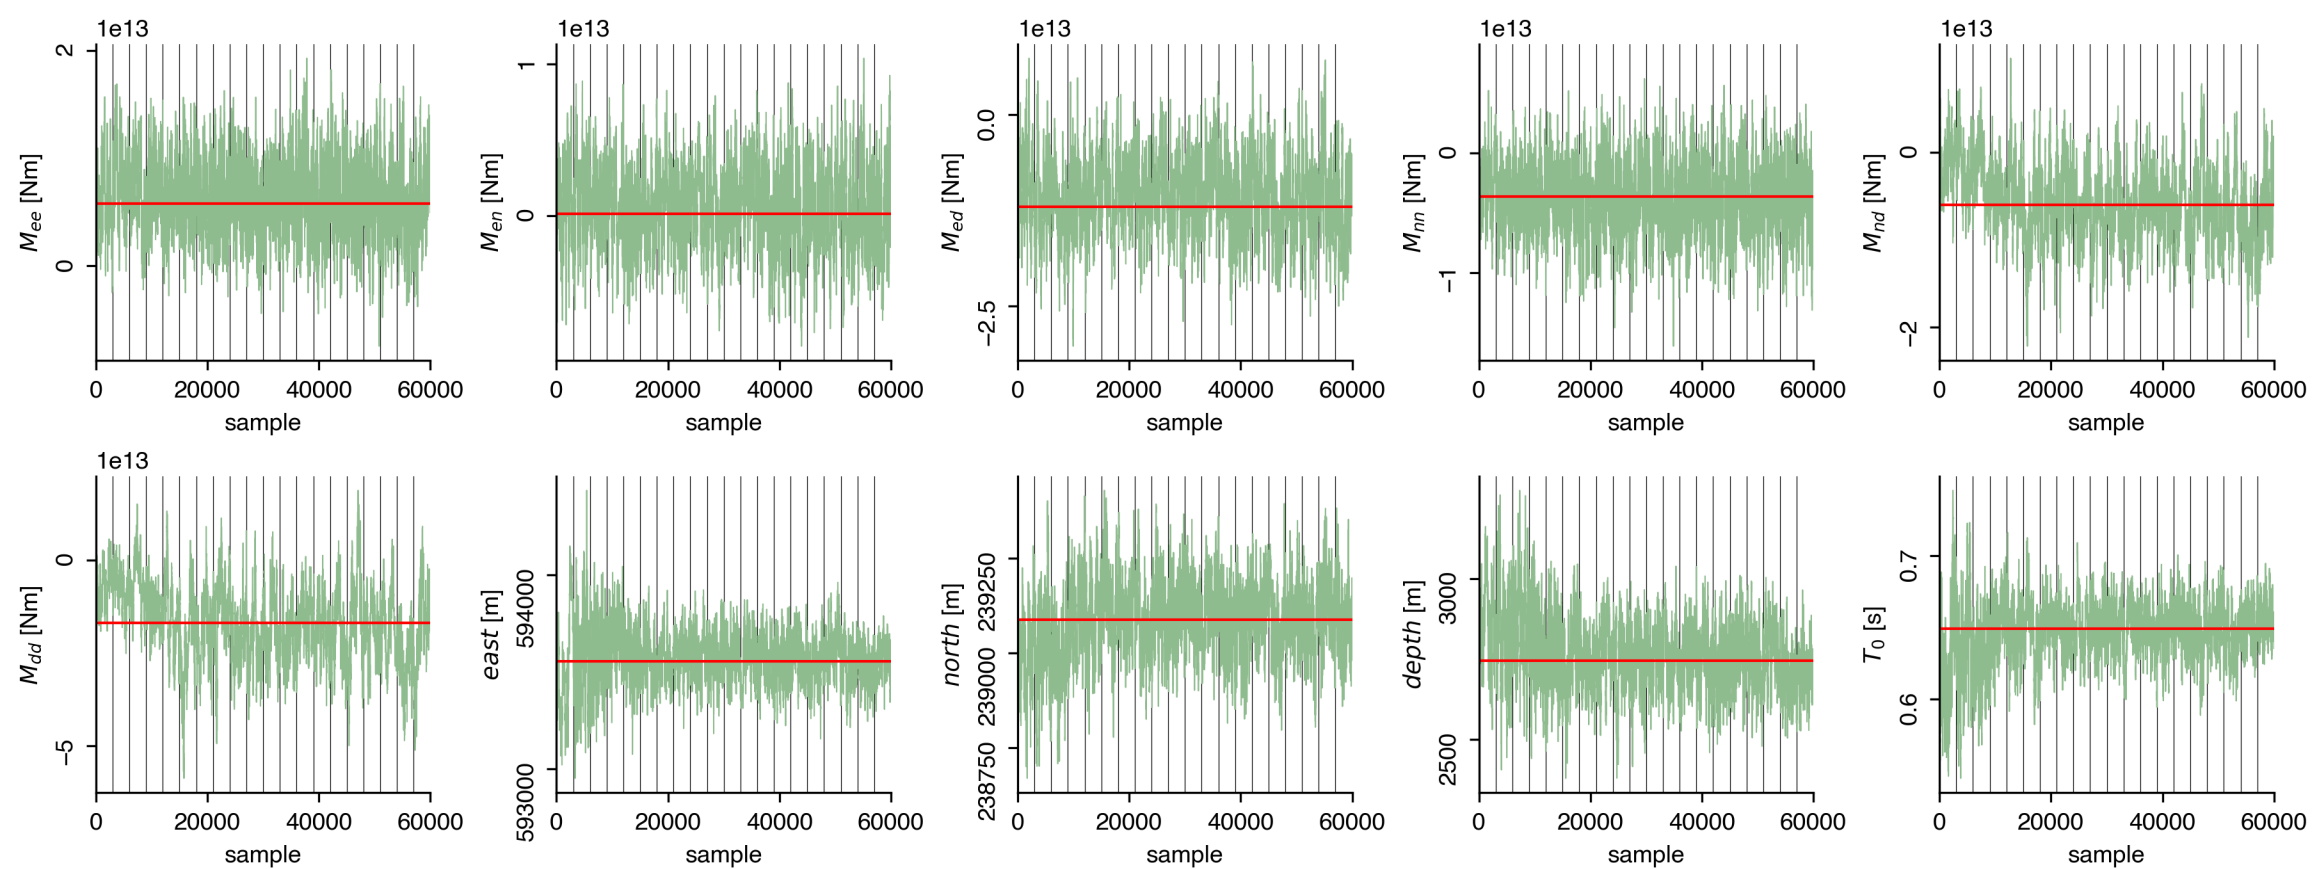


Figure 13.






Figure 14.

1C-fault



3C-fault



3C-grid




Figure 15.






Figure 16.

1C-fault



3C-fault



3C-grid

