SPECIAL ISSUE

Solving Geophysical Inversion Problems with Intractable Likelihoods: Linearized Gaussian Approximations Versus the Correlated Pseudo-marginal Method

Lea Friedli[1] · **Niklas Linde[1]**

Abstract

A geophysical Bayesian inversion problem may target the posterior distribution of geological or hydrogeological parameters given geophysical data. To account for the scatter in the petrophysical relationship linking the target parameters to the geophysical properties, this study treats the intermediate geophysical properties as latent (unobservable) variables. To perform inversion in such a latent variable model, the intractable likelihood function of the (hydro)geological parameters given the geophysical data needs to be estimated. This can be achieved by approximation with a Gaussian probability density function based on local linearization of the geophysical forward operator, thereby, accounting for the noise in the petrophysical relationship by a corresponding addition to the data covariance matrix. The new approximate method is compared against the general correlated pseudo-marginal method, which estimates the likelihood by Monte Carlo averaging over samples of the latent variable. First, the performances of the two methods are tested on a synthetic test example, in which a multivariate Gaussian porosity field is inferred using crosshole ground-penetrating radar first-arrival travel times. For this example with rather small petrophysical uncertainty, the two methods provide near-identical estimates, while an inversion that ignores petrophysical uncertainty leads to biased estimates. The results of a sensitivity analysis are then used to suggest that the linearized Gaussian approach, while attractive due to its relative computational speed, suffers from a decreasing accuracy with increasing scatter in the petrophysical relationship. The computationally more expensive correlated pseudo-marginal method performs very well even for settings with high petrophysical uncertainty.

Keywords Inverse theory · Intractable likelihood · Latent variable model · Hydrogeophysics

---

✉ Lea Friedli
  lea.friedli@unil.ch

[1] Institute of Earth Sciences, University of Lausanne, Lausanne, Switzerland

1 Introduction

This work targets a Bayesian inverse problem in which the posterior distribution of target geological or hydrogeological parameters $\boldsymbol{\theta}$ are inferred from geophysical data $\mathbf{y}$. Petrophysical relationships linking (hydro)geological variables (e.g., permeability, clay fraction, salinity) to geophysical properties (e.g., dielectric permittivity, electrical conductivity, magnetic susceptibility) must then be introduced (e.g., Hinnell et al. 2010; Kowalsky et al. 2005). Such relationships are often inherently uncertain (e.g., Mavko et al. 2009), however, in most hydrogeophysical inversion studies targeting hydrogeological properties, the predictive power of the petrophysical relationship is assumed to be perfect provided that the right parameter values are used (e.g., Lochbühler et al. 2014; Kowalsky et al. 2005). Brunetti and Linde (2017) show that this assumption may lead to bias, too narrow uncertainty bounds and overly variable parameter estimates.

Brunetti and Linde (2017) distinguish three sources of uncertainty in the petrophysical relationship: model, parameter and prediction uncertainty. While the first two result from uncertainty related to the choice of the petrophysical model and its parameter values, the latter arises from scatter and bias around the calibrated model. As in Brunetti and Linde (2017), only petrophysical prediction uncertainty is considered here, using a latent variable model formulation which expresses the geophysical properties as $\mathbf{X} = \mathcal{F}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}_P$, with $\mathcal{F}(\cdot)$ being the petrophysical relationship and $\boldsymbol{\varepsilon}_P$ the petrophysical prediction error (PPE). The inclusion of the random effect of the PPE in the latent variable $\mathbf{X}$ makes the likelihood function $p(\mathbf{y}|\boldsymbol{\theta})$ intractable. In this study, two alternative methods are investigated to approximate this likelihood function in a Metropolis–Hastings algorithm (MH algorithm; Metropolis et al. 1953; Hastings 1970). The first approach is a Gaussian approximation based on local linearization of the geophysical forward operator. Thereby, the effect of the noise in the petrophysical relationship is included by a corresponding addition to the data covariance matrix. This new approach which is similar to the so-called delta method (Van der Vaart 2000) was suggested by Linde et al. (2017), but it has remained untested to date. This approximate method is compared against the correlated pseudo-marginal (CPM) method of Deligiannidis et al. (2018), which is based on the pseudo-marginal (Beaumont 2003; Andrieu and Roberts 2009,PM) method using Monte Carlo sampling of the latent variable to estimate the likelihood. Friedli et al. (2022) introduced and adapted the CPM method to a geophysical setting and demonstrated that in data-rich geophysical settings with low noise levels, it is essential to both use a well-working importance sampling strategy for the draws of latent variables and to correlate the latent samples used in the proposed and current states of the Markov chain.

In Friedli et al. (2022), the CPM method is compared to the original formulation of lithological tomography (Bosch 1999) and the so-called full inversion approach of Brunetti and Linde (2017). This latter method avoids intractable likelihood functions by targeting the joint posterior PDF $(\boldsymbol{\theta}, \mathbf{x}) \mapsto p(\boldsymbol{\theta}, \mathbf{x}|\mathbf{y})$ of the hydrogeological and geophysical parameters. Within the original lithological tomography method, first the target variable is sampled using the proposal scheme of the MH and second, one realization of the latent variable is drawn with conditional sampling. This actually represents a simplified form of the PM method with only one latent variable sample

and without importance sampling. The original form of lithological tomography leads to high variability in the estimate of the likelihood function, with the consequence of the algorithm often being highly inefficient (Brunetti and Linde 2017). Within the full inversion of Brunetti and Linde (2017), the latent variables are treated as additional target variables and an MH proposal scheme is used to draw new realizations of both. Friedli et al. (2022) show that this approach becomes inefficient with increasing dimensionality of the target and latent space and suffers from strong (posterior) correlations between the target and latent variables. Friedli et al. (2022) present a comparison of CPM with the original lithological tomography and full inversion approaches in a weakly non-linear setting showing that the CPM method outperforms the others by greatly enhancing the posterior exploration. While the CPM method already has been tested for geophysical inversion problems, the linearized Gaussian approach has not been applied to far and in this present study, the focus is on comparing this approximate approach against the general CPM method.

As a synthetic test case, a similar setting as in Friedli et al. (2022) is considered and multi-Gaussian porosity fields are inferred using crosshole ground-penetrating radar (GPR) first-arrival travel times. As in Friedli et al. (2022), a high-dimensional parameterization of the target porosity field is used. Subsequently, a sensitivity analysis is made to explore the performances of the linearized Gaussian approach and the CPM method as a function of increasing petrophysical prediction uncertainty. To avoid the challenges of a very-high dimensional target space in this sensitivity analysis, the complexity of the porosity field is reduced and it is assumed to be layered.

This contribution is a natural extension of the study by Friedli et al. (2022), which only considered the CPM method. While the fundamental concepts of the considered problem and the CPM method are repeated, the introduction and assessment of the linearized Gaussian approach is completely new. The manuscript is structured as follows: Sect. 2 gives a methodological overview of the considered latent variable model, Bayesian inference with intractable likelihoods, the linearized Gaussian approximation approach, the CPM method and the performance assessment metrics. Section 3 presents the results of our synthetic case study inferring multi-Gaussian porosity fields. In Sect. 4, the sensitivity analysis is presented. Finally, the study finishes with a discussion and conclusions in Sects. 5 and 6.

2 Methodology

2.1 Latent Variable Model

In the considered setting, the data vector $\mathbf{y} = (y_1, y_2, \ldots, y_T)$ (geophysical data) is given by,

$$\mathbf{Y} = \mathcal{G}(\mathbf{X}) + \boldsymbol{\varepsilon}_O = \mathcal{G}(\mathcal{F}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}_P) + \boldsymbol{\varepsilon}_O, \tag{1}$$

with $\mathbf{X} = (X_1, X_2, \ldots, X_L)$ denoting the latent variable (geophysical property) and $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots \theta_d)$ the target parameters (hydrogeological parameters). The variable $\mathbf{X}$ is referred to as a latent because it includes the PPE $\boldsymbol{\varepsilon}_P$, which is unobservable (latent)

but affects the observations. $\mathcal{G} : \mathbb{R}^L \rightarrow \mathbb{R}^T$ with $\boldsymbol{\varepsilon}_O$ and $\mathcal{F} : \mathbb{R}^d \rightarrow \mathbb{R}^L$ with $\boldsymbol{\varepsilon}_P$ refer to the physical forward solver with the observational noise and the petrophysical relationship with the PPE, respectively. In what follows, random variables and random vectors are referred to with upper-case letters and realizations thereof with lower-case letters. Assuming Gaussian errors, it holds,

$$p(\mathbf{x}|\boldsymbol{\theta}) = \varphi_L(\mathbf{x}; \mathcal{F}(\boldsymbol{\theta}), \boldsymbol{\Sigma}_P), \quad p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x}) = \varphi_T(\mathbf{y}; \mathcal{G}(\mathbf{x}), \boldsymbol{\Sigma}_\mathbf{Y}), \tag{2}$$

with $\varphi_M(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denoting the PDF of a M-variate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. In the test example, the target parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots \theta_d)$ describe a Gaussian random field parameterized on a grid of size $D \times D(d = D^2)$. It holds,

$$p(\boldsymbol{\theta}) = \varphi_{D^2}(\boldsymbol{\theta}; \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta), \tag{3}$$

and it is assumed that the mean $\boldsymbol{\mu}_\theta$ and the covariance matrix $\boldsymbol{\Sigma}_\theta$ of the target field are known.

2.2 Bayesian Inference and Intractable Likelihoods

Bayesian inversion problems target the posterior probability density function (PDF) $p(\boldsymbol{\theta}|\mathbf{y})$ of the model parameters $\boldsymbol{\theta}$ given the measurements $\mathbf{y}$. In Bayes' theorem, this posterior PDF is given by,

$$p(\boldsymbol{\theta}|\mathbf{y}) = \frac{p(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})}{p(\mathbf{y})}, \tag{4}$$

with the prior PDF $p(\boldsymbol{\theta})$ of the model parameters, the likelihood function $p(\mathbf{y}|\boldsymbol{\theta})$ and the evidence $p(\mathbf{y})$. As it is not possible to sample directly from the posterior, we use the Metropolis–Hastings algorithm (MH algorithm; Metropolis et al. 1953; Hastings 1970). At iteration j, the MH algorithm proposes a new model realization using the model proposal density $q(\cdot|\boldsymbol{\theta}^{(j-1)})$, which is then accepted or rejected based on the acceptance probability,

$$\alpha_{MH}\left(\boldsymbol{\theta}^{(j-1)}, \boldsymbol{\theta}^{(j)}\right) = \min\left\{1, \frac{q(\boldsymbol{\theta}^{(j-1)}|\boldsymbol{\theta}^{(j)})p(\boldsymbol{\theta}^{(j)})p(\mathbf{y}|\boldsymbol{\theta}^{(j)})}{q(\boldsymbol{\theta}^{(j)}|\boldsymbol{\theta}^{(j-1)})p(\boldsymbol{\theta}^{(j-1)})p(\mathbf{y}|\boldsymbol{\theta}^{(j-1)})}\right\}. \tag{5}$$

To implement the MH algorithm, the likelihood function $\boldsymbol{\theta} \mapsto p(\mathbf{y}|\boldsymbol{\theta})$ has to be evaluated,

$$p(\mathbf{y}|\boldsymbol{\theta}) = \int p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x})p(\mathbf{x}|\boldsymbol{\theta})d\mathbf{x}. \tag{6}$$

In a latent variable model, this integral has generally no analytical form, leading to an intractable likelihood function.

2.2.1 Proposal Scheme

When applying the MH algorithm to generate posterior samples, it is essential to choose a well-working proposal density $q(\cdot|\boldsymbol{\theta}^{(j-1)})$. Cotter et al. (2013) showed that standard random walk MCMC algorithms entail highly inefficient performance and strong dependence on the discretization when targeting high-dimensional Gaussian random fields. As a solution, they suggest proposal schemes that preserve the prior PDF, resulting in an MH algorithm for which the acceptance ratio only depends on the likelihoods. In geophysics, this proposal scheme is known as the extended Metropolis algorithm (Mosegaard and Tarantola 1995). If the target space is high-dimensional, the prior-preserving proposal scheme still needs to be chosen carefully (Ruggeri et al. 2015). Therefore, Friedli et al. (2022) introduce a prior-preserving version of the adaptive multi-chain algorithm DREAM$_{(ZS)}$ (DiffeRential Evolution Adaptive Metropolis using an archive of past states; Laloy and Vrugt 2012).

2.3 Gaussian Approximation of the Intractable Likelihood

The new approach to approximate the intractable likelihood (Eq. 6) relies on the linearized Gaussian approximation proposed by Linde et al. (2017). Here, the data covariance matrix given by the observational noise $\boldsymbol{\Sigma_Y}$ (Eq. 2) is adjusted by adding an additional contribution accounting for petrophysical prediction uncertainty. For our latent variable model, a first-order Taylor expansion of $\mathbf{x} \mapsto \mathcal{G}(\mathbf{x})$ around $\mathcal{F}(\boldsymbol{\theta})$ is used,

$$\mathbf{Y} = \mathcal{G}(\mathcal{F}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}_P) + \boldsymbol{\varepsilon}_O \approx \mathcal{G}(\mathcal{F}(\boldsymbol{\theta})) + \mathbf{J}_{\mathcal{F}(\boldsymbol{\theta})}\boldsymbol{\varepsilon}_P + \boldsymbol{\varepsilon}_O, \tag{7}$$

with $\mathbf{J}_{\mathcal{F}(\boldsymbol{\theta})}$ denoting the Jacobian (sensitivity) matrix of the forward solver corresponding to $\mathcal{F}(\boldsymbol{\theta})$. Due to its dependence on $\mathcal{F}(\boldsymbol{\theta})$, the sensitivity matrix is evolving between iterations (with changing $\boldsymbol{\theta}$). Using Gaussian assumptions for $p(\boldsymbol{\varepsilon}_O) = \varphi_T(\boldsymbol{\varepsilon}_O; \mathbf{0}, \boldsymbol{\Sigma_Y})$ and $p(\boldsymbol{\varepsilon}_P) = \varphi_L(\boldsymbol{\varepsilon}_P; \mathbf{0}, \boldsymbol{\Sigma}_P)$ (Eq. 2), the likelihood function (Eq. 6) is approximated by,

$$\hat{p}(\mathbf{y}|\boldsymbol{\theta}) = \varphi_T(\mathbf{y}; \mu_{\mathbf{Y}}, \widetilde{\boldsymbol{\Sigma}_{\mathbf{Y}}}) \quad \text{with} \quad \mu_{\mathbf{Y}} = \mathcal{G}(\mathcal{F}(\boldsymbol{\theta})) \quad \text{and} \quad \widetilde{\boldsymbol{\Sigma}_{\mathbf{Y}}} = \mathbf{J}^T_{\mathcal{F}(\boldsymbol{\theta})}\boldsymbol{\Sigma}_P\mathbf{J}_{\mathcal{F}(\boldsymbol{\theta})} + \boldsymbol{\Sigma_Y}. \tag{8}$$

As the linearization is made around $\mathcal{G}(\mathcal{F}(\boldsymbol{\theta}))$ and not around $\mathcal{G}(\mathcal{F}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}_P)$, errors arise when the resulting Jacobians differ. For a linear geophysical problem, there are no approximation errors. Figure 1 shows a flow chart describing this approach at iteration j; in what follows, this method will be referred to as LinGau.

2.4 Correlated Pseudo-marginal Method

The pseudo-marginal method (Beaumont 2003; Andrieu and Roberts 2009) estimates the intractable likelihood (Eq. 6) by Monte Carlo averaging over samples of the latent variable,
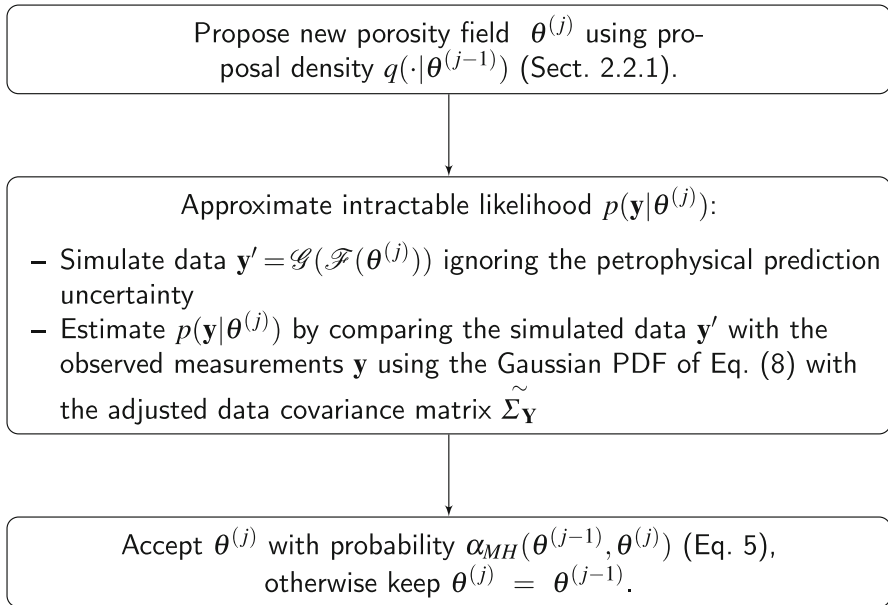
> Propose new porosity field $\theta^{(j)}$ using pro-
> posal density $q(\cdot|\theta^{(j-1)})$ (Sect. 2.2.1).

> Approximate intractable likelihood $p(\mathbf{y}|\theta^{(j)})$:
>
> – Simulate data $\mathbf{y}' = \mathscr{G}(\mathscr{F}(\theta^{(j)}))$ ignoring the petrophysical prediction
>   uncertainty
> – Estimate $p(\mathbf{y}|\theta^{(j)})$ by comparing the simulated data $\mathbf{y}'$ with the
>   observed measurements $\mathbf{y}$ using the Gaussian PDF of Eq. (8) with
>   the adjusted data covariance matrix $\widetilde{\Sigma}_{\mathbf{Y}}$

> Accept $\theta^{(j)}$ with probability $\alpha_{MH}(\theta^{(j-1)}, \theta^{(j)})$ (Eq. 5),
> otherwise keep $\theta^{(j)} = \theta^{(j-1)}$.

Fig. 1 Flow chart illustrating the LinGau method at iteration j

$$\hat{p}_N(\mathbf{y}|\boldsymbol{\theta}) = \frac{1}{N} \sum_{n=1}^{N} w(\mathbf{y}|\mathbf{X}_n, \boldsymbol{\theta}), \quad \text{with} \quad w(\mathbf{y}|\mathbf{X}_n, \boldsymbol{\theta}) = \frac{p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{X}_n) p(\mathbf{X}_n|\boldsymbol{\theta})}{m(\mathbf{X}_n|\boldsymbol{\theta})}, \quad (9)$$

where $\mathbf{X}_n \overset{i.i.d}{\sim} m(\cdot|\boldsymbol{\theta})$ for $n = 1, 2, \dots, N$ with $m(\cdot|\boldsymbol{\theta})$ denoting an importance density function. This implies that after proposing a new target parameter $\boldsymbol{\theta}$, different latent variable realizations $\mathbf{X}_n$ with the same $\boldsymbol{\theta}$ and different PPE $\boldsymbol{\varepsilon}_P$ are sampled. Then, the likelihood of each realization can be calculated and the intractable likelihood function is estimated by averaging over the obtained values. To account for the influence of importance sampling on the draws of the latent variable, weighted averaging has to be applied. Using this non-negative unbiased estimator of the likelihood leads to a MH algorithm sampling the same posterior distribution as one using the true likelihood (Beaumont 2003).

To obtain an efficient algorithm, it is crucial that the variance of the log-likelihood ratio estimator used in each MH step is low enough (Doucet et al. 2015). This can be ensured by choosing a well-working importance sampling density and by selecting a number of latent variable samples N which is high enough. Following Friedli et al. (2022), a Gaussian approximation $\mathbf{x} \mapsto \varphi_L(\mathbf{x}; \boldsymbol{\mu}_{IS}, \boldsymbol{\Sigma}_{IS})$ of $\mathbf{x} \mapsto p(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})$ is used as importance density, which implies the same linearization of the forward operator as in the LinGau approach (Eq. 7). An inappropriate linearization will lead to errors in the LinGau estimates, while it will only affect the efficiency of the pseudo-marginal method. To reduce the magnitude of N, Deligiannidis et al. (2018) introduced the correlated pseudo-marginal (CPM) method by which the samples of latent variables used in the likelihood ratio estimator are correlated. Assuming a standard-normal
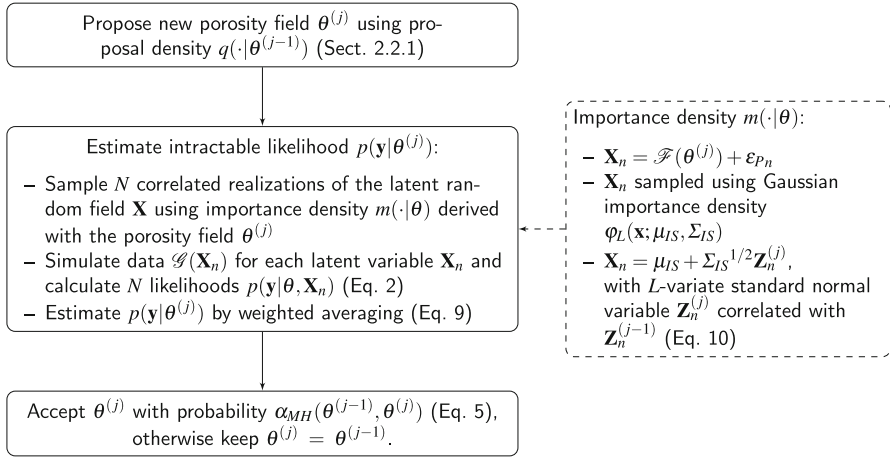
Fig. 2 Flow chart illustrating the CPM method with importance sampling at iteration j

distributed latent variable $\mathbf{X}$, the CPM method correlates one draw of iteration j with one of the former by,

$$\mathbf{X}^{(j)} = \rho\mathbf{X}^{(j-1)} + \sqrt{1-\rho^2}\boldsymbol{\epsilon}, \text{ with } \rho \in (-1, 1) \text{ and}$$

$$\boldsymbol{\epsilon} = (\epsilon_1, \epsilon_2, \ldots, \epsilon_L), \epsilon_i \overset{i.i.d.}{\sim} \mathcal{N}(0, 1). \tag{10}$$

The general applicability of the CPM method is not limited by the assumption that the latent variable has a standard-normal distribution since numerous distributions can be obtained by transformations from standard normal variates (e.g. Chen et al. 2018). The procedure of the CPM method in iteration j is illustrated in Fig. 2. Further details about the method can be found in Friedli et al. (2022).

2.5 Performance Assessment

The primarily focus of the assessment of the inversion results is on the exploration of the posterior PDF. To declare convergence of the MCMC chains, the $\hat{R}$-statistic of Gelman and Rubin (1992) is used. Subsequently, the posterior samples obtained with the considered methods are compared. For a numerical assessment of the posterior estimates, the logarithmic score (logS; Good 1952) is employed. This is a so-called scoring rule (Gneiting and Raftery 2007) evaluating the accuracy of a predictive PDF $\theta \mapsto \hat{p}(\theta)$ with respect to a true value θ_{true}. The logarithmic score is defined as $\text{logS}(\hat{p}, \theta_{\text{true}}) = -\log\hat{p}(\theta_{\text{true}})$ and kernel density estimates with manually-selected bandwidths are applied to transform the posterior samples into a PDF. Furthermore, the number of target parameters in which the true porosity value θ_{true} is in the range of the posterior samples is considered, as well as the spread of the posterior samples as quantified by their standard deviations.

3 Case Study

The new linearized Gaussian approach is compared against the CPM method using the first case study considered by Friedli et al. (2022).

3.1 Synthetic Data Generation

The considered model domain is a water-saturated subsurface area of $7.2\,\text{m} \times 7.2\,\text{m}$. As in Friedli et al. (2022), the target porosity field is assumed to be a Gaussian random field with known mean ($\mu_\theta = 0.39$) and exponential covariance function. For the latter, a sill of 2×10^{-4} is assumed such as geometric anisotropy with the main, horizontal direction having an integral scale of $5.4\,\text{m}$ and the integral scale ratio between the horizontal and vertical direction being 0.13. For the parameterization of the porosity field, a regular (50×50)-dimensional grid ($D^2 = 2500$) is used. The "true" porosity field θ_{true} for this case study is depicted in Fig. 3a. Given the porosity θ, the dielectric constant κ is predicted using the complex refractive index model (CRIM; Roth et al. 1990), from which the slowness field (our latent variable $\mathbf{X}$) can be derived,

$$\mathbf{x} = \sqrt{c^{-2}\kappa} + \boldsymbol{\varepsilon}_P = \frac{1}{c}\big(\sqrt{\kappa_s} + (\sqrt{\kappa_w} - \sqrt{\kappa_s})\theta\big) + \boldsymbol{\varepsilon}_P, \tag{11}$$

with κ_w and κ_s denoting the dielectric constants of water [81] and mineral grains [5], respectively, and c referring to the speed of light in vacuum [0.3 m/ns]. A PPE $\boldsymbol{\varepsilon}_P$ is added (Fig. 3b), which is a realization of a centred Gaussian random field over a regular two dimensional grid of size 50×50. Thereby, the same correlation structure as for the porosity field is used with a sill of $7.11 \times 10^{-2}\,\text{ns}^2/\text{m}^2$. The "true" slowness field is depicted in Fig. 3c while Fig. 3d shows a scatter plot depicting the influence of the PPE on the slowness values.

The 625 first-arrival travel times are generated using 25 equidistant GPR transmitters located on the left side and 25 receivers on the right side of the model domain (the transmitter-receiver layout is shown in Fig. 3c). As forward solver $\mathbf{y} \mapsto \mathcal{G}(\mathbf{y})$, the non-linear (eikonal) solver time2D of Podvin and Lecomte (1991) is used. The observational noise $\boldsymbol{\varepsilon}_O$ is assumed to be $i.i.d.$ centered normal with a standard deviation of 1 ns. The noise-affected synthetic first-arrival traveltimes are depicted in Fig. 3e.

3.2 Inversion Setting and Prior Assumptions

As the considered target space is high-dimensional with 2500 unknown parameters, it is crucial to choose a well-working proposal scheme for the MH algorithm. Due to its convincing performance in Friedli et al. (2022), this study relies on prior-preserving DREAM$_{(ZS)}$ proposals (Sect. 2.2.1) and four MH chains are run in parallel. For both the LinGau and the CPM method, the linearization used to adjust the data covariance matrix and the importance sampling, respectively, is updated every tenth MCMC iteration. For the CPM method, a configuration of $N = 10$ and $\rho = 0.95$ is used as
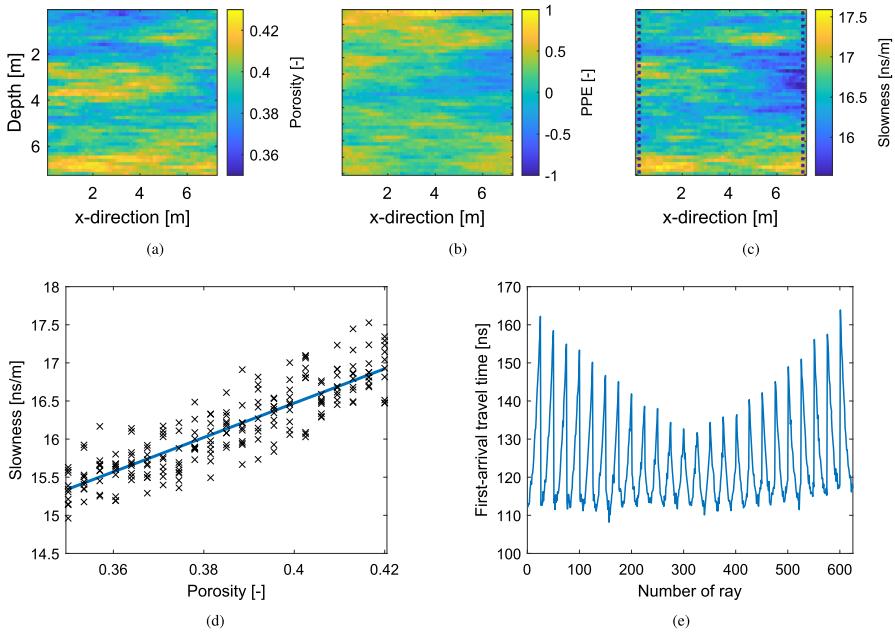
Fig. 3 The synthetic "true" model adapted from Friedli et al. (2022): (**a**) porosity field θ_{true}, (**b**) PPE field $\varepsilon_{P\,\text{true}}$, (**c**) slowness field $\mathbf{x}_{\text{true}}$, (**d**) scatter plot of porosity and slowness values per grid cell, the line depicts the petrophysical relationship (Eq. 11) without considering PPE and (**e**) set of noise-affected first-arrival travel times $\mathbf{y}$

this choice guarantees an appropriate variance of the log-likelihood ratio estimator (see, Friedli et al. 2022).

For the prior on porosity, a Gaussian PDF $p(\theta) = \varphi_{2500}(\theta; \mu_\theta, \Sigma_\theta)$ with known mean μ_θ and covariance structure Σ_θ (the same values as for the data generation) is assumed. Then, the 2500-dimensional vector $\mathbf{Z}$ defining the porosity by $\theta = \mu_\theta + \Sigma_\theta^{1/2}\mathbf{Z}$ is inferred, with $\mathbf{Z}$ having a multivariate standard-normal prior PDF. For the PPE ε_P also a Gaussian prior PDF $p(\varepsilon_P) = \varphi_{2500}(\theta; 0, \Sigma_P)$ with known covariance structure Σ_P is used. For the likelihood function, centred independent Gaussian measurement errors ε_O with a standard deviation of 1 ns as in the data generation process are assumed.

The LinGau and CPM methods are compared with an inversion ignoring the petrophysical prediction uncertainty (No PPE). In this case, the intractable likelihood $p(\mathbf{y}|\theta)$ is estimated by $\varphi_T(\mathbf{y}; \mathcal{G}(\mathcal{F}(\theta)), \Sigma_Y)$. As prior-preserving DREAM$_{(ZS)}$ proposals lead to an unfeasible low acceptance rate in this case, standard DREAM$_{(ZS)}$ proposals are employed.

3.3 Results

Figure 4a, b display the posterior mean estimates of the LinGau and the CPM method, respectively. Both images look very similar and the structural resemblance to the true
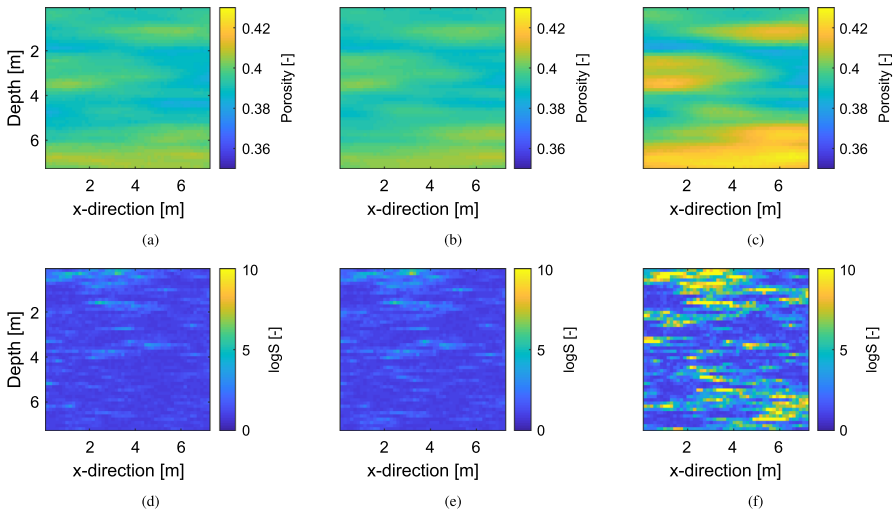
Fig. 4 Estimated posterior means of the porosity field θ obtained with (**a**) the LinGau method, (**b**) the CPM method and (**c**) an inversion ignoring the petrophysical prediction uncertainty. Corresponding pixel-wise logarithmic scores assessing the estimated posterior PDFs for (**d**) the LinGau method, (**e**) the CPM method and (**f**) the inversion ignoring the petrophysical prediction uncertainty

porosity field in Fig. 3a is high. Both methods need about 100,000 MCMC iterations to converge and have a very similar acceptance rate (AR) of 10%. They also lead to a similar performance in terms of posterior exploration as both methods sample the true porosity value in all of the pixels and have very similar median standard deviations and logarithmic scores (Fig. 4d, e; Table 1). By comparing the estimated posterior means of LinGau and CPM (Fig. 4a, b) with the one of the inversion ignoring the petrophysical prediction uncertainty (Fig. 4c), it is found that the mean estimate of the inversion ignoring the petrophysical prediction uncertainty has larger amplitudes even if the mean estimates are structurally similar. Its posterior exploration is less extensive, leading to a higher median logarithmic score (Fig. 4f), roughly half the median posterior standard deviation and many pixels that never sample the corresponding true porosity value (about one eighth of the pixels; Tab. 1). For this example exhibiting weak non-linearity and rather small petrophysical prediction uncertainty, it is concluded that the performance for LinGau and CPM are similar. On the other hand, ignoring petrophysical prediction uncertainty leads to biased estimates and too small uncertainty bounds.

4 Sensitivity Analysis

In the previous synthetic test case, both the LinGau and the CPM method perform comparably well. As the computational cost of the LinGau method is lower (no need for N repeated sampling of the latent variable at each iteration), the use of the LinGau method would be recommended in such a setting. However, the degree of non-linearity in the geophysical forward operator is rather low for this test case. This is illustrated

Table 1 Summary of the results for the Gaussian porosity field example: the acceptance rates (AR), the convergence (Conv) showing the number of iterations needed for the 99th percentile of the parameter's $\hat{R}$-statics to be below 1.2, the percentage of pixels in which the true porosity value θ_{true} lies within the range of posterior samples, the median logarithmic score (logS) and the median posterior standard deviation (Post SD)

Method	AR (%)	Conv	θ_{true} (%)	logS	Post SD
LinGau	10	108,000	100.00	1.10	10.7×10^{-3}
CPM	10	96,000	100.00	1.16	10.5×10^{-3}
No PPE	15	104,000	87.24	2.45	5.2×10^{-3}
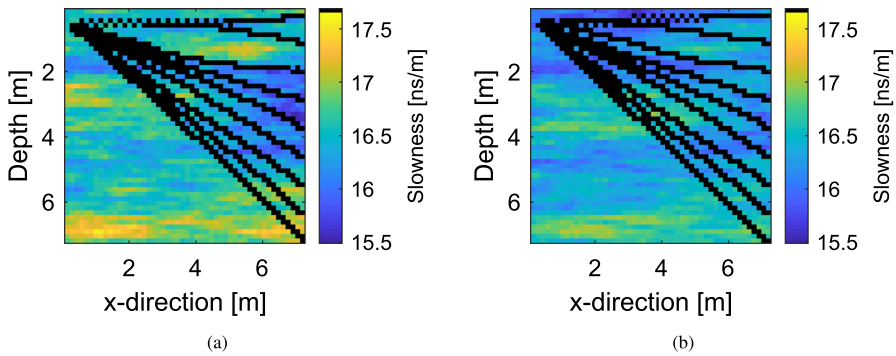
Fig. 5 Exemplary GPR ray paths for (**a**) the true slowness field $\mathbf{x}_{\text{true}}$ of the test case presented in Sect. 3 and (**b**) the slowness field resulting from the true porosity field θ_{true} when ignoring the PPE ε_P

in Fig. 5 showing exemplary ray paths for the true slowness field $\mathbf{x}_{\text{true}}$ (Fig. 5a) and the slowness field based on the true porosity field θ_{true} but ignoring the PPE (Fig. 5b): most ray paths are close to linear and they are very similar for both fields. Since the LinGau approximation of the likelihood (Eq. 8) relies on a first-order Taylor expansion of the physical forward solver (Eq. 7), it deteriorates when adding PPE ε_P realizations lead to different ray-paths than the field around which the linearization is made. For a linear geophysical relationship, the first-order Taylor expansion is exact for any degree of Gaussian petrophysical prediction uncertainty.

In practice, the question is how to decide, for a given setting, if the LinGau approximation of the likelihood is accurate enough. To shed light on this, a sensitivity analysis exploring the performances of the LinGau and CPM methods for different levels of petrophysical prediction uncertainty is performed. To make the comparison more didactic and to avoid unrelated challenges associated with a very-high dimensional target space, the porosity field is assumed to be layered. Generally, the same setup as in Sect. 3 is considered, but only with 13 transmitters and receivers (169 data points). The observational error is assumed to have a standard deviation of 1 ns.
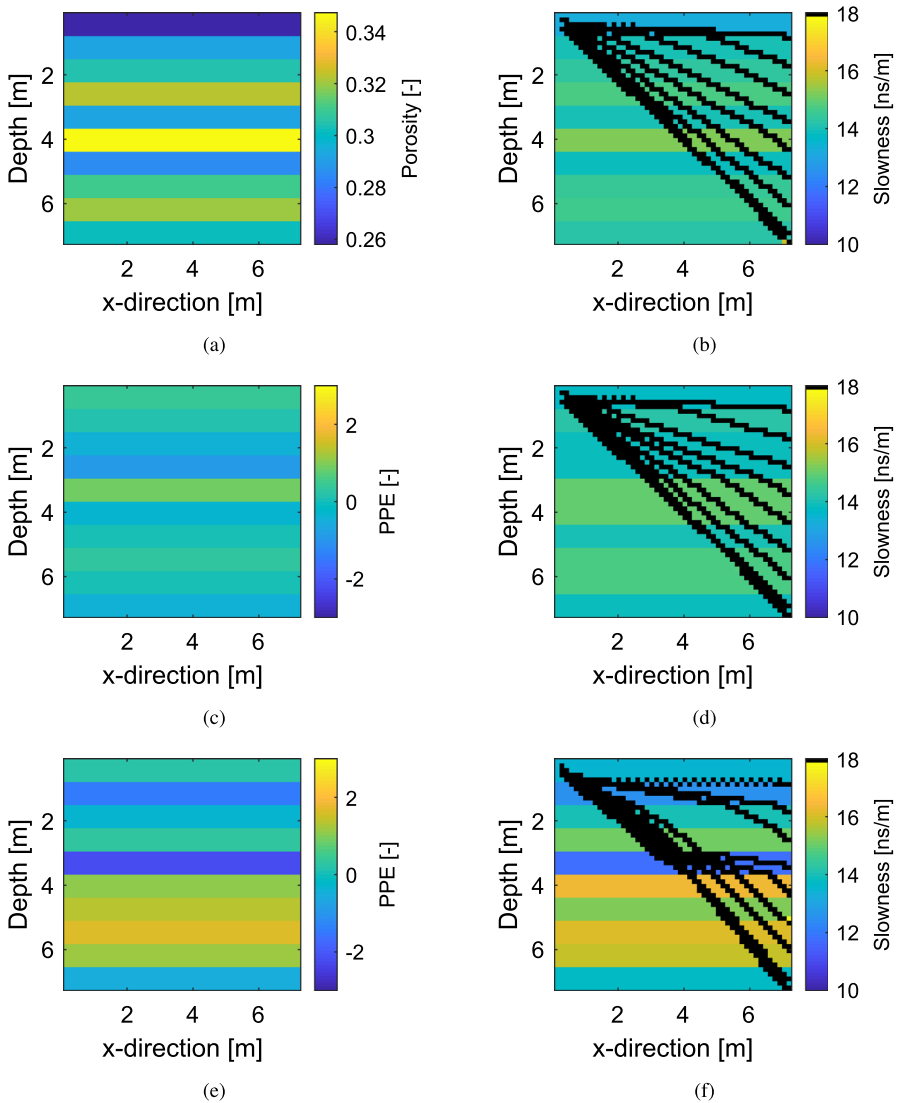
Fig. 6 Synthetic "true" models for the layered test cases with the (**a**) same porosity field θ_{true} and different PPE for (**c**) Setting 1 and (**e**) Setting 2. The right column shows the slowness fields with exemplary ray paths for (**b**) the true porosity θ_{true} without adding PPE, (**d**) Setting 1 and (**f**) Setting 2

4.1 Likelihood Estimation

The "true" porosity field $\theta_{\text{true}} = (\theta_1, \ldots, \theta_{10})$ is generated by assuming 10 horizontal layers of equal thickness and drawing independently from a Gaussian distribution with mean 0.3 and standard deviation 0.03 (Fig. 6a). The resulting slowness field (Eq. 11) is distorted with a layered PPE field having zero mean and independent layers with standard deviation σ_P. For each of sixteen different standard deviation values σ_P
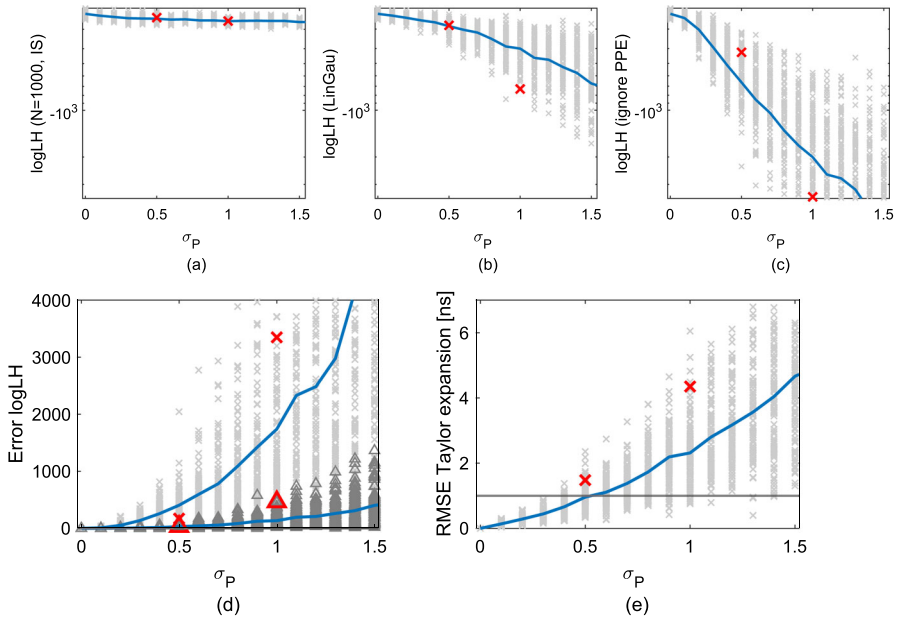
Fig. 7 Estimators of log-likelihood $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ of the true porosity field in the layered setting as a function of petrophysical prediction uncertainty (σ_P): (**a**) importance sampling estimate using 1000 samples of the PPE, (**b**) LinGau estimate and (**c**) estimate ignoring the PPE. For each value of σ_P, the same underlying porosity field $\boldsymbol{\theta}_{\text{true}}$ is used with 100 different realizations of the PPE, leading to 100 data sets each. The crosses indicate the values for each data set and the solid lines their mean. (**d**) The mean absolute difference between the values of (**a**) and (**b**) are shown in darkgrey triangles and (**a**) and (**c**) in lightgrey crosses. (**e**) The RMSEs in the first-order Taylor expansion (Eq. 7) of the data sets; the black horizontal line shows the standard deviation of the observational noise (1 ns). The red symbols in (**a**)–(**e**) indicate the errors of the settings used in the subsequent inversion examples; thereby in (**d**), the red triangles refer to the errors obtained when using LinGau and the red crosses refer to the errors obtained when ignoring PPE

ranging from 0.0 to 1.5, one hundred data sets are generated using the same porosity field $\boldsymbol{\theta}_{\text{true}}$ (Fig. 6a) but different realizations of the PPE $\boldsymbol{\varepsilon}_P$. Thereby, the true log-likelihood value $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ of $\boldsymbol{\theta}_{\text{true}}$ has a different value depending on the realization of the PPE and the observational noise, even for the same σ_P. Two exemplary PPE fields and resulting slowness fields with $\sigma_P = 0.5$ and $\sigma_P = 1.0$ are depicted in Fig. 6c–f, respectively.

For each of the one hundred data sets per value of σ_P, $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ is approximated by using the LinGau approach. The corresponding values and their mean for the different σ_P are shown in Fig. 7b. Those values are compared with the log-likelihood estimates obtained under the assumption of no PPE (Fig. 7c). Eventually, the aim is to assess these estimates by comparing them to the true log-likelihood values $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ of the corresponding data set. Due to the intractability inherited by the latent variable model, the analytical solution remains elusive and the unbiased importance sampling estimate of Eq. 9 employing 1000 realizations of the PPE (Fig. 7a) is applied. Due to the high number of samples, a well-specified importance density and the low dimensionality of the problem at hand, this estimator is close to the true log-likelihood value. To show that

this assumption is valid, the standard deviation of the importance sampling estimator of $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ for both the exemplary settings with $\sigma_P = 0.5$ and $\sigma_P = 1.0$ of Fig. 6 is evaluated (indicated with red symbols in Fig. 7). For both $\sigma_P = 0.5$ and $\sigma_P = 1.0$, the standard deviation is below one for log-likelihood values around -260. With the underlying assumption of the estimators in Fig. 7a being close to the true value, they are compared with the ones of the LinGau method (Fig. 7d, darkgrey triangles). The absolute errors in the log-likelihood estimation of LinGau increase with increasing σ_P. While the mean absolute error in the log-likelihood estimator for $\sigma_P = 0.5$ is around 30, it grows to 150 for $\sigma_P = 1.0$ and to 400 for $\sigma_P = 1.5$. However, in comparison to the errors of the method ignoring the PPE (Fig. 7d, lightgrey crosses), the errors of the LinGau approach are comparably small. When ignoring the PPE completely, the mean absolute error for $\sigma_P = 0.5$ is about 400, for $\sigma_P = 1.0$ it is 1700 and for $\sigma_P = 1.5$ even 5000. Thereby, the method ignoring the PPE almost always underestimates the true log likelihood values $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ as it cannot account for the true PPE and therefore gives reduced likelihood to the true porosity field $\boldsymbol{\theta}_{\text{true}}$. The same holds true for the LinGau approach with increasing σ_P as the method accounts for an increasingly wrong PPE.

The LinGau approximation relies on a first-order Taylor expansion of the physical forward solver (Eq. 7). This approximation deteriorates with increasing degree of PPE as can be seen in Fig. 7e depicting the root mean square errors (RMSE) of the Taylor expansions for the data sets with increasing σ_P. While the mean of the RMSEs is comparable to the observational noise for $\sigma_P = 0.5$, it is twice as large for $\sigma_P = 1.0$. To establish the influence of the discussed errors on the inversion results, the setups introduced in Fig. 6 that are indicated with red symbols in Fig. 7 are considered: The first with $\sigma_P = 0.5$ employs rather small errors, however, the standard deviation of the PPE is twice as high as in the case-study in Sect. 3. In the second setting with $\sigma_P = 1.0$ the errors are doubled and a realization is targeted where the error in the LinGau likelihood approximation is especially high (about 460 while the mean error for $\sigma_P = 1.0$ is 150).

4.2 Inversion

Using the two layered synthetic data sets (Fig. 6), the MH algorithms are run with three chains in parallel. Due to the simplicity of the problem, a basic Gaussian random walk is used as proposal scheme, within which, for comparison purposes, a step width (standard deviation) of 0.005 is applied for all methods. Furthermore, the linearizations used in the LinGau and CPM methods are updated at every MCMC iteration to prevent any errors resulting from less frequent updates. For the CPM method, importance sampling and a configuration of $N = 3$, $\rho = 0.9$ for the first and $N = 50$, $\rho = 0.975$ for the second setting is used. For the layers of the porosity field, independent Gaussian prior PDFs with mean 0.3 and standard deviation 0.03 and for the PPE layers, independent centred Gaussian priors with a standard deviation of 0.5 and 1.0 are assumed, respectively. Finally, for the likelihoods uncorrelated Gaussian observational noise with a standard deviation of 1 ns is assumed.

The estimated marginal posterior PDFs for both settings are shown in Fig. 8 for three of the ten layers. These results are representative of the other layers, but a summary
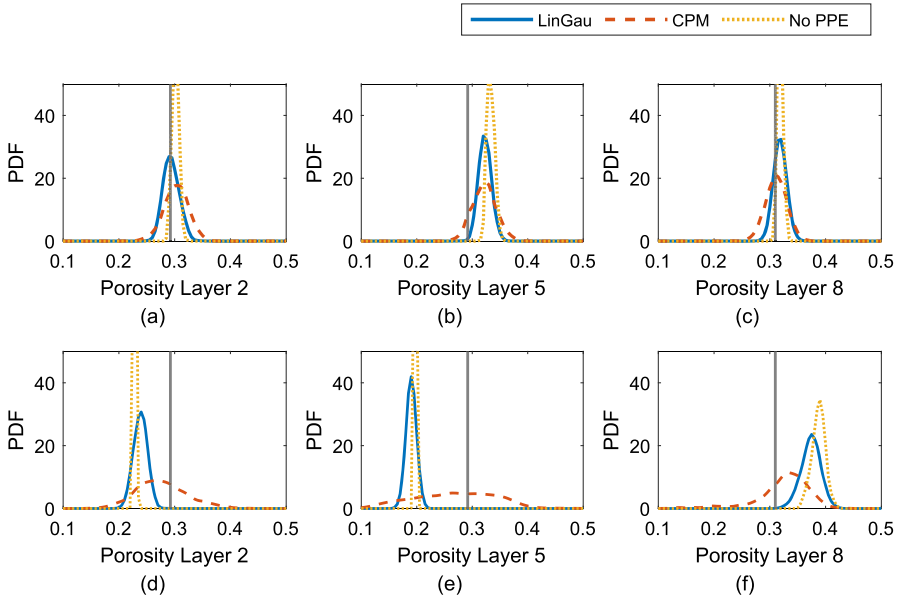
Fig. 8 Estimates of the marginal posterior PDFs for the layered test cases: Setting 1 for (**a**) layer 2, (**b**) layer 5 and (**c**) layer 8. Setting 2 for (**d**) layer 2, (**e**) layer 5 and (**f**) layer 8. The solid vertical lines indicate the true porosity values

of the performance for all layers is for completeness provided in Table 2. For the first setting (Fig. 8a–c), the modes of all considered methods are located around rather similar values. While the method ignoring the PPE generates samples that do not include the true value of the porosity in half of the layers, both LinGau and CPM capture the true value in all of the ten layers (Table 2). However, while the LinGau method already doubles the standard deviation of the posterior samples compared to the method ignoring the PPE, the CPM method leads to a further doubling compared to LinGau (Table 2). Compared to the method ignoring the PPE, the LinGau method reduces the median logarithmic score from 3.53 to −2.37 and the CPM method reduces it even further to −2.71. For the second setting (Fig. 8d–f), rather different posterior estimates are obtained for the three methods. The method ignoring the PPE generates posterior samples which are far from the true porosity values and with a small standard deviation (Table 2). The realizations obtained with the LinGau method have a twice as high standard deviation but are also centered far from the true values. The CPM method yields posteriors three times wider than those obtained with LinGau. Thereby, while the posterior samples obtained with the CPM method include the true values of the porosity in all layers, the LinGau method misses them for two layers and the method ignoring the PPE in seven layers. The median logarithmic score of the CPM method (−2.00) is distinctly lower than the one of the LinGau method (−1.35), which in turn is dramatically lower than for the method ignoring the PPE (378.77).

Table 2 Summary of the results obtained for the study targeting a layered porosity field: the acceptance rates (AR), convergence with respect to the $\hat{R}$-statics (Conv), the number of layers in which the true porosity value lies within the range of posterior samples (θ_{true}), the median logarithmic score (logS) and the median posterior standard deviation (Post SD)

Method	Setting	AR (%)	Conv	θ_{true}	logS	Post SD
LinGau		40	2500	10/10	−2.37	11.8×10^{-3}
CPM	1 ($\sigma_P = 0.5$)	40	12,000	10/10	−2.71	21.9×10^{-3}
No PPE		2	13,000	5/10	3.53	4.2×10^{-3}
LinGau		40	6000	8/10	−1.35	15.0×10^{-3}
CPM	2 ($\sigma_P = 1.0$)	30	51,000	10/10	−2.00	46.4×10^{-3}
No PPE		2	20,000	3/10	378.77	7.2×10^{-3}

5 Discussion

In this work, non-linear geophysical inversion problems involving uncertain petrophysical relationships are targeted. Two different approaches to account for the corresponding intractable likelihood function are explored: a linearized Gaussian approximation (LinGau) and the correlated pseudo-marginal (CPM) method. The performance of these two methods for examples with increasing petrophysical prediction uncertainty is assessed. This work is a continuation of Friedli et al. (2022) as it is the first time the LinGau method is employed and compared with the CPM method. A synthetic crosshole travel-time tomography is first considered, with the aim of inferring a water-saturated multivariate Gaussian porosity field in a situation of only moderate non-linearity and petrophysical prediction uncertainty (adapted from Friedli et al. 2022). The results obtained with the LinGau and the CPM methods are very similar (Fig. 4; Table 1), thereby, indicating that the approximate LinGau method works well in this setting. The LinGau method is then to be preferred as it only necessitates one forward simulation for each MCMC chain and iteration, while CPM with $N = 10$ (number of latent variable samples) uses ten times as many forward simulations. If parallelization with $N = 10$ more processors are employed, then the computational time can still be made similar to the LinGau approach even if the overall computational cost is significantly higher. For this test case, one iteration for one MCMC chain on a standard laptop takes about 0.3 s for LinGau and 2.1 s for CPM ($N = 10$ and no parallelization). Importantly, this example demonstrated that an inversion that ignores petrophysical prediction uncertainty (the most common case in the literature) leads to biased estimates and underestimations of the widths of the posterior estimates by a factor of two.

In Sect. 4, the effect of the degree of petrophysical prediction uncertainty on the likelihood estimation and inversion results is studied for the different methods. To simplify the comparison, layered porosity and PPE fields are used. By comparing log-likelihood estimates, it is shown that the errors in the LinGau method increase with increasing degree of petrophysical prediction uncertainty, even if these errors are much smaller than for the method ignoring petrophysical prediction uncertainty (Fig. 7). By

selecting one data set with twice as high petrophysical prediction uncertainty as in the previous multivariate Gaussian example ($\sigma_P = 0.5$, Fig. 6c) and one with four times as high uncertainty ($\sigma_P = 1.0$, Fig. 6e), the influence of the likelihood estimation errors on the inversion results is investigated. For the first setting (moderate degree of PPE), the LinGau method performs less well than CPM (underestimation of posterior uncertainty by a factor of two; higher logS scores). However, employing the LinGau method still enables a reasonable approximation of the posterior modes of the layers. Although the true log-likelihood value $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ is underestimated by about 30 with LinGau (Fig. 7d), the RMSE of the Taylor approximation is roughly on the same order of magnitude as the observational error (Fig. 7e). That is different in the second setting with a higher degree of PPE, where the true log-likelihood value $p(\mathbf{y}|\boldsymbol{\theta}_{\text{true}})$ is underestimated by about 460 and the RMSE of the Taylor approximation is four times as high as the observational noise. The resulting deterioration in performance for the LinGau method is drastic: underestimation of the posterior uncertainty by a factor of three and twice as high difference in the logS score to CPM than for the first setting. The growing error in the LinGau estimate is due to the Jacobian $\mathbf{J}_{\mathcal{F}(\boldsymbol{\theta})}$ being increasingly different when considering or not considering the PPE (Fig. 6b, f). For the CPM method, the growing petrophysical prediction uncertainty is accompanied by a very important increase in the posterior standard deviation of the samples (Fig. 8d–f), guaranteeing that the method samples the true porosity values even if the mode of the posterior may be located at the wrong place. Even if the LinGau method still provides much better results than the common approach of ignoring PPE altogether, the linearization on which this method is based (Eq. 7) is unable to account for non-linear effects associated with specific PPE realizations. In contrast, the CPM method only relies on the linearization to derive an importance sampling distribution and errors in this distribution will only lead to slower convergence of the CPM method while still targeting the true posterior distribution.

The approximation error of the first-order Taylor expansion used in the LinGau method (Eq. 7) grows with increasing non-linear effects related to the PPE. By comparing the forward operators for latent variables with $(\mathcal{F}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}_P)$ and without considering the PPE $\boldsymbol{\varepsilon}_P$ ($\mathcal{F}(\boldsymbol{\theta})$), this relation can be established (as in Figs. 5 and 6). It is seen that high standard deviations in the $\boldsymbol{\varepsilon}_P$ have an adverse effect on the approximation accuracy by highly influencing the forward operator. If the PPE $\boldsymbol{\varepsilon}_P$ strongly influences the Jacobian (as in Fig. 6e), caution is advised when applying the LinGau method (Fig. 7d). To choose between the LinGau and the CPM method, this study recommends to investigate the RMSE of the first-order Taylor expansion (Fig. 7e) and to compare it with the observational noise. For a setting when the RMSE is significantly lower than the observational error (as for $\sigma_P = 0.25$ in Sect. 3), the recommendation is to use the LinGau method due to its lower computational cost (the effect of the petrophysical prediction uncertainty is incorporated into the likelihood function and no importance sampling of latent variables is needed). If the RMSE is in the same order of magnitude as the observational error (as for $\sigma_P = 0.5$ in Sect. 4), the LinGau method can be applied if moderate errors in the posterior estimation are acceptable. If the RMSE is clearly higher than the observational error (as for $\sigma_P = 1.5$ in Sect. 4), the exact CPM method should be used. If this is computationally too expensive, one could consider

accounting for the approximation error in the LinGau method by treating it as a model error term (see e.g. Hansen et al. 2014).

The approximation underlying the LinGau approach could easily be incorporated in deterministic inversion methods (using gradients), while this is impossible for CPM. On the other hand, while the LinGau method requires Gaussian assumptions for the PPE, the CPM method can sample the latent variables from a variety of possible distributions. But in practice the CPM method needs a well-working importance sampling strategy to be efficient. For the presented test cases, a decreasing acceptance rate for the CPM method is observed with increasing petrophysical prediction uncertainty if the number of latent variables N and correlation ρ are fixed. This occurs as our importance sampling scheme gets more and more inaccurate, which can be compensated by using larger N or ρ. Although the efficiency is reduced, the accuracy of the posterior samples remains the same as the importance sampling is only used to decrease the variance of the likelihood estimator. In contrast, the LinGau method does not only loose efficiency with increasing non-linearity, but also leads to overconfident and biased estimates of the target parameters as the approximation of the likelihood is getting increasingly inaccurate (Fig. 8).

In the presented test cases, the mean and the covariance structure of the target and petrophysical prediction error fields are assumed to be known. In a setting where they are unknown, hierarchical Bayes could be employed (as e.g. in Laloy et al. 2015; Brunetti and Linde 2017). A recent study by Friedli et al. (2023) targets a setting where the problem is formulated differently: only the posterior mean and the covariance structure are derived, while the small-scale variations in the model domain are accounted for but not inferred. In this study, a crosshole ray-based setting that is only weakly non-linear is considered. More non-linear problems such as electrical resistivity tomography or surface-based seismic refraction tomography might exhibit an even stronger sensitivity to petrophysical prediction uncertainty and the applicability of the LinGau might be reduced compared with the present study. Indeed, the errors in the LinGau approximation are not due to the petrophysical predication uncertainty as such, but rather how individual realizations of it affects the sensitivity patterns compared to the Taylor expansion that is performed in absence of petrophysical prediction uncertainty. Even if the focus herein is on how to account for petrophysical prediction uncertainty, both the LinGau and the CPM methods could be adapted to, for instance, account for 3-D effects in 2-D inversions, variations in porosity in tracer-test tomography targeting the permeability field, or to account for hydraulic storativity fluctuations in hydraulic tomography studies. Indeed, this type of latent variable problem arises as soon as the measured data do not only depend on the main parameters of interest but also on some other variables influencing the response.

6 Conclusions

This work focus on geophysical inversion problems targeting the posterior distribution of (hydro)geological parameters while accounting for uncertain petrophysical relationships and non-linear physics. The resulting intractable likelihood function is accounted for by either the linearized Gaussian approximation (LinGau) method or

the correlated pseudo-marginal (CPM) method. The LinGau method, so far untested in geophysical inversion, is an approximate method that is computationally cheaper than CPM as it does not necessitate a Monte Carlo estimation of the likelihood at each MCMC iteration. In a first case study, a water-saturated multivariate Gaussian porosity field is considered for which the scatter in the petrophysical relationship and the non-linearity is comparatively small. In this setting involving crosshole first-arrival travel times, the LinGau method succeeds equally well as CPM in exploring the posterior distribution. For comparison, an inversion ignoring petrophysical uncertainty provides biased results and too narrow uncertainty estimates. In a subsequent sensitivity analysis considering layered fields, the degree of petrophysical prediction uncertainty is increased, thereby introducing increasing inaccuracies associated with the Taylor expansion on which the LinGau method is based. Consequently, the Lin-Gau method produces increasingly inaccurate results as the petrophysical prediction uncertainty grows such that the true values are more and more often unsampled and the logarithmic scores are high. In contrast, the CPM method performs very well for all settings and accommodate the growing uncertainty in the petrophysical uncertainty, while this is only partially achieved by the LinGau method. The computationally less intensive LinGau method is attractive when the impact of the scatter of petrophysical prediction uncertainty is small compared to the observational noise. In comparison, the computationally more costly CPM method is an exact and much more general method that clearly outperforms the LinGau method when the petrophysical uncertainty grows in magnitude, but it needs an efficient importance sampling distribution to work well in practice. If the CPM method is computationally too expensive for a given application and if petrophysical uncertainty is significant, it is still better to use the LinGau method than inversions ignoring petrophysical prediction errors as the resulting results are less biased and the underestimation of posterior uncertainty is less pronounced.

Declarations

Conflict of interest The authors declare that there is no financial/non-financial conflict of interest.

References

Andrieu C, Roberts GO (2009) The pseudo-marginal approach for efficient Monte Carlo computations. Ann Stat 37(2):697–725. https://doi.org/10.1214/07-AOS574

Beaumont MA (2003) Estimation of population growth or decline in genetically monitored populations. Genetics 164(3):1139–1160. https://doi.org/10.1093/genetics/164.3.1139

Bosch M (1999) Lithologic tomography: from plural geophysical data to lithology estimation. J Geophys Res Solid Earth 104(B1):749–766. https://doi.org/10.1029/1998JB900014

Brunetti C, Linde N (2017) Impact of petrophysical uncertainty on Bayesian hydrogeological inversion and model selection. Adv Water Resour 111:346–359. https://doi.org/10.1016/j.advwatres.2017.11.028

Chen V, Dunlop MM, Papaspiliopoulos O, Stuart AM (2018) Dimension-robust MCMC in Bayesian inverse problems. arXiv preprint. https://doi.org/10.48550/arXiv.1803.03344

Cotter SL, Roberts GO, Stuart AM, White D (2013) MCMC methods for functions: modifying old algorithms to make them faster. Stat Sci. https://doi.org/10.1214/13-STS421 arXiv:1202.0709

Deligiannidis G, Doucet A, Pitt MK (2018) The correlated pseudo-marginal method. J R Stat Soci Ser B Stat Methodol 80(5):839–870. https://doi.org/10.48550/arXiv.1511.04992

Doucet A, Pitt MK, Deligiannidis G, Kohn R (2015) Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. Biometrika 102(2):295–313. https://doi.org/10.1093/biomet/asu075

Friedli L, Linde N, Ginsbourger D, Doucet A (2022) Lithological tomography with the correlated pseudo-marginal method. Geophys J Int 228–2:839–856. https://doi.org/10.1093/gji/ggab381

Friedli L, Linde N, Ginsbourger D, Visentini AF, Doucet A (2023) Inference of geostatistical hyperparameters with the correlated pseudo-marginal method. Adv Water Resour. https://doi.org/10.1016/j.advwatres.2023.104402

Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences. Stat Sci 7(4):457–472. https://doi.org/10.1214/ss/1177011136

Gneiting T, Raftery AE (2007) Strictly proper scoring rules, prediction, and estimation. J Am Stat Assoc 102(477):359–378. https://doi.org/10.1198/016214506000001437

Good IJ (1952) Rational decisions. J R Stat Soc Ser B 14:107–114. https://doi.org/10.1111/j.2517-6161.1952.tb00104.x

Hansen TM, Cordua KS, Jacobsen BH, Mosegaard K (2014) Accounting for imperfect forward modeling in geophysical inverse problems-exemplified for crosshole tomography. Geophysics 79(3):H1–H21. https://doi.org/10.1190/geo2013-0215.1

Hastings WK (1970) Monte Carlo sampling methods using Markov Chains and their applications. Biometrika 57(1):97–109. https://doi.org/10.2307/2334940

Hinnell A, Ferré T, Vrugt J, Huisman J, Moysey S, Rings J, Kowalsky M (2010) Improved extraction of hydrologic information from geophysical data through coupled hydrogeophysical inversion. Water Resour Res 46(4):W00D40. https://doi.org/10.1029/2008WR007060

Kowalsky MB, Finsterle S, Peterson J, Hubbard SS, Rubin Y, Majer E, Ward A, Gee G (2005) Estimation of field-scale soil hydraulic and dielectric parameters through joint inversion of GPR and hydrological data. Water Resour Res 41(11):W11425. https://doi.org/10.1029/2005WR004237

Laloy E, Vrugt JA (2012) High-dimensional posterior exploration of hydrologic models using multiple-try DREAM (ZS) and high-performance computing. Water Resour Res 48(1):W01526. https://doi.org/10.1029/2011WR010608

Laloy E, Linde N, Diederik J, Vrugt JA (2015) Probabilistic inference of multi-Gaussian fields from indirect hydrological data using circulant embedding and dimensionality reduction. Water Resour Res 51:4224–4243. https://doi.org/10.1002/2014WR016395

Linde N, Ginsbourger D, Irving J, Nobile F, Doucet A (2017) On uncertainty quantification in hydrogeology and hydrogeophysics. Adv Water Resour 110:166–181. https://doi.org/10.1016/j.advwatres.2017.10.014

Lochbühler T, Breen SJ, Detwiler RL, Vrugt JA, Linde N (2014) Probabilistic electrical resistivity tomography of a CO_2 sequestration analog. Geophysics 107:80–92. https://doi.org/10.1016/j.jappgeo.2014.05.013

Mavko G, Mukerji T, Dvorkin J (2009) The rock physics handbook: tools for seismic analysis of porous media, 2nd edn. Cambridge University Press, Cambridge

Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E (1953) Equation of state calculations by fast computing machines. J Chem Phys 21(6):1087–1092. https://doi.org/10.1063/1.1699114

Mosegaard K, Tarantola A (1995) Monte Carlo sampling of solutions to inverse problems. J Geophys Res Atmos 1001(B7):12431–12448. https://doi.org/10.1029/94JB03097

Podvin P, Lecomte I (1991) Finite difference computation of traveltimes in very contrasted velocity models: a massively parallel approach and its associated tools. Geophys J Int 105(1):271–284. https://doi.org/10.1111/j.1365-246X.1991.tb03461.x

Roth K, Schulin R, Flühler H, Attinger W (1990) Calibration of time domain reflectometry for water content measurement using a composite dielectric approach. Water Resour Res 26(10):2267–2273. https://doi.org/10.1029/WR026i010p02267

Ruggeri P, Irving J, Holliger K (2015) Systematic evaluation of sequential geostatistical resampling within MCMC for posterior sampling of near-surface geophysical inverse problems. Geophys J Int 202:961–975. https://doi.org/10.1093/gji/ggv196

Van der Vaart A (2000) Asymptotic statistics. Cambridge University Press, Cambridge