Use of machine learning to estimate statistics of the posterior distribution in probabilistic inverse problems - an application to airborne EM data.

T. M. Hansen[1] and C. C. Finlay[2]

[1]Department of Geoscience, Aarhus University, Denmark
[2]DTU Space, Technical University of Denmark, Kgs. Lyngby, Denmark

Key Points:

- A machine learning approach for the probabilistic solution of inverse problems by directly estimating posterior statistics of any continuous or discrete feature of the posterior distribution.
- Allows the use of complex prior information and noise models.
- Demonstrated on nonlinear probabilistic inversion of airborne electromagnetic; enables analysis of more than 100000 1D soundings per second.

Corresponding author: Thomas Mejer Hansen, `tmeha@geo.au.dk`

Abstract

The solution to a probabilistic inverse problem is the posterior probability distribution for which a full analytic expression is rarely possible. Sampling methods are therefore often used to generate a sample from the posterior. Decision-makers may be interested in the probability of features related to model parameters (for example existence of a pollution or the cumulative clay thickness) rather than the individual realizations themselves. Such features and their associated uncertainty, are simple to compute once a sample from the posterior distribution has been generated. However, sampling methods are often associated with high computational costs, especially when the prior and posterior distribution is non-trivial (non-Gaussian), and when the inverse problem is non-linear. Here we demonstrate how to use a neural network to directly estimate posterior statistics of any continuous or discrete feature of the posterior distribution. The method is illustrated on a probabilistic inversion of airborne EM data, where the forward problem is nonlinear and the prior information is non-Gaussian. Once trained the application of the network is fast, with results similar to those obtained using much slower sampling methods.

1 Introduction

A key challenge in geoscience is that of combining different kinds of geo-information into one geo-model, typically describing the subsurface. This information can be direct information about geological processes, spatial variability, or it can be indirect information from measurements of properties related to the subsurface, such as geophysical data. Ideally, when such a geo-model has been established, one should be able to quantify information about specific features related to the geo-model, consistent with all information.

This integration of geo-information is typically solved using inverse problem theory (Tarantola & Valette, 1982a; Menke, 2012). Fast deterministic methods exist and have been widely used. For such methods, the goal is to obtain one optimal model, such as the simplest possible model, consistent with available information, typically in the form of observed data (Tikhonov, 1963; Menke, 2012; Constable et al., 1987). In practice, in part due to noise on data and model nonlinearities and imperfections, infinitely many models exist that will be consistent with data, and the deterministic approach can in general not account properly for such uncertainty.

Probabilistic inversion methods can, in principle, take into account arbitrarily complex information, and integrate the information into one consistent model, as given by the posterior probability distribution. A full analytic expression of the posterior distribution is rarely possible. Instead, sampling methods can be used to generate a sample of the posterior, which is a collection of realizations of the posterior distribution. From such a sample, the posterior statistics of any feature related to the model parameters can be computed. The probabilistic approach is therefore ideal for decision-makers for uncertainty quantification, as it allows probabilistic analysis and risk assessment consistent with available information.

The main obstacle to applying the probabilistic methodology in practice is that sampling methods are computationally very demanding (Hastings, 1970; Mosegaard & Tarantola, 1995). Sampling-based methods typically require both sampling or evaluation of a prior model, and evaluation of the physical forward response(s), many times.

One approach for reducing the computational requirements is to make use of fast approximate forward modeling. This can be related to using simplified 1D forward modeling as opposed to 3D forward models, or by using approximate physical models, which leads to modeling errors that should be accounted for (Hansen et al., 2014; Madsen & Hansen, 2018; Köpke et al., 2018).

Machine learning algorithms, which are fast to evaluate once trained, have also been used to approximate the forward modeling (Hansen & Cordua, 2017; Conway et al., 2019; Moghadas et al., 2020; Bording et al., 2021). Unsupervised machine learning methods, for example Generative adversarial neural networks (GANs), have been used more generally as a means of representing features in a prior dataset; once trained, these provide an efficient means of rapidly generating many prior realizations (Mosser et al., 2017; Laloy et al., 2018; Mosser et al., 2020).

Attempts have also been made to use machine learning methods to learn a mapping from data to model that can directly solve the inverse problem. Röth and Tarantola (1994) were among the first to solve an inverse problem in this way using a multilayer perceptron neural network, and demonstrated an application of inversion of reflection seismic data to obtain single estimates of 1D velocity profiles. Recently, several authors have further explored this approach for directly solving a geophysical inverse problem, making use of convolutional neural networks (Puzyrev & Swidinsky, 2019; Moghadas, 2020; Bai et al., 2020). A drawback of such methods is that, as in the deterministic solution of an inverse problem, they estimate only a single model, typically without accounting for uncertainty on geophysical data, and do not quantify the uncertainty on the predicted model parameters.

An important step towards finding probabilistic solutions to inverse problems using neural networks was made by Devilee et al. (1999) who considered training data sets consisting of realizations from the prior distribution and the corresponding forward simulated data with and without noise. They then used neural networks to learn a set of probabilistic estimators about each model parameter, including median and equidistant histogram estimators. Meier et al. (2007) extended this work and used mixture density network (MDN) to estimate the parameters of a Gaussian mixture model representing a parametric distribution approximating the 1D marginal posterior distribution, and applied it to the problem of estimating global crustal thickness maps and comparing to results obtained using a Monte Carlo based sampling method. Several other applications of MDN to approximate the posterior distribution, for different geophysical problems, have followed (Shahraeeni & Curtis, 2011; de Wit et al., 2013; Earp & Curtis, 2020; Earp et al., 2020).

Zhang and Curtis (2020a) argue that it may be problematic to apply such MDN's for higher dimensional inverse problems, and suggest to use variational inference (Blei et al., 2017) to estimate the mean and standard deviation of the (non-Gaussian) posterior distribution in a seismic tomographic inverse problem. This method has been developed further for variational full waveform inversion and tomographic inversion using normalizing flows (Zhang & Curtis, 2020b; Zhao et al., 2022). In all these cases a uniform prior was assumed.

Attempts have also been made to use so-called invertible neural networks to simultaneously estimate both the forward and inverse mapping between data and model parameters (Ardizzone et al., 2018). This approach, which has recently been applied to geophysical data by Zhang and Curtis (2021), allows the generation of multiple realizations of the posterior distribution, from which properties of the posterior distribution can be estimated, although constructing invertible neural networks involves more work than traditional neural networks and involves compromises related to the flexibility of the network.

Here we present a method where the goal is not primarily to estimate the marginal 1D posterior distribution (as in works based on Meier et al. (2007)). Instead, we propose and demonstrate a machine learning-based method that provides direct estimates of any desired statistical property (continuous or discrete) of the posterior distribution, including any feature or property that can be computed from realizations of an, in prin-

ciple arbitrarily complex, prior model. This is done without generating realizations of the posterior distribution.

Following Devilee et al. (1999) and Meier et al. (2007) we construct a finite size training data set, representing the information available in any probabilistic formulation of the inverse problem, namely prior information and information about the forward model and the noise.

This is then used to train a neural network whose output parameterizes any desired statistical property of the posterior distribution for which a log-likelihood can be computed. These properties can for example represent a Gaussian, generalized Gaussian, log-normal, or a mixture model distribution, representing continuous model parameters. The output can also refer to the posterior probability of defined classes of model features of discrete model parameters. The neural network is designed to ensure that the estimated statistical properties of the posterior are similar to the same statistics derived from a sample of the posterior. Given a suitable training set the method provides accurate information regarding the specific properties of the posterior distribution that are of scientific interest in a given problem at a fraction of the time used by traditional sampling-based approaches.

In the following the method is first presented for probabilistic inverse problems in general, which can be considered as a generalization of the ideas proposed by Devilee et al. (1999) and later derived work e.g. Meier et al. (2007); Earp et al. (2020). Then, we demonstrate the method, applying it to non-linear probabilistic inversion of airborne electromagnetic data with non-Gaussian prior models of varying complexity. We show the neural network approach can be used to accurately estimate statistical properties of the posterior, related to both discrete and continuous model parameters, using regression and classification networks. The results are compared to results obtained by calculating the same statistical properties from a sample of the posterior obtained using the extended rejection sampler.

2 Method

Let $\mathbf{m} = [m_1, m_2, \ldots, m_{N_M}]$ represent N_M model parameters that define some properties of a system, such as for example physical properties of a geo-model. $\mathbf{m}$ is typically represented on a grid in a Cartesian or spherical coordinate system. For example, $\mathbf{m}$ might represent geophysical properties such as resistivity, velocity, or any other geological/geophysical/geochemical parameter.

A key issue in geosciences is how to infer information about $\mathbf{m}$ from different types of available information, such as geological expert knowledge, geophysical data, well log data, etc. This is generally referred to as an inverse problem. Tarantola and Valette (1982b) describe the inverse problem as a problem of probabilistic conjunction of information. Available information about $\mathbf{m}$ is described in the form of probability densities and then combined using conjunction of information to obtain a single probability density that describes the combined information. For example, consider a case when a specific type of information about structural properties is quantified by $\rho(\mathbf{m})$, and that information from observed electromagnetic data and well logs is quantified through $L(\mathbf{m})$. Then the conjunction of this information is given by the posterior probability distribution $\sigma(\mathbf{m})$, which, under the assumption that the individual types of information have been obtained independently, is given by

$$\sigma(\mathbf{m}) \;\propto\; \rho(\mathbf{m}) \cdot L(\mathbf{m}). \tag{1}$$

In other words, the conjunction of the independent information is proportional to the product of probability densities describing each independent set of information. The likelihood $L(\mathbf{m})$ is a measure of how well the data $\mathbf{d}$ computed from a specific model matches

observed data $\mathbf{d}_{obs}$ given noise with a specified probability distribution. Noise-free data can be computed by evaluating the forward model

$$\mathbf{d} = g(\mathbf{m}), \tag{2}$$

where g is a non-linear operator that maps the model parameters into data. g typically refers to some numerical algorithm solving some physical equations (such as Maxwell's equations).

The probabilistic inverse problem is then to infer information about $\sigma(\mathbf{m})$, which contains the combined information of, for example, both structural prior information, through the prior $\rho(\mathbf{m})$, and information from observed geophysical data, through $L(\mathbf{m})$.

A general approach (that allows using a non-linear forward model and non-Gaussian prior) for solving probabilistic formulated inverse problems is use of sampling methods to sample the posterior distribution, Eqn. 1, (Metropolis et al., 1953; Hastings, 1970; Geman & Geman, 1984; Green, 1995; Mosegaard & Tarantola, 1995; Laloy & Vrugt, 2012; Hansen et al., 2013, 2016). Unfortunately, such sampling methods can be extremely computationally demanding, to the point where they cannot be practically applied. They rely on solving the forward problem, Eqn. 2, many (often millions of) times.

Some algorithms make implicit assumptions about the prior model, such as a layered subsurface (Malinverno, 2002; Sambridge et al., 2013), while others, such as the classical rejection sampler and Metropolis algorithm (Hastings, 1970) require that both the prior and likelihood can be evaluated. This typically leads to using relatively simple prior models.

The extended variations of the Metropolis algorithm (Mosegaard & Tarantola, 1995) and the rejection sampler (Hansen et al., 2016; Hansen, 2021) do not require that an analytical description of the prior exists, as evaluation of the prior is not needed. It is sufficient that an algorithm exists that can generate a realization from the prior. This opens up the possibility of using a variety of more complex prior models, based on for example geostatistical simulation-based methods (Hansen et al., 2008, 2012).

2.1 Properties related to geophysical model parameters.

The model parameters $\mathbf{m}$ typically refer to physical parameters (e.g. resistivity when dealing with electromagnetic (EM) data, or elastic properties when dealing with seismic data). In practice, decision makers may be more interested in related features, or specific questions, such as "What is the chance of penetrating a specific lithology when drilling'?" (Scales & Snieder, 1997). Such features or occurences of events will be referred to through $\mathbf{n}$.

In general the relation between $\mathbf{m}$ and $\mathbf{n}$ can be complex and is formally described by a joint prior distribution $\rho(\mathbf{m}, \mathbf{n})$. This can for example be the case if $\mathbf{n}$ refers to subsurface lithology, and $\mathbf{m}$ to a geophysical property. This has been widely studied in the inversion of reflection seismic data, where information about geophysical properties is often assumed dependent on lithology, such that $\rho(\mathbf{m}, \mathbf{n}) = \rho(\mathbf{n})\rho(\mathbf{m}|\mathbf{n})$ (Bosch et al., 2010; Grana & Della Rossa, 2010; Rimstad et al., 2012). A more general formulation of Eqn. 1, describing information on both $\mathbf{m}$ and $\mathbf{n}$ is then

$$\sigma(\mathbf{m}, \mathbf{n}) \;\propto\; \rho(\mathbf{m}, \mathbf{n}) \cdot L(\mathbf{m}, \mathbf{n}), \tag{3}$$

given the available joint prior information, the forward model, and the noise. The corresponding forward problem, generalizing Eqn. 2, takes the form

$$\mathbf{d} = g(\mathbf{m}, \mathbf{n}). \tag{4}$$

Sometimes the relation between $\mathbf{m}$ and $\mathbf{n}$ is so simple that $\mathbf{n}$ can be computed from $\mathbf{m}$ through a mapping function $\mathbf{n} = h(\mathbf{m})$. For example, $\mathbf{n}$ can refer to the volume of

a reservoir (a scalar) obtained from a high dimensional set of geophysical model parameters $\mathbf{m}$. Or, $\mathbf{n}$ can refer to the cumulative thickness of layers with a resistivity ($\mathbf{m}$) above some threshold. Another example is when $\mathbf{m}$ refers to properties of a groundwater model. Then flow modeling based on a set of realizations from the posterior, can be used to propagate uncertainties into for example, the arrival time of polluted groundwater ($\mathbf{n}$) at a specific location (Vilhelmsen et al., 2019). Such a focus on related features and properties derived from the posterior distribution, rather than the posterior distribution over the geophysical parameter $\sigma(\mathbf{m})$ itself, is discussed by Scheidt et al. (2015).

The sampling algorithms described above can be used to generate a sample from $\sigma(\mathbf{m}, \mathbf{n})$ from which statistical analysis of any feature related to $\sigma(\mathbf{m}, \mathbf{n})$ can be computed. This is computationally demanding, and in many cases decision-makers are more interested in the statistical analysis of features and properties of the posterior distribution rather than the actual realizations.

Here a method is proposed that allows direct computation of properties and features of $\sigma(\mathbf{m}, \mathbf{n})$, using a neural network trained on a data set representing a sample of known information (including the prior, forward, noise and modeling errors), without ever generating realizations of $\sigma(\mathbf{m}, \mathbf{n})$. The approach follows the basic strategy suggested by Devilee et al. (1999), and consists of two steps: A) construction of a training data set, and B) construction and training of a neural network. This is done once. Then, the trained machine learning algorithm can be applied, very efficiently to compute desire properties of the posterior distribution, for potentially many sets of observed data.

2.2 A: Construction of training data set

Eqn. 4 describes the forward problem of computing noise free data. The forward problem describing simulation of data including noise, $\mathbf{d}_{sim}$ is

$$\mathbf{d}_{sim} \;=\; g(\mathbf{m}, \mathbf{n}) + r(\mathbf{m}, \mathbf{n}) \;=\; \mathbf{d} + r(\mathbf{m}, \mathbf{n}), \tag{5}$$

where $r(\mathbf{m}, \mathbf{n})$ is a realization of an assumed noise model. Often geophysical data $\mathbf{d}$ depends only directly on the physical parameters, in which case $g(\mathbf{m}, \mathbf{n}) = g(\mathbf{m})$.

Let $\mathbf{M}^* = [\mathbf{m}^{1*}, \mathbf{m}^{2*}, ..., \mathbf{m}^{N_T*}]$ and $\mathbf{N}^* = [\mathbf{n}^{1*}, \mathbf{n}^{2*}, ..., \mathbf{n}^{N_T*}]$ represent N_T realizations of $\rho(\mathbf{m}, \mathbf{n})$. Let $\mathbf{D}^* = [\mathbf{d}^{1*}, \mathbf{d}^{2*}, ..., \mathbf{d}^{N_T*}]$ represent the corresponding N_T noise free data, obtained by evaluating Eqn. 4. Finally let $\mathbf{D}^*_{sim} = [\mathbf{d}^{1*}_{sim}, \mathbf{d}^{2*}_{sim}, ..., \mathbf{d}^{N_T*}_{sim}]$ represent N_T corresponding realizations of simulated noisy data, following Eqn. 5. This constitutes a training data set

$$\mathbf{T} = [\mathbf{N}^*; \mathbf{M}^*; \mathbf{D}^*; \mathbf{D}^*_{sim}], \tag{6}$$

that can be obtained by 1) sampling the prior, 2) solving the forward problem, 3) simulation of the noise.

The sample $\mathbf{T}$ in Eqn. 6 represents the available information (prior, physics of the forward model, noise) in so far as it can be represented by a finite sample of size N_T. The larger the sample, the more complete the representation of the available information.

2.2.1 Infinite training data

Consider first the hypothetical limiting case when $\mathbf{T}$ is infinitely large ($N_T \to \infty$). In this case, $\mathbf{T}$ represents not just a subset of the available information, but all available information. The full probability distribution over any sets of parameters $\mathbf{n}$, $\mathbf{m}$, $\mathbf{d}$, and $\mathbf{d}_{sim}$ can be fully reconstructed from $\mathbf{T}$. Say some data have been measured as $\mathbf{d}_{obs}$. The corresponding inverse problem can then be solved simply by locating all the sets of models $[\hat{\mathbf{m}}, \hat{\mathbf{n}}] = ([\hat{\mathbf{m}}^{1*}, \hat{\mathbf{m}}^{2*}, ...], [\hat{\mathbf{n}}^{1*}, \hat{\mathbf{n}}^{2*}, ...])$ in $\mathbf{T}$ for which $\mathbf{d}^{i*}_{sim} = \mathbf{d}_{obs}$. $[\hat{\mathbf{m}}, \hat{\mathbf{n}}]$ will then represent a sample of $\sigma(\mathbf{m}, \mathbf{n})$ consisting of all possible realizations.

The goal here is not to generate realizations of the posterior, but instead to compute statistical properties of the posterior. In other words, given a sample $\hat{\mathbf{n}}$ of the posterior, $\sigma(\mathbf{n})$, the goal is to compute parameters Θ that define a desired statistical property of $\sigma(\mathbf{n})$. For example, if $\mathbf{n}$ refers to a discrete parameter with N_o possible outcomes, then $\Theta = [\theta_1, ..., \theta_{N_o}]$ could refer to the probability of realizing each possible outcome. If $\mathbf{n}$ refers to a continuous parameter, $\Theta = [\theta_0, \mathbf{C}_\theta]$ could represent the mean and covariance of a multivariate Gaussian distribution. $\Theta = [\theta_0, \theta_1, \theta_2]$ could represent the mean, variance and power of a generalized 1D Gaussian distribution. $\Theta = [\theta_0]$ could represent the rate of a Poisson distribution. $\Theta = [\theta_0, \theta_1]$ could represent a Binomial distribution.

The optimal values of Θ, given a sample $\hat{\mathbf{n}}$ of $\sigma(\mathbf{n})$, can be found maximizing the likelihood, L_Θ that each realization of the posterior is a realization of the probability distribution (defined by the parameter Θ)

$$L_\Theta = f(\hat{\mathbf{n}}|\Theta) = \prod_{i=1}^{N_\sigma} f(\hat{\mathbf{n}}^{i*}|\Theta), \tag{7}$$

where N_σ is the number of realizations of $\hat{\mathbf{n}}$. Maximization of Eqn. 7 is equivalent to minimizing the negative log-likelihood (which we refer to as the loss J_Θ):

$$J_\Theta = -\log(\prod_{i=1}^{N_\sigma} f(\hat{\mathbf{n}}^{i*}|\Theta)) \tag{8}$$

$$= -\sum_{i=1}^{N_\sigma} \log(f(\hat{\mathbf{n}}^{i*}|\Theta)). \tag{9}$$

Minimization of the loss function, Eqn. 9, can be used to obtain estimates of Θ representing a desired statistical property of $\sigma(\mathbf{n})$.

2.2.2 Finite training data

Use of an infinite training dataset is obviously unrealistic. Instead, we use a finite-sized training data $\mathbf{T}$ to design and train a neural network to estimate Θ directly from realizations of simulated data including noise $\mathbf{d}^{i*}_{sim}$. It is trained by maximization of the probability that $\mathbf{n}^{i*}$ are realizations of the chosen probability distribution, described by Θ, which are the result of evaluating the neural network $\mathbf{d}^{i*}_{sim} \mapsto \Theta$.

To achieve this, we make use of the fact that a particular feature or derived property in the training data set, $\mathbf{n}^{i*}$, is a realization of the posterior distribution $\sigma(\mathbf{n})$ one would get from probabilistic inversion of $\mathbf{d}^{i*}_{sim}$, as discussed above in 2.2.1.

If enough training data are available and a neural network complex enough to encompass the available information can be trained, then the network will estimate the statistical parameters Θ characterizing the desired properties of the posterior distribution $\sigma(\mathbf{m}, \mathbf{n})$.

2.3 B: Construct and train a neural network to estimate relevant statistics of $\boldsymbol{\sigma}(\mathbf{m}, \mathbf{n})$

In principle any machine learning method capable of regression and/or classification, such as regression trees and support vector machines (Bishop et al., 1995), can be used to estimate the mapping $\mathbf{d}^*_{sim} \mapsto \Theta$ which after training can be used to evaluate $\mathbf{d}_{obs} \mapsto \Theta$. Here we use make use of a fully connected artificial neural network. The presented approach builds on earlier work by Röth and Tarantola (1994), Devilee et al. (1999) and Meier et al. (2007).

A neural network is arranged into a number of layers, each consisting of a number of neurons. Each neuron has two adjustable parameters, the weight w, and the bias b, as well as an activation function Ψ. All neurons in one layer are fully connected to all neurons in the following layer. The input for a neuron (except for the first layer where the input is $\mathbf{d}^*_{sim}$) is the output of the neurons in the previous layer, and the output y_j of a neuron in response to inputs x_i, is given by

$$y_j = \Psi\left(\sum_i (w_i * x_i) + b\right)$$

For a specific network, with specified values for the weights and biases, one can compute the output, given some input, simply by evaluating the neurons layer by layer, starting from the input layer. See e.g. Bishop et al. (1995) for more details. To learn a particular mapping, the parameters of the neural network are adjusted to minimize a loss function that quantifies the performance of the network based on the training dataset. The choice of loss function is key to how the output of the neural network can be interpreted.

2.3.1 The structure of the neural network

A neural network can be described in terms of an input layer, the central inner part of the neural network (which can consist of many layers, referred to as hidden layers), and an output layer.

The input layer here represents the training data, which include noise, and consists of N_d nodes. The output layer has N_θ nodes representing the statistical parameters describing a distribution characterizing the features or properties of the posterior distribution that one wishes to predict.

The inner part of the network can be either simple or complex, and it can consist of either (fully) connected layers, convolutional layers, or combinations of these and other types of layers depending on the application. Here a fully connected neural network is considered as it has been demonstrated that such a neural network, with at least one hidden layer, can approximate any continuous function with arbitrary accuracy, when the number of hidden units is large enough (Hornik et al., 1990).

2.3.2 The loss function

When a neural network is trained using the training data set, its free parameters (the weight and bias of each node for a fully connected network) are adjusted to minimize a specific loss function. In the present case, the training data set consists of (when properties of $\sigma(\mathbf{n})$ are of interest) $\mathbf{T} = [\mathbf{N}^*; \mathbf{D}^*_{sim}]$. The goal is to estimate $\mathbf{d}^*_{sim} \mapsto \Theta$ rather than simply $\mathbf{d}^*_{sim} \mapsto \mathbf{n}$.

This is achieved by constructing a loss function where the unknown parameters Θ describe statistical properties of the desired probability distribution, Eqn. 7, whose parameters can be found by minimizing the loss function, Eqn. 9. The key here is to use a loss function that represents the negative log-likelihood of the probability distribution whose parameters Θ one wishes to estimate.

At each iteration of training the neural network, the loss is computed by applying the following steps for each dataset $T^i = [\mathbf{n}^{i*}, \mathbf{d}^{i*}_{sim}]$ in the training data set $\mathbf{T}$:

1. Evaluate the network using $\mathbf{d}^{i*}_{sim}$ as input. This provides as output an estimate $\hat{\Theta}_i$
2. Evaluate the corresponding loss, J^i, as $J^i = -\log(f(\mathbf{n}^{i*}|\hat{\Theta}_i))$.

The total loss is then given by

$$\mathbf{J} = \sum_{i=1}^{N_T} J^i. \tag{10}$$

$\mathbf{n}^{i*}$ is a realization of $\sigma(\mathbf{n})$, given the data $\mathbf{d}^{i*}_{sim}$, and therefore, minimizing the loss in Eqn. 10 leads to estimates of statistical parameters Θ that describe $\sigma(\mathbf{n})$, in the same manner as would minimizing Eqn. 9 given a sample, $\hat{\mathbf{n}}$, of $\sigma(\mathbf{n})$. The difference is that the proposed method achieves this without the need to first realize the sample $\hat{\mathbf{n}}$ of $\sigma(\mathbf{n})$.

Minimizing the loss function thus maximizes the probability that each $\mathbf{n}^{i*}$ can be seen as a realization of the probability distribution whose parameters Θ_i are the result of evaluating the neural network $\mathbf{d}^{i*}_{sim} \mapsto \Theta_i$.

In general, $\mathbf{n}$ (and/or $\mathbf{m}$) can refer to a continuous parameter (such as velocity, resistivity, temperature, or related properties) or a discrete parameter (such as lithology type and event type). Continuous model parameters typically lead to a regression type problem, whereas discrete model parameters lead to a classification problem.

2.3.2.1 Continuous model parameters - regression Here we first consider the case when $\mathbf{n}$ represents continuous parameters. Say we wish to estimate the mean and covariance, $\hat{\Theta}_0$ and $\hat{\mathbf{C}}_\theta$, of the posterior distribution $\sigma(\mathbf{n})$ given a set of observed data $\mathbf{d}_{obs}$.

Assume a neural network exists that outputs a set of parameters describing $\Theta = [\hat{\Theta_0}^i, \hat{\mathbf{C}^i}_\theta]$, given the input $\mathbf{d}^i_{sim}$. The likelihood that a set of parameters from the training dataset $\mathbf{n}^{i*}$ is a realization from the multivariate Gaussian distribution $\mathcal{N}(\hat{\Theta_0}^i, \hat{\mathbf{C}^i}_\theta)$ as obtained from evaluating the neural network using $\mathbf{d}^{i*}_{sim}$ as input, is given by

$$f(\mathbf{n}^{i*}|\hat{\Theta_0}^i, \hat{\mathbf{C}^i}_\theta\ ) = k_C\ \exp(-0.5\ (\mathbf{n}^{i*} - \hat{\Theta_0}^i)^T\ \hat{\mathbf{C}^i_\theta}^{-1}\ (\mathbf{n}^{i*} - \hat{\Theta_0}^i)), \tag{11}$$

where $k_C = ((2\pi)^{N_d}|\hat{\mathbf{C}^i_\theta}|)^{-.5}$ is a normalization factor. The corresponding loss function J^i is the negative log-likelihood loss function, that is

$$J^i = -\log(f(\mathbf{n}^{i*}|\hat{\Theta_0}^i, \hat{\mathbf{C}^i}_\theta)) \tag{12}$$

$$= -0.5\ (\mathbf{n}^{i*} - \hat{\Theta_0}^i)^T\ \hat{\mathbf{C}^i_\theta}^{-1}\ (\mathbf{n}^{i*} - \hat{\Theta_0}^i) \tag{13}$$

The total average loss is then given by Eqn. 10.

Any machine learning method that minimizes this loss function, will lead to a neural network that provides an estimate of the parameters of interest, here $\mathbf{\Theta} = [\hat{\Theta}_0, \hat{\mathbf{C}}_\theta]$, that are computed directly without ever computing realizations of $\sigma(\mathbf{n})$.

To represent the posterior mean and full covariance, given N_m model parameters, an output layer of $N_\Theta = N_m + N_m^2$ nodes must be used. If only the posterior mean and variance are estimated, an output layer of $N_\Theta = N_m + N_m$ nodes is needed. If only the posterior mean is of interest an output layer of $N_\Theta = N_m$ nodes is needed and minimizing Eqn. 13 is then similar to minimizing the widely used mean squared error loss function (Bishop et al., 1995), as utilized for example in e.g. Röth and Tarantola (1994).

Recall, that the above scheme does not impose any assumptions on either the prior or the posterior distribution which may be complex. The estimated mean and covariance are simply statistical parameters of the posterior distribution, that may or may not be useful for a specific use case.

The quality of the obtained estimate naturally depends on the complexity of the machine learning model used, and the size of the training data set, which will be considered in more detail in the application presented below.

Other statistical parameters of the posterior can be estimated by minimizing the appropriate log-likelihood function for the corresponding probability distribution.

For example, a 1D generalized probability distribution is defined by three parameters $\Theta = [\theta_1, \theta_2, \theta_3]$, and its probability distribution given by (Tarantola, 2005)

$$f(n^i|\Theta) = \frac{1}{2\theta_2\Gamma(1+1/\theta_3)} \exp\left(- \left(\frac{|n^i - \theta_1|}{\theta_2}\right)^{\theta_3}\right) \tag{14}$$

A 1D Gaussian mixture model based on a mixture of Nc 1D Gaussian distribution, as considered by e.g. Meier et al. (2007), is defined by $\Theta = [\mathbf{t}_1, \mathbf{t}_2, \mathbf{t}_3] = [t_1^1, ..., t_1^{Nc}, t_2^1, ..., t_2^{Nc}, t_3^1, ..., t_3^{Nc},]$, where $\mathbf{t}_1$ refer to the mean, $\mathbf{t}_2$ refer to the standard deviation of Nc Gaussian distribution, each with weight $\mathbf{t}_3$, and its probability distribution given by

$$f(n^i|\mathbf{t}_1, \mathbf{t}_2, \mathbf{t}_3) = \sum_{i=1}^{Nc} t_3^i \left(t_2^i\sqrt{2\pi}\right)^{-1} \exp\left(- 0.5\left(\frac{n^i - t_1^i}{t_2^i}\right)^2\right) \tag{15}$$

The corresponding log-likelihood of Eqn. 14 and 15 can trivially be obtained and used as a loss function in a neural network to estimate Θ. In principle, any statistical parameters, for which the associated log-likelihood can be computed, and used as a loss function, can be estimated using the proposed methodology.

2.3.2.2 Discrete model parameters - classification Say n_i represents a discrete parameter with N_o possible outcomes (classes). One's aim is then to estimate the posterior probability of each of the N_o classes given some data $\mathbf{d}_{obs}$.

Let $\theta_i^* = [{p_i}^{1*}, {p_i}^{2*}, ..., {p_i}^{N_o*}]$ represent the true probabilities of n_i^* belonging to a specific class. In practice the true probability of one (the correct) class will be one, and the others zero. Further $\hat{\theta}_i = [\hat{p_i}^1, \hat{p_i}^2, ..., \hat{p_i}^{N_o}]$ represent the corresponding predictions by the neural network of the probabilities of each class for a specific model parameter, n_i.

The likelihood of observing θ_i given $\hat{\theta}_i$ is then

$$f(\theta_i|\hat{\theta}_i) = \prod_{j=1}^{N_o} (\hat{p_i}^j)^{{p_i}^{j*}}. \tag{16}$$

The corresponding loss function J^i is then

$$J^i = -\log(f(\theta_\mathbf{i}|\hat{\theta}_\mathbf{i})) = -\sum_{j=1}^{N_o} p_i^{j*} \log(\hat{p}_i^j). \tag{17}$$

The choice of class probabilities $\hat{\theta}_\mathbf{i}$ that maximizes Eqn. 16 can be found by minimizing the negative log-likelihood given by the loss function, Eqn. 17, which is equivalent to the categorical cross-entropy between the two probability distributions (Bishop et al., 1995). Usually, the softmax activation is used for multi-class classification problems (and the sigmoid activation function for binary classification problems), as it forces all probabilities to be in the range 0 to 1, and ensures that $\sum_{j=1}^{N_o} \hat{p}_i^j = 1$, such that the output parameters can be interpreted as a probability. A neural network that estimates the mapping $\mathbf{d}_{sim}^i \mapsto \hat{\theta}^i$ by minimizing Eqn. 17, using the softmax activation function in the output layer, therefore locates the maximum-likelihood of Eqn. 16, which directly estimates $\sigma(p_i^*)$.

In other words, this method can be used to compute the posterior class probability of a discrete model parameter, without generating a sample of the posterior distribution.

To summarize, our proposed method involved first constructing a training data set (Eqn. 6) that represents (limited by the size of the used training data set) the known information (the prior, the forward, and the noise model), and specifically contains prior knowledge regarding any feature $\mathbf{n}$, that may be directly or indirectly related to the model parameters $\mathbf{m}$, about which one wishes to infer information. A neural network is then designed and trained by minimizing a specific loss function (that expresses the log-likelihood of the parameters Θ describing the probability distribution of desired features $\mathbf{n}$ that may be either continuous or discrete.

3 Results / Application to airborne EM data

The methodology described above is applied to the inversion of airborne electromagnetic (AEM) data. This inverse problem has been widely studied by deterministic linearized least-squares methods using both a 1D and 3D forward model (Christensen, 2002; Auken & Christiansen, 2004; Viezzoli et al., 2008; Cox et al., 2010; Grayver et al., 2013; Auken et al., 2014).

The full non-linear 1D inverse problem has also been addressed using Markov chain Monte Carlo (McMC) sampling methods, based on for example the reversible-jump sampling method relying on a prior model representing a 1D layered subsurface (B. J. Minsley, 2011; B. J. Minsley, Foks, & Bedrosian, 2021; Brodie & Sambridge, 2012). Hansen and Minsley (2019) proposed the use of extended Metropolis algorithm, also an McMC method, that allows the use of any prior model that can be sampled. The 1D nonlinear inverse EM problem leads to a non-trivial sampling problem, due to the existence of model equivalences (significantly different models lead to the same forward response). Sufficient sampling of the 1D posterior distribution of resistivity values, to obtain a limited set of independent realizations, may require hundreds of thousands of McMC iterations, and hence forward model evaluations. For a single sounding this may take at least 10 minutes per sounding, requiring access to supercomputers for application of real-world data sets (Foks & Minsley, 2020). Hansen (2021) proposed 1D probabilistic inversion based on the extended rejection sampler (using lookup tables, similar to $[\mathbf{N}^*,\mathbf{M}^*,\mathbf{D}^*]$) that rely on the construction of a large sample for the prior along with the forward responses (generated once). This is then used to generate independent realizations of the posterior distribution numerically more efficiently than is possible using Markov Chain based algorithms, and at the same time avoids issues related to model equivalences. This sampling approach is used for comparison below.

The size of airborne EM surveys is becoming larger, so the use of any of the inversion methods discussed above will lead to considerable computational demands. Currently, two major airborne EM surveys are being carried out. The AusAEM20 project, by Geoscience Australia, is expected to collect around 65000 flight-line-kilometer of data, leading to many hundreds of thousands of EM measurements (Howard, 2020). USGS has collected more than 43000 flight-line-kilometer data in the Mississippi Alluvial Plain, and another 25000 flight-line-kilometer is planned for 2021, leading to significantly more than 1.000.000 data points to be inverted in the Mississippi Alluvial Plain (B. J. Minsley, Rigby, et al., 2021).

3.1 AEM data from Morrill, Nebraska

As an example we consider the inversion of airborne electromagnetic (AEM) data from Morrill, Nebraska (Smith et al., 2010; Abraham et al., 2012). The same profile of data obtained at 451 locations along a West-East profile is used here as described in B. J. Minsley (2011). Each observed data set consists of 12 measurements (in-phase and quadrature measurements from 6 pairs of transmitter and receiver coils).

Three different types of prior models will be defined, that represent different information about the subsurface resistivities ($\mathbf{m}$) and related (both discrete and continuous) properties $\mathbf{n}$ at Morrill. For each of the three prior models considered, a unique posterior probability distribution exists. Various properties of the posterior distribution will be computed using the proposed machine learning method and compared to results obtained from a finite sample of the posterior distributions obtained using the extended rejection sampler with a lookup table of size $N_T = 2 \cdot 10^6$.

3.2 A priori models and noise

3.2.1 Parameterization

In this example, the subsurface is parameterized into 125 layers of $dz = 1$ m thickness. Prior models based on up to four sets of parameters, $\rho(\mathbf{m}, \mathbf{n}_1, \mathbf{n_2}, \mathbf{n_3})$ are considered.

Resistivity. $\mathbf{m} = [m_1, m_2, ..., m_{N_M}]$ represents the resistivity of each of the 125 layers.

Layer interface. $\mathbf{n_1}$ represents the existence of a sharp boundary between two neighboring layers ($n_{1i} = 0$ when there is no boundary and $n_{1i} = 1$ in case of a boundary). A sharp boundary is defined when two neighboring resistivity values differ more than 20%. $\mathbf{n_1}$ refers to 125 discrete parameters and can be directly computed from $\mathbf{m}$.

Thickness of highly resistive layer. $\mathbf{n_2}$ represent the cumulative thickness of resistivity values above 225 ohmm. $\mathbf{n_2}$, which can be directly computed from $\mathbf{m}$ using

$$\mathbf{n}_2 = \sum_{i}^{N_M} dz * I(m_i),$$

where $I(m_i) = 1$ when $m_i > 225$ ohmm, and $I(m_i) = 0$ when $m_i \leq 225$ ohmm. $\mathbf{n}_2$ refers to a single continuous parameter.

Lithology. $\mathbf{n_3}$ represents a category ('1', '2', and '3', representing three distinct lithologies) in each layer. $\mathbf{n_3}$ cannot be computed from $\mathbf{m}$, but $\mathbf{n_3}$ and $\mathbf{m}$ are linked through a conditional prior distribution $\rho(\mathbf{m}|\mathbf{n_3})$ (see example below). $\mathbf{n_3}$ refers to 125 discrete parameters with 3 possible outcomes.

For brevity, all model parameters combined will be referred to as $\mathbf{p} = [\mathbf{m}, \mathbf{n}_1, \mathbf{n}_2, \mathbf{n}_3]$. To illustrate the potential of the method 3 different non-Gaussian prior models are considered that vary in complexity and information content.

3.2.2 Prior information

$\rho_A(\mathbf{p}) = \rho_A(\mathbf{m}, \mathbf{n_1}, \mathbf{n_2})$*, a uniform prior model.* $\rho_A(\mathbf{p})$ represents a choice of independence between model parameters, $\rho_A(m_i, m_j) = \rho_A(m_i)\rho_A(m_i)\ \forall(i, j)$. The resistivity of each resistivity model parameter is assumed to be log-uniform distributed in the range $\mathcal{U}[2, 2800]$ ohmm. This is the least informative prior model considered. 11 independent realizations of $\rho_A(\mathbf{m}, \mathbf{n_1})$ are shown in Figure 1a.

$\rho_B(\mathbf{p}) = \rho_B(\mathbf{m}, \mathbf{n_1}, \mathbf{n_2})$*, Discrete layered model.* $\rho_B(\mathbf{p})$ represents a layered subsurface consisting of 1 to 8 layers (uniformly distributed), each with a constant resistivity. The resistivity in a specific layer is assumed to be log-uniform distributed in the range $\mathcal{U}_{\updownarrow}[2, 2800]$ ohmm.

A realization $\mathbf{p}^*$ of $\rho_B(\mathbf{p})$ is generated by first choosing the number of layers as a random number, Nl, between 1 and 8. Then $Nl - 1$ layer interfaces are randomly selected from a uniform distribution of $\mathcal{U}[0, 125]$ m. Then the resistivity within each layer

is realized from a uniform distribution $\mathcal{U}_{\updownarrow}[2, 2800]$ ohmm. This type of prior model is similar to the transdimensional prior considered by (B. J. Minsley, 2011). 11 independent realizations of $\rho_B(\mathbf{m}, \mathbf{n_1})$ are shown in Figure 1b.

$\rho_C(\mathbf{p})$, Trimodal mixture Gaussian. $\rho_C(\mathbf{p})$ represents a subsurface with three possible lithologies ('1', '2' and '3') each with a distinct resistivity distribution. See discussion about the prior geological knowledge in Morrill in Abraham et al. (2012) and Hansen and Minsley (2019).

To sample $\rho_C(\mathbf{p}) = \rho_A(\mathbf{m}, \mathbf{n_1}, \mathbf{n_2}, \mathbf{n_3})$, first a realization of $\rho_C(\mathbf{n}_3)$ is generated as $\rho_C(\mathbf{n_3}^*)$, which represents an example of the distribution of the lithologies. This is achieved by generating a realization of a multivariate normal distribution with a Gaussian type covariance model with a range of 30 m, followed by a simple truncation to obtain 40% of lithology A, 40% of lithology B, and 20% of lithology C. Then a realization of the resistivity $\mathbf{m}^*$ is generated, conditional to the lithology type from $\rho_C(\mathbf{m}|\mathbf{n}_3^*)$. The resistivity, within each lithology, is generated as a realization of a multivariate normal distribution in $\log_{10}$-resistivity space with a range of 30 m, a specific mean, m_0 and standard deviation, m_{std}. For lithology '1', $m_0 = 1.1$ and $m_{std} = 0.14$. For lithology '2', $m_0 = 2$ and $m_{std} = 0.2$. For lithology '3', $m_0 = 2.75$ and $m_{std} = 0.25$. Finally, $\mathbf{n_1}^*$ and $\mathbf{n_2}^*$ are computed from $\mathbf{m}^*$. In this way a realization $\mathbf{p}^* = [\mathbf{m}^*, \mathbf{n_1}^*, \mathbf{n_2}^*, \mathbf{n_3}^*]$ of $\rho_C(\mathbf{p})$ is generated. 11 independent realizations of $\rho_C(\mathbf{p})$ are shown in Figure 1b.

$\rho_C(\mathbf{p})$ is designed to reflect available information related to the subsurface at Morrill (Abraham et al., 2012; Hansen & Minsley, 2019). $\rho_A(\mathbf{p})$ and $\rho_B(\mathbf{p})$ are considered here to investigate how the proposed methodology reacts to a uniform (maximum entropy) prior such as $\rho_A(\mathbf{p})$, and a simple prior as $\rho_B(\mathbf{p})$.

3.2.3 Noise

The noise of the EM data is assumed to be independent uncorrelated zero-mean Gaussian noise, with a standard deviation of 5 ppm (parts per million) plus 5 percent noise relative to the noise-free data value. This is the same noise model as considered in previous works on the EM data from Morrill (B. J. Minsley, 2011; Hansen & Minsley, 2019; Hansen, 2021).

3.3 Sampling of the posterior distribution

For reference, the extended rejection sampler, with a lookup table of size $N_T = 2{\cdot}10^6$, is used to sample the posterior distribution, as detailed in Hansen (2021). 11 independent realizations of the three posterior distribtions ($\sigma_A(\mathbf{p})$, $\sigma_A(\mathbf{p})$, and $\sigma_C(\mathbf{p})$) are shown in Figures 1b,d,f.

The goal of the proposed machine learning approach is to directly compute statistical properties of the posterior distribution similar to obtaining the same statistical properties from a sample of the posterior obtained using sampling, such as shown in Figures 1b,d,f.

3.4 Neural network design

Two fully connected neural networks are designed to allow characterizing the 1D marginal posterior distribution of continuous and discrete parameters. The input layer, in both cases, consists of the observed data $\mathbf{d}_{obs}$, or simulated data with noise. For this specific case, it consists of 12 neurons.

The inner network is designed using a number of hidden layers, each with 40 neurons with the Exponential Linear Unit (ELU) activation function (Bishop et al., 1995). This inner part of the network should be complex enough that the desired mapping can

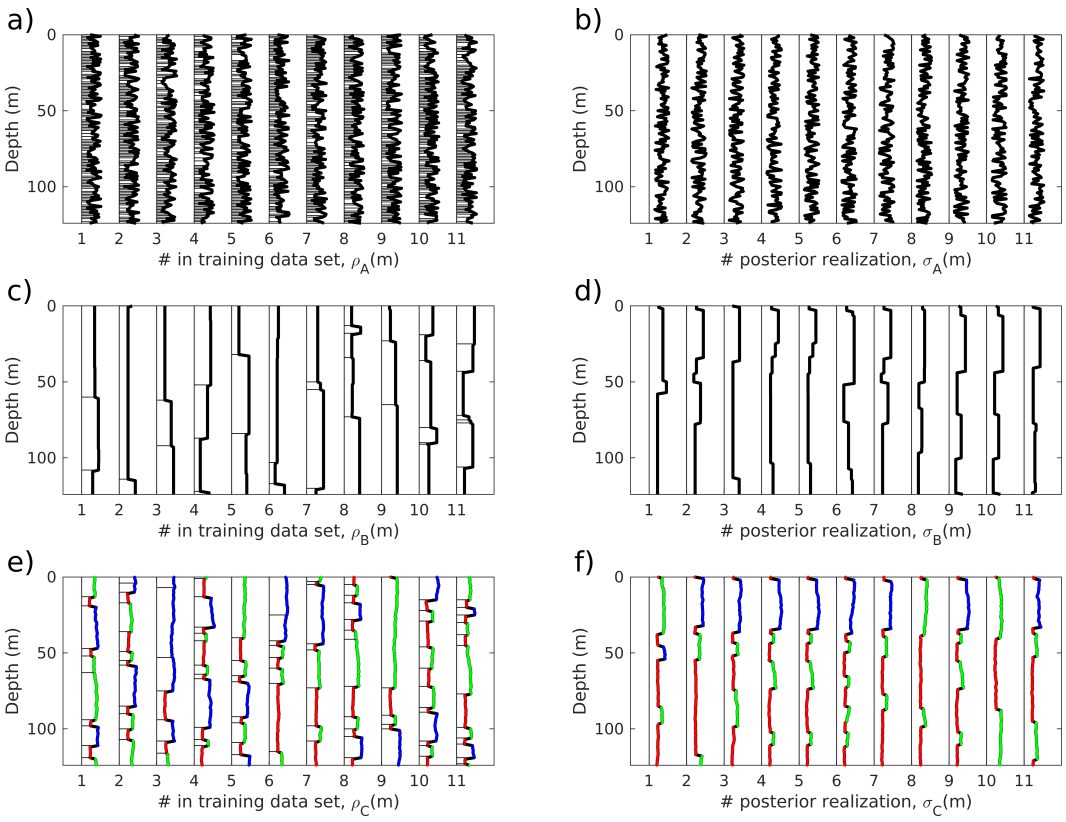

Figure 1. First 11 models from the lookup table for three prior models a) $\rho_a(\mathbf{m}, \mathbf{n_1})$, c) $\rho_b(\mathbf{m}, \mathbf{n_1})$,and e) $\rho_c(\mathbf{m}, \mathbf{n_1}, \mathbf{n_2})$, as well as 11 independent realizations from the posterior distribution obtained for the data at x=15km for a) $\sigma_a(\mathbf{m}, \mathbf{n_1})$, c) $\sigma_b(\mathbf{m}, \mathbf{n_1})$,and e) $\sigma_c(\mathbf{m}, \mathbf{n_1}, \mathbf{n_2})$. Thin black lines indicate the existence of a layer interface ($\mathbf{n}_1$). The thick line indicates variation in resistivity ($\mathbf{m}$). In c) the colors of the thick line represent lithology A (red), B(blue), and C (green) when defined.

be represented, but simple enough to avoid overfitting, as discussed also by (Meier et al., 2007). Network design is highly problem-dependent, and for the present problem, we found this network design provides results on par with, and in some cases better than, sampling-based approaches, while at the same time being relatively easy to optimize.

The choice of loss function, and to some extent the activation function, defines the specific property of the posterior distribution that will be estimated. This leads to two specific types of output layers for regression and classification type problems[1].

3.4.1 Regression type neural network

The first neural network type is designed to estimate parameters θ of a probability distribution describing the 1D marginal posterior distribution of a continuous parameter (such as $\mathbf{m}$ and $\mathbf{n_3}$). If N_θ is the number of parameters needed to describe a specific 1D distribution, then in total $N_{out} = N_\theta N_m$ neurons are needed in the output layer if the target is properties of $\sigma(\mathbf{m})$, and $N_{out} = N_\theta$ if the target is $\sigma(\mathbf{n_3})$.

3.4.2 Classification type neural network

The second neural network type is designed to estimate the posterior probability of possible classes for the discrete type model parameters $\mathbf{n_1}$ and $\mathbf{n_3}$, i.e. of $\sigma(\mathbf{n_1})$ $\sigma(\mathbf{n_3})$.

If the goal is to estimate the 1D marginal distribution of a discrete parameter with N_{cat} possible outcomes, this can be achieved by selecting an output layer with $N_{out} = N_m$ when $N_{cat} = 2$ (using a sigmoid activation function), and $N_{out} = N_{cat} N_m$ when $N_{cat} > 2$ (using the softmax activation function). As discussed above, using the cross-entropy loss function, Eqn. 17, will lead to direct estimation of the 1D posterior marginal probabilities in this case.

3.5 Network training

Using the nonlinear forward model and the noise model, a training data set of size $N_T = 1000000$ is constructed (one for each type of prior model) and used for training. Both networks are trained using 67% of the training data set, while 33% is reserved for validation. In both cases, the loss function is minimized using the Adam optimizer (Kingma & Ba, 2014) using a learning rate of 0.001, for a maximum of 2000 epochs. Early stopping is utilized which stops the training if the loss function evaluated on the validation data does not decrease for 50 epochs. This is done to avoid over-fitting, where the loss on the training data will decrease, but where the loss on the validation data increases. TensorFlow with Keras and TensorFlow-probability have been used to implement and train the neural networks (Abadi et al., 2015; Chollet, 2015; Dillon et al., 2017).

The two considered networks, and the training of the networks, only differ concerning the definition of the output layer (the number of nodes and activation function), the choice of loss function, and the chosen number of hidden layers.

3.6 Estimation of properties of $\sigma(\mathbf{m})$

First, properties related to the posterior distribution of resistivity, $\sigma(\mathbf{m})$, are considered.

[1] Example implementations of these two types of neural networks can be found at `http://github.com/cultpenguin/ip_and_ml/`.

3.6.1 Estimation of mean and standard deviation of $\sigma(\mathbf{m})$

A neural network is set up and trained to estimate the pointwise mean and standard deviation of $\sigma(\mathbf{m})$, using 8 hidden layers, by minimizing the loss function in Eqn 13.

Figures 2a-d shows the pointwise mean of the posterior distribution $\sigma_C(\mathbf{m})$ obtained using the machine learning approach with a training data set of size $N = [1000, 10000, 100000, 1000000]$, compared to the same statistics computed from a sample of the posterior obtained using the sampling method, Figure 2e. The corresponding standard deviation controls the transparency in the plot, with high transparency corresponding to high standard deviation[2]

It is clear from Figure 2a that using $N_T = 1000$ provides very poor results, as compared to the results obtained using sampling, Figure 2e. But even using $N_T = 10000$ leads to results close to the sampling-based results. For $N_T \geq 100000$ the quality of the direct estimates of the mean and standard deviation does not seem to differ much.

One notable difference when comparing Figure 2d (N_T=1000000) and 2e (sampling), is that sampling results in more small scale variability in the estimated parameters, as opposed to the more smooth result obtained using machine learning. The reason is simply that the sampling-based approach is based on inferring the statistics from a finite-sized sample of the posterior, whereas in the machine learning approach these statistics are estimated directly.

Figure 3 shows a comparison between the posterior mean (with standard deviation used to control transparency) obtained using the sampling approach and using the machine learning approach ($N_T = 1000000$), for $\sigma_A(\mathbf{m})$ (Figures 3a-b) and $\sigma_B(\mathbf{m})$ (Figures 3c-d) respectively.

$\rho_A(\mathbf{m})$ refers to the least informed prior model, and hence one should expect the least resolution in the corresponding posterior distribution. This is what can be seen in results from both the machine learning and the sampling approach, Figures 3a-b, where only the top high resistive layer is somewhat resolved.

While $\rho_B(\mathbf{m})$ is somewhat simpler than $\rho_C(\mathbf{m})$, the mean and standard deviation of the corresponding posterior distribution are rather similar, with most difference related to the posterior standard deviation (as illustrated by the transparency in Figures 2 and 3c-d).

A key point from Figures 2 and 3 is that the use of the machine learning based approach seems to provide results at least on par with the results obtained using sampling, when the goal is to estimate the mean and standard deviation of the (non-Gaussian) posterior distribution. This is the case using both informed and uninformed prior models.

3.6.1.1 Computational efficiency Figure 4 shows the computation time[3] needed to train the neural networks for the results presented in Figure 2. The training time increase with the size of the training data set, N_T. Both training and validation loss is reduced when N_T increases. It is also clear that the relative difference in loss decreases when comparing the use of $N_T = 100000$ to $N_T = 1000000$, to when comparing the use of $N_T = 1000$ to $N_T = 10000$. Hence, using $N_T > 100000$ leads to a substantial longer training time, but only to a minor loss reduction.

[2] The mean and standard deviation without transparency are available in the supplementary material.

[3] a workstation with an Intel Core(TM) i7-8700K CPU, Nvidia RTX 3090 GPU, and 64 Gb RAM was used

Once set up and trained, the prediction of the network is very fast. For all the networks presented above, the prediction time for all 451 data locations was less than 5ms. This means that more than 100000 soundings can be analyzed per second.

3.6.2 Estimation of multiple 1D properties of $\sigma(m_i)$

As described above, any parameter of a probability distribution for which a loss function can be described through Eqn. 10 can be estimated using the machine learning method. To demonstrate this, 4 independent networks have been trained to estimate properties (Θ) of the 1D marginal posterior distribtion $\sigma(m_i)$ given by a) a normal distribution (Eqn. 11, as in Figure 2), b) a generalized normal distribution (Eqn. 14), c) a mixture distribution based on two Gaussian distributions (Eqn. 14), and d) a mixture distribution based on three Gaussian distributions (Eqn. 14) . The loss functions used are the negative log-likelihood of the probability distribution in Eqns. 11, 14, and 15 respectively.

The number of parameters to estimate for the 4 cases, and hence neurons in the output layer, is $N_\theta = [2{*}N_m, 3{*}N_m, 2{*}N_m{*}N_c, 3{*}N_m{*}N_c] = N_\theta = [250, 375, 750, 1125]$, where N_c is the number of distributions in the mixture model.

Figure 5a shows the posterior 1D marginal distribution of resistivity values obtained using sampling, based on a finite set of realizations, obtained at x=15 km. One can clearly identify a bimodal to trimodal distribution at depth representing the three possible lithologies from the prior model $\rho_C(\mathbf{m})$ with different resistivity values.

Figures 5b-f, shows the probability distributions representing the estimated statistical properties of the 4 considered distributions. These distributions do not represent assumptions about the posterior distribution (which can be arbitrarily complex) but reflect the statistical properties one would get if the particular choice of distribution is used to represent a sample of the posterior.

If the goal is to compute a representation of the 1D posterior marginal distribution, as considered by (Meier et al., 2007; Shahraeeni & Curtis, 2011), then care should be taken to use a parameterizations for the chosen 1D distribution complex enough to allow describing the variability of the posterior. From Figure 5 it is evident that only in case using the mixture model with 3 Gaussian distributions, the estimated marginal probability density represents the actual 1D marginal posterior distribution well.

The statistical parameters of the posterior distribution which it is relevant to compute for a specific inverse problem, is naturally problem-dependent. This example nonetheless demonstrates that the machine learning methodology is capable of estimating parameters of different types of probability distributions, for which a probability density, and hence the corresponding loss function, can be computed.

3.7 Estimation of properties of $\boldsymbol{\sigma}(\mathbf{n_2})$

We consider the simpler problem of inferring information about a single continuous parameter, $\mathbf{n_2}$, representing the cumulative thickness of layers with a resistivity above 225 ohmm. The same neural network is used as considered above to estimate properties related to $\mathbf{m}$, except here only 4 hidden layers are used.

Figure 6 shows the mean of $\sigma_C(\mathbf{n_2})$ (black line), as well as the probability distribution reflecting the mean and standard deviation estimated using the machine learning approach for $N_T = [1000, 10000, 100000, 1000000]$ in figures 6a-d. The mean computed using the machine learning approach compares well to the mean obtained using sampling methods for $N_T \geq 100000$.

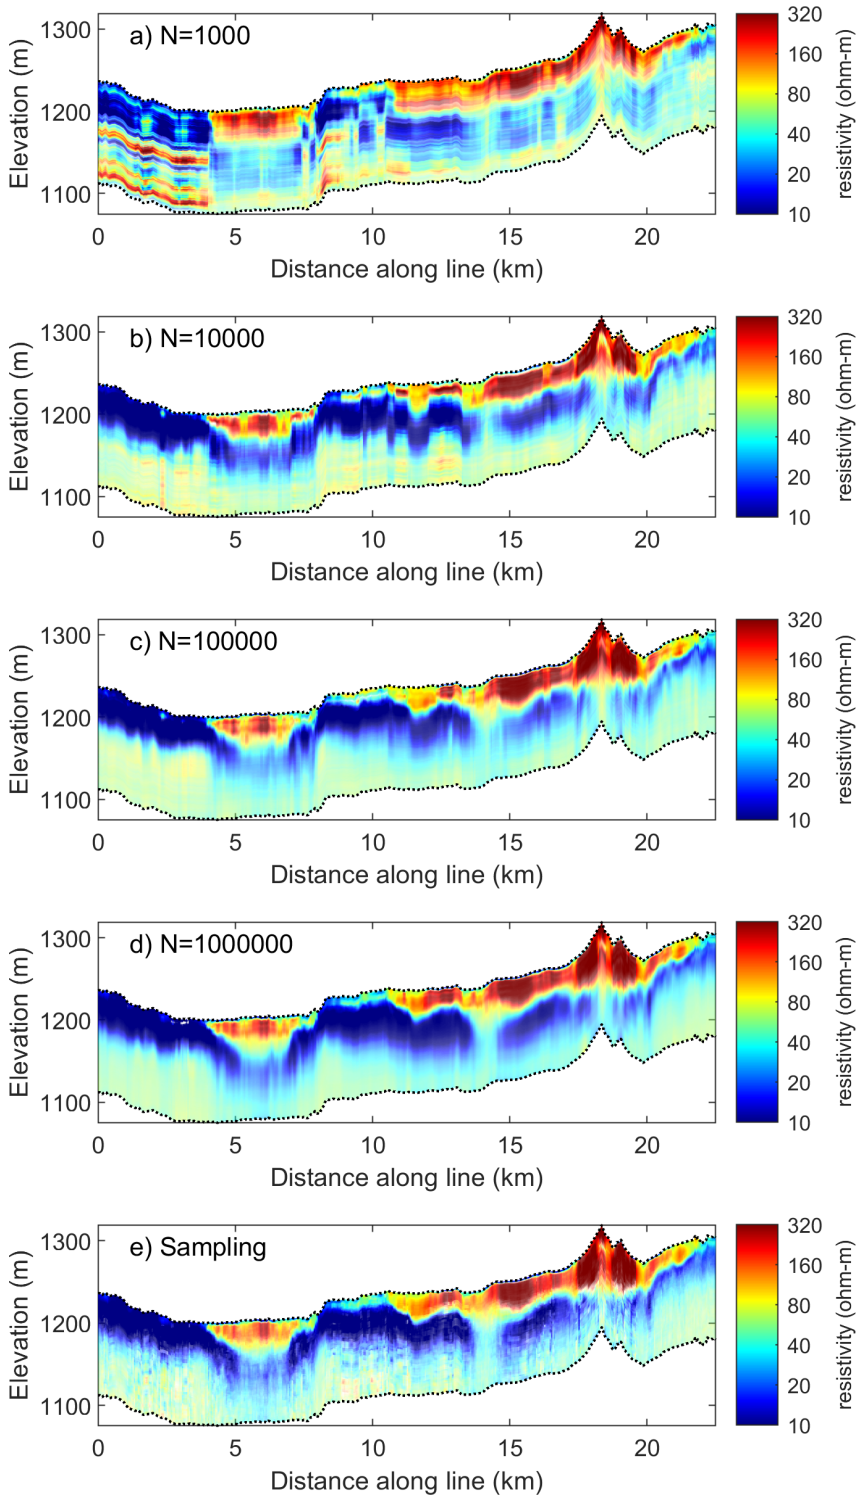

Figure 2. Pointwise mean obtained from $\sigma_C(\mathbf{m})$ obtained using machine learning based on a training data set of size a) 1000, b) 10000, b) 100000, and b) 1000000, and using e) the extended rejection sampler. Transparency based on pointwise posterior standard deviation. The mean and standard deviation without transparency are available in the supplementary material.

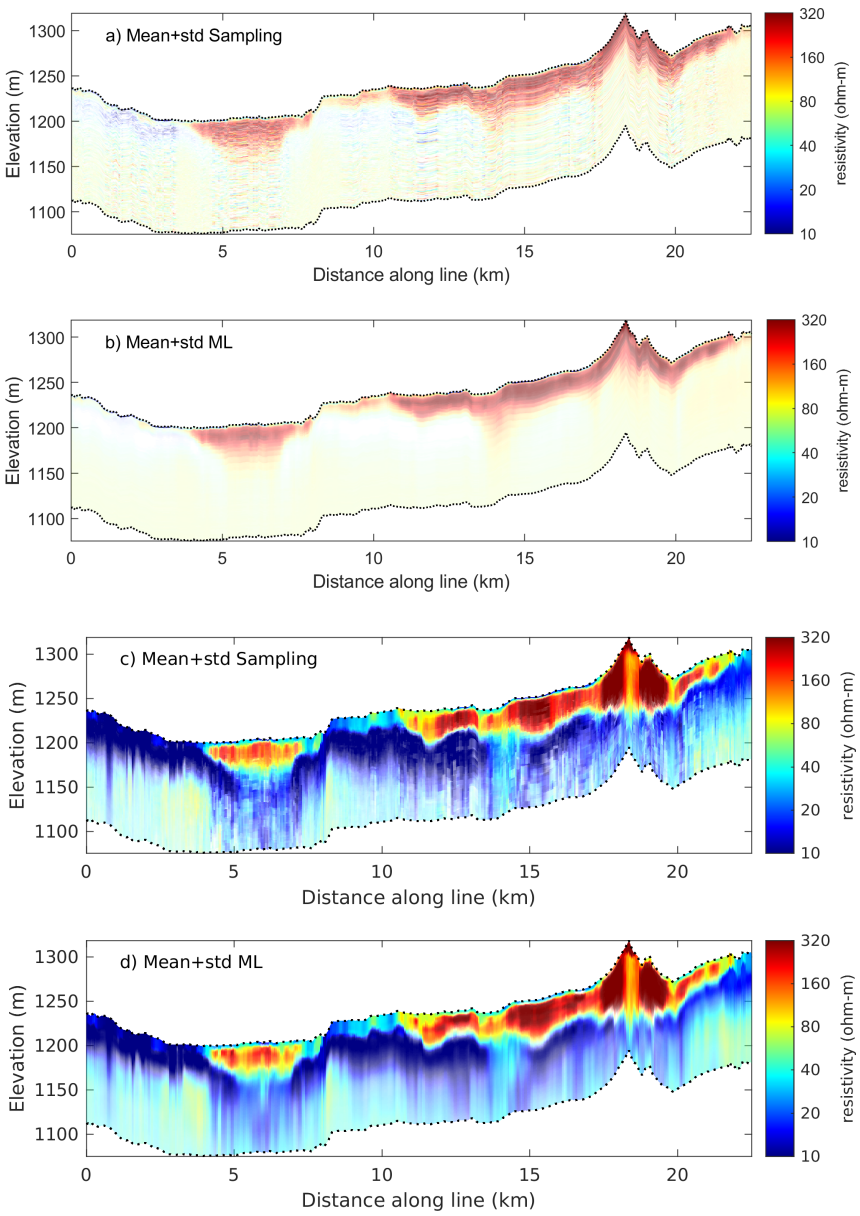

Figure 3. a-b) Pointwise mean obtained from $\sigma_A(\mathbf{m})$ obtained using the extended rejection sampler (a) and machine learning (b) based on a training data set of size 1000000. c-d) As a-b) but for $\sigma_B(\mathbf{m})$. Transparency based on pointwise posterior standard deviation. The mean and standard deviation without transparency are available in the supplementary material.

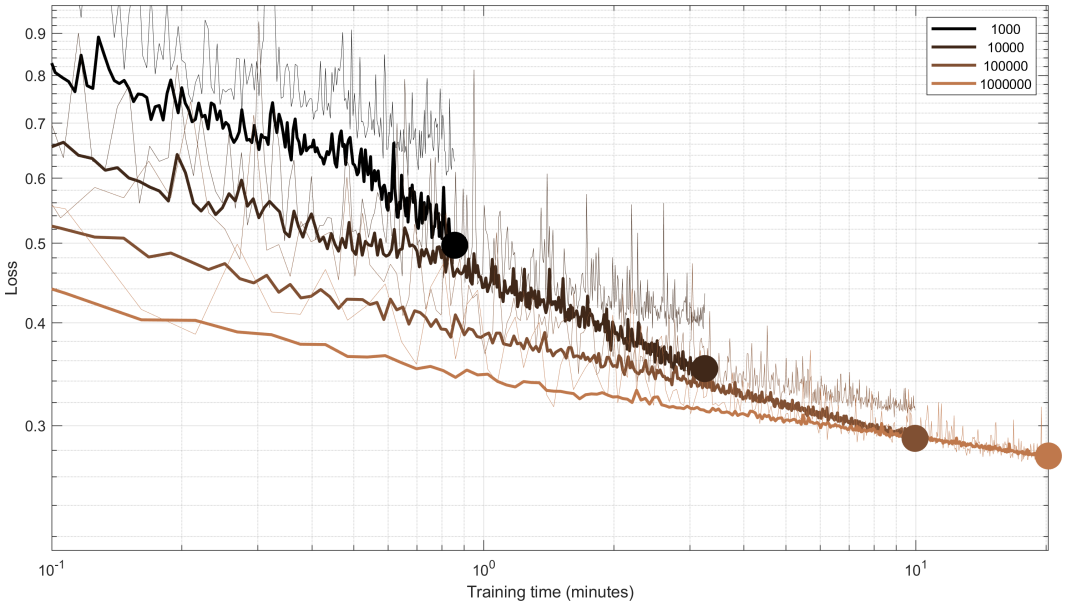

Figure 4. Training (thick lines) and validation (thin lines) loss as a function of training time for $N_t = [1000, 10000, 100000, 1000000]$.

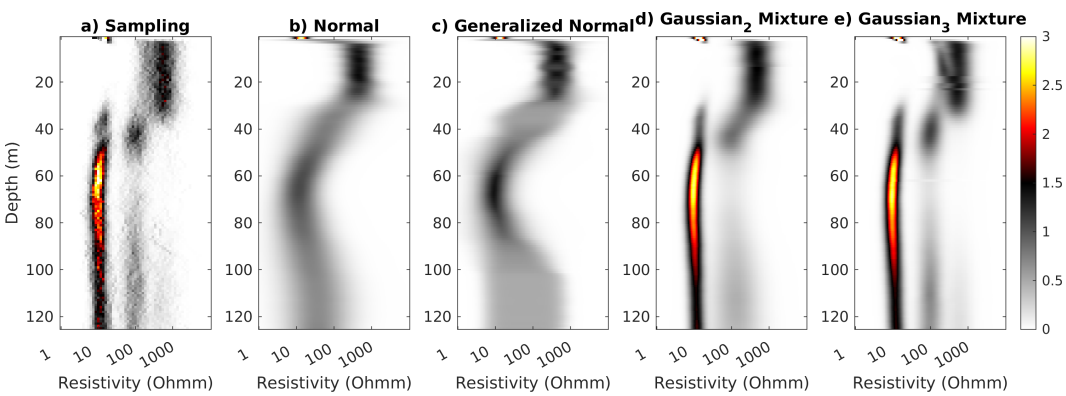

Figure 5. 1D posterior probability density with depth using data at X=6.2 km a) obtained using sampling, and constructed from statistical properties inferred for a b) normal distribution, c) generalized normal distribution, d)-e) a mixture model based on 2 and 3 1D normal distributions.

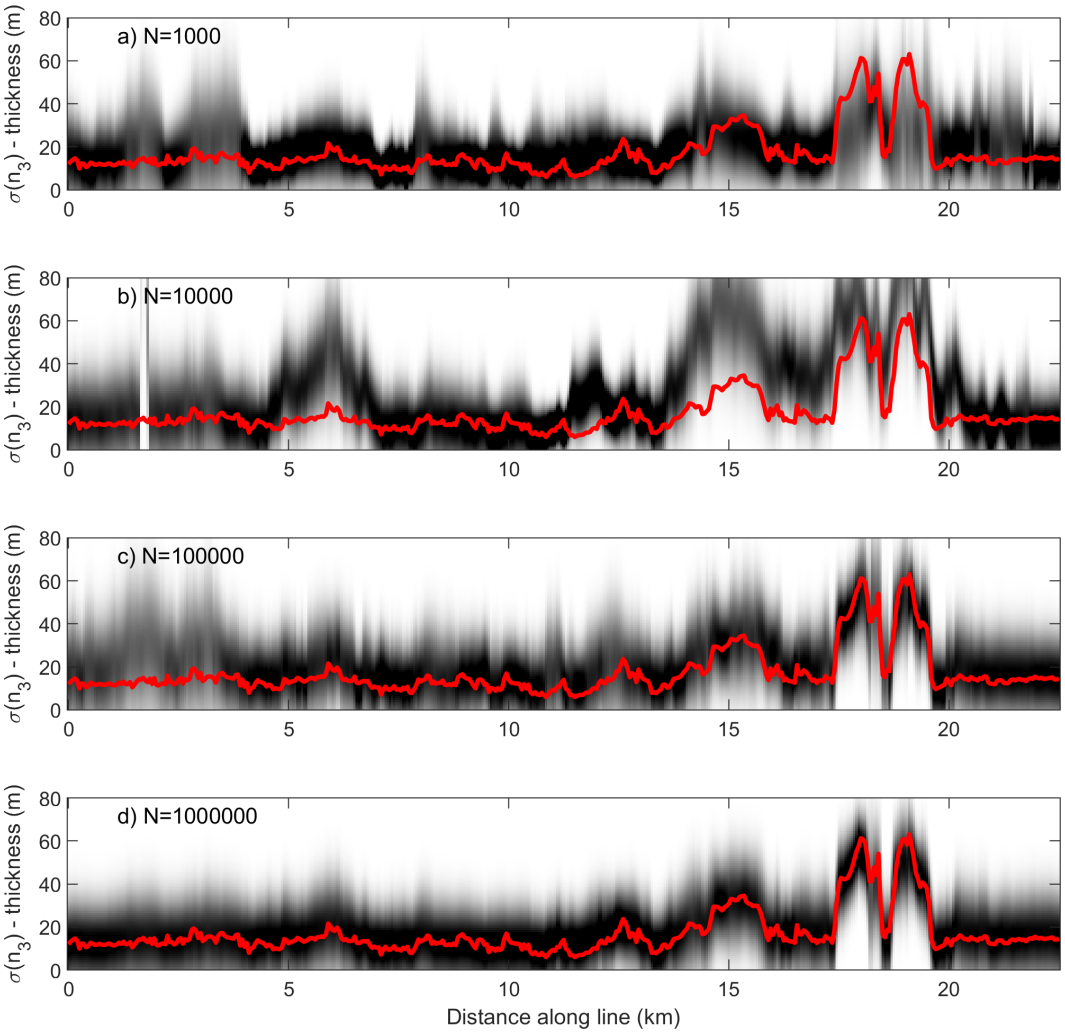

Figure 6. Mean of the posterior distributions $\sigma_C(\mathbf{n_3})$ estimated using sampling (red line) compared with the estimated mean and standard deviation of $\sigma_C(\mathbf{n_3})$ (probability density as grayscale) estimated using the machine learning approach using a training data set of size $N_T = [1000, 10000, 100000, 1000000]$ in a)-d).

3.8 Estimation of properties of $\boldsymbol{\sigma}(\mathbf{n_1})$

$\sigma(\mathbf{n_1})$ refers to the existence (or lack of) a layer interface, which can be formulated as a binary classification problem. Therefore, a classification type network is constructed using a sigmoid activation function, and the loss function in Eqn. 17. 4 hidden layers are used.

Figures 7a and 7c refer to the pointwise posterior probability of locating a layer interface, as computed from a sample from the posterior distribution of $\sigma_B(\mathbf{n_1})$ and $\sigma_C(\mathbf{n_1})$. The corresponding results obtained as the output of a trained neural network based on a training data set of size $N_T = 1000000$ are shown in Figures 7b and 7d. The prior probability of a layer interface is around 0.1, and hence a posterior probability of 0.25 is indicative of a layer interface.

The results using sampling and the machine learning approach are in both cases very similar with a bit more variability in the results obtained using sampling, due to the use of a finite-sized sample of the posterior distribution.

3.9 Estimation of properties of $\boldsymbol{\sigma}(\mathbf{n_3})$

Finally, we consider the discrete parameter $\mathbf{n_3}$ which refer to lithology type, which can be of type '1', '2' and '3'. The outcome at each model parameter is then a multi-class (three classes) classification problem. Therefore, a classification type network is constructed using a softmax activation function, and the loss function in Eqn. 17. 4 hidden layers is used.

Figures 8a,c,e show the posterior probability for each of the three classes obtained using sampling, while Figures 8b,d,f show the corresponding results obtained by evaluating the trained network. Except for some small-scale variations in the sampling results, due to using finite sample size, the obtained posterior statistics are strikingly similar.

4 Discussion

A typical application of probabilistic inversion is to use some sampling method to generate a large sample from the posterior distribution. Then some appropriate statistic, computed from the sample of the posterior distribution, is chosen and visualized.

The theory presented above proposes how one can construct a neural network that can directly estimate any statistical property of the posterior distribution (for discrete and continuous parameters) for which a probability distribution can be evaluated, without ever generating realizations of the posterior distribution. This can be achieved by

1. Construct a training data set, in the style of Devilee et al. (1999), $\mathbf{T}^* = [\mathbf{N}^*, \mathbf{D}^*_{sim}]$, where $\mathbf{N}^*$ represents a set of features/properties of interest, and $\mathbf{D}^*_{sim}$ represents a corresponding set of simulated data with noise, using both the forward and the noise model.
2. Design a neural network whose output layer represents the relevant statistical parameters Θ of the posterior distribution $\sigma(\mathbf{n})$ of interest.
3. Train the neural network by minimizing a loss function that is the negative log-likelihood of the probability density, $f(\Theta)$, whose properties one wishes to estimate.

Practical application of the methodology requires a) a neural network structure complex enough to be able to estimate the mapping $\mathbf{d}_{sim} \mapsto \Theta$, and b) a training data set large enough to allow inferring the mapping.

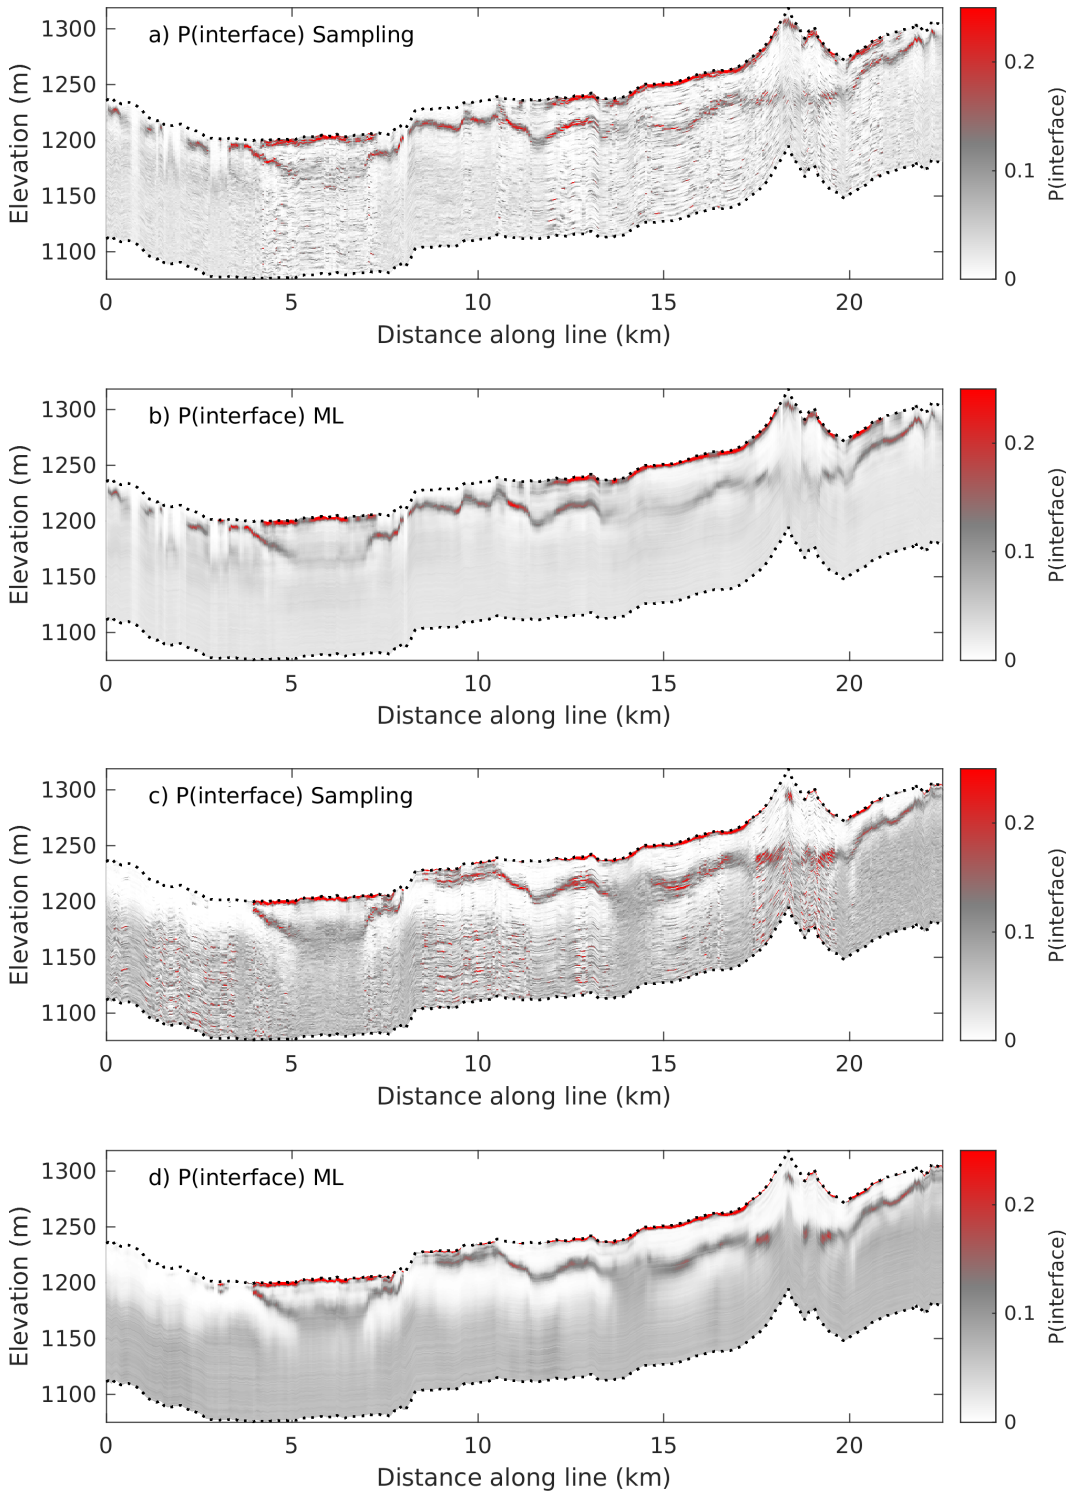

Figure 7. a-b) Posterior probability of a layer interface obtained using extended rejection sampling (a), and machine learning (b), for $\sigma_B(\mathbf{n}_1)$. c-d) Posterior probability of a layer interface obtained using extended rejection sampling (c), and machine learning (d), for $\sigma_C(\mathbf{n}_1)$.

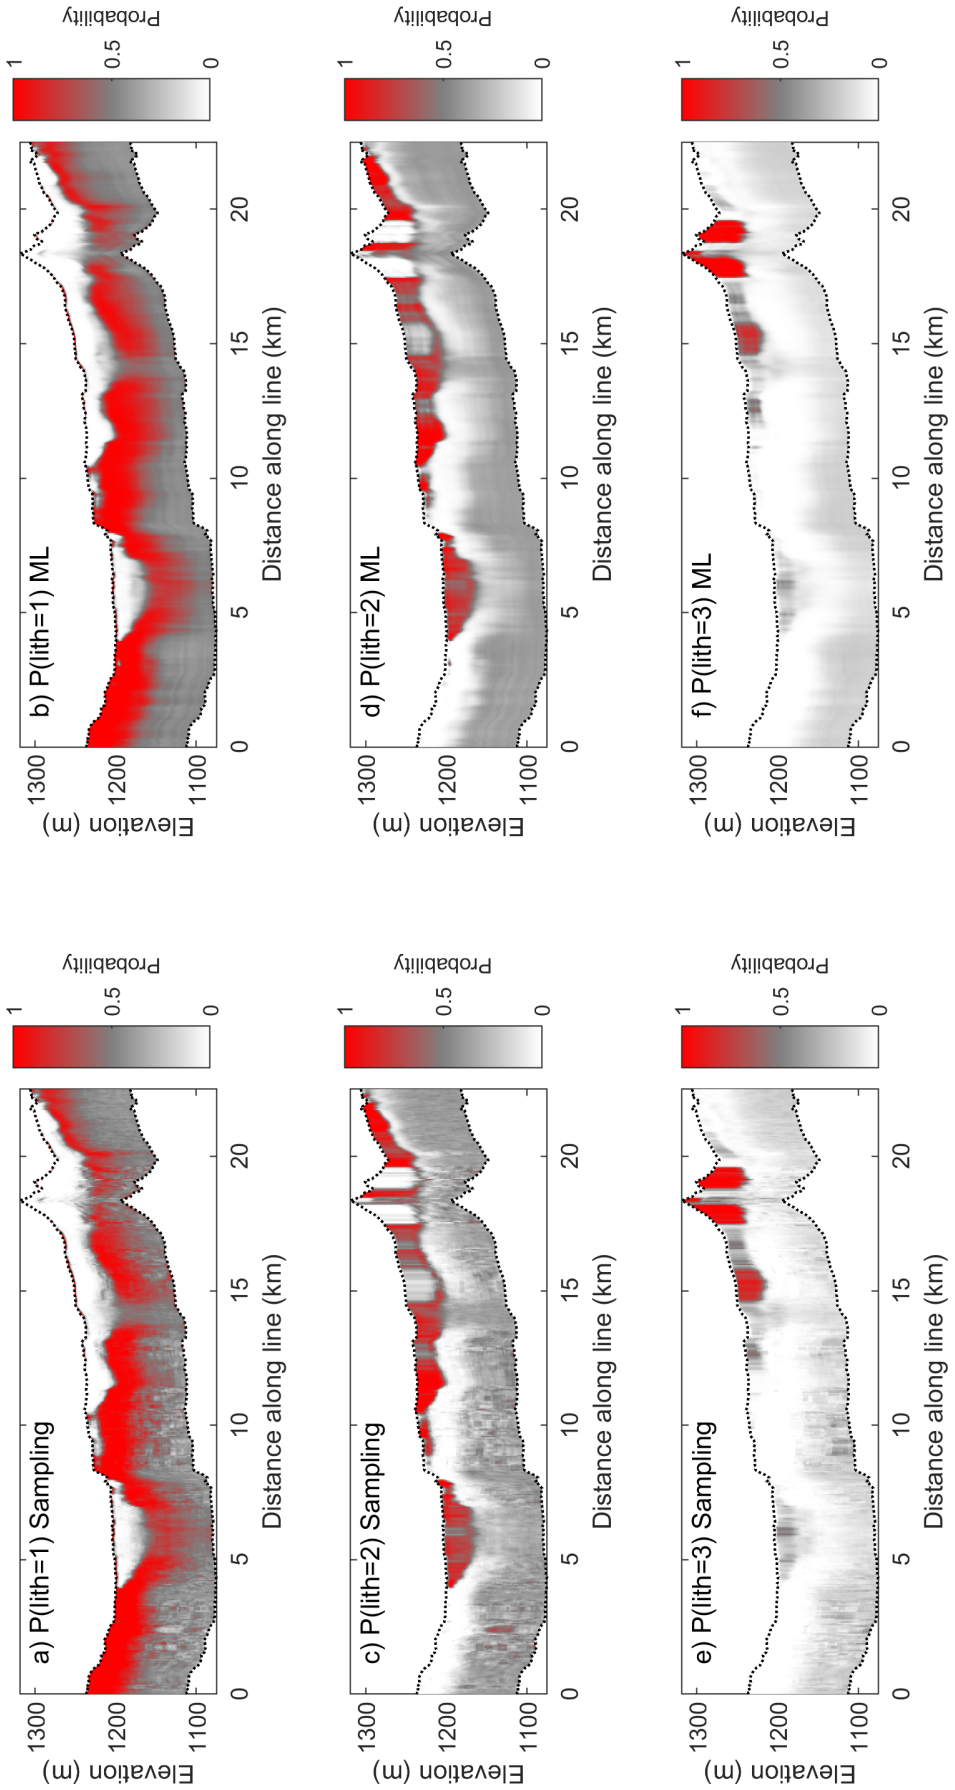

Figure 8. Posterior probability of lithology, using a-c) sampling and d-f) machine learning for $\sigma_C(\mathbf{n}_3)$.

The methodology was applied and demonstrated in a case study using airborne EM data from Morrill, Nebraska. Several (uninformed to more informed) prior models were considered, describing both subsurface resistivity (a continuous parameter, $\mathbf{m}$) and lithology (a discrete parameter, $\mathbf{n_2}$) and the considered forward problem was nonlinear. In addition, the method was used to estimate posterior statistics of low-dimensional features of the prior models, such as the existence of a layer interface, $\mathbf{n_1}$, and the thickness of layers with resistivity above 225 ohmm, $\mathbf{n_3}$. Results showed that using a training data set of size $N_T \geq 100000$ leads to a trained neural network that provides estimates of posterior statistics similar to those obtained using sampling methods, using a fraction of the computational power (about 5ms per sounding).

Relation to previous work

The methodology proposed here is based on the ideas originally proposed by Devilee et al. (1999) and extended by e.g. Meier et al. (2007); Earp et al. (2020). The explicit goal of Meier et al. (2007), and following related work (Shahraeeni & Curtis, 2011; de Wit et al., 2013; Earp et al., 2020), is to model the marginal posterior distribution as a mixture of Gaussian distributions.

The key goal of this manuscript is to show that one can construct a neural network that can estimate any desired statistical parameter describing the posterior distribution $\sigma(\mathbf{m}, \mathbf{n})$ related to both the main set of model parameters $\mathbf{m}$ and any set of parameters $\mathbf{n}$ related to the main parameters, for which an appropriate loss function can be defined. Such statistical parameters can, as a special case, represent the 1D marginal posterior distribution (as shown in Figure 5e) as in Meier et al. (2007).

Limitations

The proposed method does not generate realizations of the posterior distribution, as do other sampling-based methods (B. J. Minsley, 2011; Brodie & Sambridge, 2012; Hansen & Minsley, 2019; Hansen, 2021). Instead, statistics of the posterior distribution of features of interest are estimated directly by applying a trained neural network.

In some use cases, one may need the realizations, for example to propagate flow responses from of a set of realizations from the posterior representing hydraulic parameters, (Vilhelmsen et al., 2019). But, in many applications, where one is primarily interested in some statistical parameter describing the posterior, such as the posterior probability of a lithology type, the presented methodology may be very useful.

The methodology is particularly promising for localized inverse problems, where the trained neural network can be set up and trained once, but applied many times. It is less obviously suited to 3D inversions with very large model dimensions because 1) construction an adequately large training data set will be difficult and CPU intensive, 2) solving the 3D forward problem may be CPU intensive, and 3) it may be very difficult to train a neural network with millions of parameters in the output layer. The use of variational inference has been suggested as a more efficient approach to estimate marginal statistics for higher dimensional inverse problems, see e.g. Zhang and Curtis (2020a).

Potential

The immediate appeal of the proposed methodology is that it leads to fast prediction times. One can get similar results, but much faster, compared with using sampling-based methods to analyze the posterior distribution.

The presented method is faster than linearized least squares based deterministic inversion of EM data (Auken et al., 2017), which have been widely used for inversion of large surveys (B. Minsley et al., 2021) because they require much less computational re-

sources than sampling-based methods. With the computational efficiency of the proposed method, the computational benefits of linearized methods are no longer so substantial that one should ignore the benefits of using the probabilistic methods that allow the use of site-specific prior information, a non-linear forward model, and full exploration of the space of uncertainty.

The more general appeal is that the proposed methodology allows the use of in principle arbitrarily complex prior models. The only requirement is that one must be able to generate independent realizations of the prior model. This allows an end-user to actively choose a prior model based on available information, as opposed to being forced to use the implicit prior assumptions in most available inversion algorithms, such as the assumptions of a layered subsurface (B. J. Minsley, Foks, & Bedrosian, 2021) or a Gaussian type smooth prior (Auken & Christiansen, 2004). The prior can be constructed according to site-specific information, and can then estimate posterior statistics of any parameter that can be computed from the prior model, as illustrated with the parameters $\mathbf{n}_1$ and $\mathbf{n}_2$ in the case study

The main challenge then becomes the construction of realistic prior models that represent geological realistic information.

Workflow for decision makers.

A key feature/limitation of the proposed methodology is that a realization of the posterior is never realized. Instead, statistics of features of the posterior distribution are computed directly.

We argue that for decision-makers it is the statistics of features related to model parameters that are most often of interest, not the individual realizations themselves, as discussed by Scales and Snieder (1997). Even if a large sample exists of the posterior distribution for each data location along the considered profile of data at Morrill, then one will have to compute some feature/statistics of this sample to make it useable by decision-makers.

The examples presented here demonstrate that similar results are obtained using either sampling of the posterior distribution, followed by analyses of the sample of the posterior distribution, or directly using the proposed machine learning methodology. The key practical difference to using sampling methods is then that with our proposed methodology one has to quantify the feature that one is interested in and specify an appropriate loss function before running the inversion. Whereas using sampling methods to sample the posterior, one can convert the realizations of the posterior into a specific feature, and perform the posterior analysis, after the sampling algorithm has run.

5 Conclusions

A simple, yet powerful, approach to probabilistic inversion has been proposed. Its application requires that one can simulate sets of examples capturing the known information. That is 1) sample from an arbitrarily complex prior model, 2) solving the forward problem, and 3) adding realistic noise to the simulated data. From each of these sets of models and data, a set of corresponding features related to the model parameters can be obtained. Together these represent, up to the limit of the finite set of models, all known information about these features of interest.

From such sets of features and corresponding noisy input data, posterior statistics describing the features given the data can be obtained by minimizing an appropriate loss function. This provides the ability to carry out a fast and accurate estimation of relevant posterior statistics given an observed dataset.

A case study of the methodology applied to a nonlinear probabilistic inversion of EM data demonstrates it is possible to obtain posterior statistics similar to those obtained using sampling methods, using a fraction of the computation time. This approach allows the use and testing of multiple prior models, and to consider multiple features related to the prior distributions, in a fully probabilistic setting using only modest computational resources. The method has most appeal for localized inverse problems, where the same trained neural network can be applied on many datasets with little computational effort.

Data Availability Statement

The airborne EM data used in this study is freely available (Smith et al., 2010) and can be accessed through `https://doi.org/10.3133/ofr20101259`. Training data sets and python notebooks for training and prediction will be made available upon publication at `https://github.com/cultpenguin/probabilistic-inverse-problems_and_ml`

Acknowledgments
This work is funded by the Danish Free Research Council, project 717-00160B.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng, X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems.* Retrieved from `https://www.tensorflow.org/` (Software available from tensorflow.org)

Abraham, J. D., Cannia, J. C., Bedrosian, P. A., Johnson, M. R., Ball, L. B., & Sibray, S. S. (2012). *Airborne electromagnetic mapping of the base of aquifer in areas of western nebraska.* U.S. Geological Survey Scientific Investigations Report 2011-5219, 38 p.

Ardizzone, L., Kruse, J., Wirkert, S., Rahner, D., Pellegrini, E. W., Klessen, R. S., ... Köthe, U. (2018). Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*.

Auken, E., Boesen, T., & Christiansen, A. V. (2017). A review of airborne electromagnetic methods with focus on geotechnical and hydrological applications from 2007 to 2017. *Advances in Geophysics*, *58*, 47–93.

Auken, E., & Christiansen, A. V. (2004). Layered and laterally constrained 2d inversion of resistivity data. *Geophysics*, *69*(3), 752–761.

Auken, E., Christiansen, A. V., Kirkegaard, C., Fiandaca, G., Schamper, C., Behroozmand, A. A., ... others (2014). An overview of a highly versatile forward and stable inverse algorithm for airborne, ground-based and borehole electromagnetic and electric data. *Exploration Geophysics*, *46*(3), 223–235.

Bai, P., Vignoli, G., Andrea, V., Jouni, N., & Vacca, G. (2020). (quasi-) real-time inversion of airborne time-domain electromagnetic data via artificial neural network. *Remote Sensing*.

Bishop, C. M., et al. (1995). *Neural networks for pattern recognition.* Oxford university press.

Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American statistical Association*, *112*(518), 859–877.

Bording, T. S., Asif, M. R., Barfod, A. S., Larsen, J. J., Zhang, B., Grombacher, D. J., ... others (2021). Machine learning based fast forward modelling of ground-based time-domain electromagnetic data. *Journal of Applied Geophysics*, 104290.

Bosch, M., Mukerji, T., & Gonzalez, E. F. (2010). Seismic inversion for reservoir

properties combining statistical rock physics and geostatistics: A review. *Geophysics*, *75*(5), 75A165–75A176.

Brodie, R. C., & Sambridge, M. (2012). Transdimensional Monte Carlo inversion of AEM data. *ASEG Extended Abstracts*, *2012*(1), 1–4.

Chollet, F. (2015). *Keras.* `https://github.com/fchollet/keras`. GitHub.

Christensen, N. B. (2002). A generic 1-D imaging method for transient electromagnetic data. *Geophysics*, *67*(2), 438–447.

Constable, S. C., Parker, R. L., & Constable, C. G. (1987). Occam's inversion: A practical algorithm for generating smooth models from electromagnetic sounding data. *Geophysics*, *52*(3), 289–300.

Conway, D., Alexander, B., King, M., Heinson, G., & Kee, Y. (2019). Inverting magnetotelluric responses in a three-dimensional earth using fast forward approximations based on artificial neural networks. *Computers & geosciences*, *127*, 44–52.

Cox, L. H., Wilson, G. A., & Zhdanov, M. S. (2010). 3D inversion of airborne electromagnetic data using a moving footprint. *Exploration Geophysics*, *41*(4), 250–259.

Devilee, R., Curtis, A., & Roy-Chowdhury, K. (1999). An efficient, probabilistic neural network approach to solving inverse problems: inverting surface wave velocities for eurasian crustal thickness. *Journal of Geophysical Research: Solid Earth*, *104*(B12), 28841–28857.

de Wit, R. W., Valentine, A. P., & Trampert, J. (2013). Bayesian inference of earth's radial seismic structure from body-wave traveltimes using neural networks. *Geophysical Journal International*, *195*(1), 408–422.

Dillon, J. V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., ... Saurous, R. A. (2017). Tensorflow distributions. *arXiv preprint arXiv:1711.10604*.

Earp, S., & Curtis, A. (2020). Probabilistic neural network-based 2d travel-time tomography. *Neural Computing and Applications*, *32*(22), 17077–17095.

Earp, S., Curtis, A., Zhang, X., & Hansteen, F. (2020, 08). Probabilistic neural network tomography across Grane field (North Sea) from surface wave dispersion data. *Geophysical Journal International*, *223*(3), 1741-1757. Retrieved from `https://doi.org/10.1093/gji/ggaa328` doi: 10.1093/gji/ggaa328

Foks, N., & Minsley, B. (2020). *Geophysical Bayesian inference in Python (GeoBiPy).* https://github.com/usgs/geobipy. GitHub. doi: 10.5066/P9K3YH9O

Geman, S., & Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*(6), 721–741.

Grana, D., & Della Rossa, E. (2010). Probabilistic petrophysical-properties estimation integrating statistical rock physics with seismic inversion. *Geophysics*, *75*(3), O21–O37.

Grayver, A. V., Streich, R., & Ritter, O. (2013). Three-dimensional parallel distributed inversion of CSEM data using a direct forward solver. *Geophysical Journal International*, *193*(3), 1432–1446.

Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, *82*(4), 711–732.

Hansen, T. M. (2021). Efficient probabilistic inversion using the rejection sampler—exemplified on airborne EM data. *Geophysical Journal International*, *224*(1), 543–557.

Hansen, T. M., Cordua, K. C., & Mosegaard, K. (2012). Inverse problems with non-trivial priors - efficient solution through sequential Gibbs sampling. *Computational Geosciences*, *16*(3), 593-611. doi: 10.1007/s10596-011-9271-1

Hansen, T. M., & Cordua, K. S. (2017). Efficient monte carlo sampling of inverse problems using a neural network-based forward—applied to gpr crosshole traveltime inversion. *Geophysical Journal International*, *211*(3), 1524–1533.

Hansen, T. M., Cordua, K. S., Jacobsen, B. H., & Mosegaard, K. (2014). Accounting for imperfect forward modeling in geophysical inverse problems—exemplified for crosshole tomography. *Geophysics*, *79*(3), H1–H21.

Hansen, T. M., Cordua, K. S., Looms, M. C., & Mosegaard, K. (2013, 03). SIPPI: a Matlab toolbox for sampling the solution to inverse problems with complex prior information: Part 1, methodology. *Computers & Geosciences*, *52*, 470-480. doi: 10.1016/j.cageo.2012.09.004

Hansen, T. M., Cordua, K. S., Zunino, A., & Mosegaard, K. (2016). Probabilitsic Integration of Geo-Information. In M. Moorekamp (Ed.), *Joint inversion* (Vol. 218). John Wiley & Sons.

Hansen, T. M., & Minsley, B. J. (2019). Inversion of airborne EM data with an explicit choice of prior model. *Geophysical Journal International*, *218*(2), 1348–1366.

Hansen, T. M., Mosegaard, K., & Cordua, K. C. (2008). Using geostatistics to describe complex a priori information for inverse problems. In J. M. Ortiz & X. Emery (Eds.), *Viii international geostatistics congress* (Vol. 1, p. 329-338). Mining Engineering Department, University of Chile.

Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, *57*(1), 97.

Hornik, K., Stinchcombe, M., & White, H. (1990). Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural networks*, *3*(5), 551–560.

Howard, D. (2020). Geological survey of western australia: Ausaem20-wa project. *Preview*, *2020*(205), 18–18.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Köpke, C., Irving, J., & Elsheikh, A. H. (2018). Accounting for model error in Bayesian solutions to hydrogeophysical inverse problems using a local basis approach. *Advances in Water Resources*, *116*, 195–207.

Laloy, E., Hérault, R., Jacques, D., & Linde, N. (2018). Training-image based geostatistical inversion using a spatial generative adversarial neural network. *Water Resources Research*, *54*(1), 381–406.

Laloy, E., & Vrugt, J. A. (2012). High-dimensional posterior exploration of hydrologic models using multiple-try dream (zs) and high-performance computing. *Water Resources Research*, *48*(1).

Madsen, R. B., & Hansen, T. M. (2018). Estimation and accounting for the modeling error in probabilistic linearized amplitude variation with offset inversion. *Geophysics*, *83*(2), N15–N30.

Malinverno, A. (2002). Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem. *Geophysical Journal International*, *151*(3), 675–688.

Meier, U., Curtis, A., & Trampert, J. (2007). Global crustal thickness from neural network inversion of surface wave data. *Geophysical Journal International*, *169*(2), 706–722.

Menke, W. (2012). *Geophysical data analysis: Discrete inverse theory* (Vol. 45). Academic Press.

Metropolis, N., Rosenbluth, M., Rosenbluth, A., Teller, A., & Teller, E. (1953). Equation of state calculations by fast computing machines. *J. Chem. Phys.*, *21*, 1087–1092.

Minsley, B., Rigby, J., James, S., Burton, B., Knierim, K., Pace, M., ... Kress, W. (2021). *Closing a scale-gap in earth observation using regional-scale airborne geophysics in the lower mississippi valley* (Tech. Rep.). Copernicus Meetings.

Minsley, B. J. (2011). A trans-dimensional Bayesian Markov chain Monte Carlo algorithm for model assessment using frequency-domain electromagnetic data. *Geophysical Journal International*, *187*(1), 252–272.

Minsley, B. J., Foks, N. L., & Bedrosian, P. A. (2021). Quantifying model structural uncertainty using airborne electromagnetic data. *Geophysical Journal International*, *224*(1), 590–607.

Minsley, B. J., Rigby, J. R., James, S. R., Burton, B. L., Knierim, K. J., Pace, M. D., ... Kress, W. H. (2021). Airborne geophysical surveys of the lower mississippi valley demonstrate system-scale mapping of subsurface architecture. *Communications Earth & Environment*, *2*(1), 1–14.

Moghadas, D. (2020). One-dimensional deep learning inversion of electromagnetic induction data using convolutional neural network. *Geophysical Journal International*, *222*(1), 247–259.

Moghadas, D., Behroozmand, A. A., & Christiansen, A. V. (2020). Soil electrical conductivity imaging using a neural network-based forward solver: Applied to large-scale bayesian electromagnetic inversion. *Journal of Applied Geophysics*, 104012.

Mosegaard, K., & Tarantola, A. (1995). Monte Carlo sampling of solutions to inverse problems. *J. geophys. Res*, *100*(B7), 12431–12447.

Mosser, L., Dubrule, O., & Blunt, M. J. (2017). Reconstruction of three-dimensional porous media using generative adversarial neural networks. *Physical Review E*, *96*(4), 043309.

Mosser, L., Dubrule, O., & Blunt, M. J. (2020). Stochastic seismic waveform inversion using generative adversarial networks as a geological prior. *Mathematical Geosciences*, *52*(1), 53–79.

Puzyrev, V., & Swidinsky, A. (2019). Inversion of 1d frequency-and time-domain electromagnetic data with convolutional neural networks. *arXiv preprint arXiv:1912.00612*.

Rimstad, K., Avseth, P., & Omre, H. (2012). Hierarchical Bayesian lithology/fluid prediction: A North Sea case study. *Geophysics*, *77*(2), B69–B85.

Röth, G., & Tarantola, A. (1994). Neural networks and inversion of seismic data. *Journal of Geophysical Research: Solid Earth*, *99*(B4), 6753–6768.

Sambridge, M., Bodin, T., Gallagher, K., & Tkalčić, H. (2013). Transdimensional inference in the geosciences. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *371*(1984), 20110547.

Scales, J. A., & Snieder, R. (1997). To [b]ayes or not to [b]ayes? *Geophysics*, *62*(4), 1045–1046.

Scheidt, C., Renard, P., & Caers, J. (2015). Prediction-focused subsurface modeling: investigating the need for accuracy in flow-based inverse modeling. *Mathematical Geosciences*, *47*(2), 173–191.

Shahraeeni, M. S., & Curtis, A. (2011). Fast probabilistic nonlinear petrophysical inversion. *Geophysics*, *76*(2), E45–E58.

Smith, B. D., Abraham, J. D., Cannia, J. C., Minsley, B. J., Deszcz-Pan, M., & Ball, L. B. (2010). *Helicopter electromagnetic and magnetic geophysical survey data, portions of the North Platte and South Platte Natural Resources Districts, Western Nebraska* (Tech. Rep. No. 2010-1259). U.S. Geological Survey Scientific Investigations Report. doi: 10.3133/ofr20101259

Tarantola, A. (2005). *Inverse problem theory and methods for model parameter estimation*. SIAM.

Tarantola, A., & Valette, B. (1982a). Generalized nonlinear inverse problems solved using the least squares criterion. *Rev. Geophys. Space Phys*, *20*(2), 219–232.

Tarantola, A., & Valette, B. (1982b). Inverse problems= quest for information. *J. Geophys*, *50*(3), 150–170.

Tikhonov, A. N. (1963). On the solution of ill-posed problems and the method of regularization. In *Doklady akademii nauk* (Vol. 151, pp. 501–504).

Viezzoli, A., Christiansen, A. V., Auken, E., & Sørensen, K. (2008). Quasi-3D modeling of airborne TEM data by spatially constrained inversion. *Geophysics*, *73*(3), F105–F113.

Vilhelmsen, T. N., Auken, E., Christiansen, A. V., Barfod, A. S., Marker, P. A., & Bauer-Gottwein, P. (2019). Combining clustering methods with mps to estimate structural uncertainty for hydrological models. *Frontiers in Earth Science*, *7*, 181.

Zhang, X., & Curtis, A. (2020a). Seismic tomography using variational inference methods. *Journal of Geophysical Research: Solid Earth*, *125*(4), e2019JB018589.

Zhang, X., & Curtis, A. (2020b). Variational full-waveform inversion. *Geophysical Journal International*, *222*(1), 406–411.

Zhang, X., & Curtis, A. (2021). Bayesian geophysical inversion using invertible neural networks. *Journal of Geophysical Research: Solid Earth*, *126*(7), e2021JB022320.

Zhao, X., Curtis, A., & Zhang, X. (2022). Bayesian seismic tomography using normalizing flows. *Geophysical Journal International*, *228*(1), 213–239.