Probabilistic inverse problems using machine learning - applied to inversion of airborne EM data.

Thomas Mejer Hansen

[1]Department of Geoscience, Aarhus University, Denmark

Key Points:

- A machine learning approach for solving probailistic inverse problems by directly estimating properties of the posterior distribution
- Allow the use of arbitrarily complex prior and noise models as long as they can be sampled
- Exemplified on inversion of airborne electromagnetic data, allowing analysis of more than 100000 1D soundings per second

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

Abstract

Probabilistic inversion methods allow, in principle, to combine probabilistic geo-information from diverse sources into one consistent statistical model (the posterior distribution) containing all available information. In practice, however, they rely on Monte Carlo sampling methods, which can be extremely computationally demanding. Here a general, and simple to apply, method is presented, utilizing machine learning, which allows fast direct estimation of properties of the posterior distribution. The fundamental idea is to construct a training data set that represents all the information represented in the probabilistic formulation of inverse problems. From such a training data set, it is demonstrated how regression and classification type neural networks can be designed, with specific choices of output layer and loss functions, that allows direct characterization of the posterior distribution using regression and classification type The methodology is demonstrated on probabilistic inversion of airborne electromagnetic data and compared to results obtained computationally more expensive sampling methods.

Plain Language Summary

Probabilistic inversion is in principle ideal as a method for combining available information about geo-models, in a way that will allow detailed risk analysis and hypothesis testing. In practice, however, such methods have historically been limited for practical use on realistic size geo models, because they 1) typically require computationally expensive Monte Carlo sampling algorithms, 2) typically rely on relatively simple prior/structural models. Here a general method for probabilistic inversion, based on machine learning, is presented, that allows direct estimation of properties of the posterior distribution, that represents the combined available information, without the need to generate actual realizations from the posterior distribution. It is simply to apply and requires only that one can 1) quantify and sample from a chosen prior distribution, 2) compute expected noise-free data for any given model, and 3) simulate noise from a chosen noise model. As an example, results of inversion of airborne electromagnetic data obtained using the machine learning-based method and a sampling-based method are compared. The machine learning-based method is accurate, fast, allowing analysis of about 100000 1D soundings per second, making it applicable to analysis of very large electromagnetic data sets.

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 certain features related to the geo-model, consistent with all information. For example, one may wish to validate different hypotheses against the available information, or, one may wish to determine the cumulative thickness of clay, that suggests protection of drinking water reservoirs.

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; P. C. Hansen, 1992; Constable et al., 1987). In practice, in part due to noise on data and model 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 methods can, in principle, take into account arbitrarily complex information, and integrate the information into one consistent model. The goal of probabilistic methods is to construct a posterior probability density, or realizations of such a probability density, that describe the full uncertainty, and for which all realizations are consistent with all available information. Any property of the posterior distribution can then be approximated by analysis of the sample of the posterior distribution.

The probabilistic approach is therefore ideal for decision-makers, as it will allow probabilistic analysis and risk assessment consistent with, in principle, all available information.

The main obstacle to applying the probabilistic methodology in practice is that it is computationally very demanding (Hastings, 1970; Mosegaard & Tarantola, 1995). Sampling-based methods typically require both sampling 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 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 (T. M. Hansen et al., 2014; Madsen & Hansen, 2018; Köpke et al., 2018). Fast to evaluate machine learning algorithms have been used to approximate forward modeling (T. M. Hansen & Cordua, 2017; Conway et al., 2019; Moghadas et al., 2020; Bording et al., 2021). Generative adversarial networks (GAN) can be used to represent prior information, which, once trained, allows for a very efficient way of sampling a prior (Mosser et al., 2017; Laloy et al., 2018; Mosser et al., 2020).

Others have used machine learning methods to directly solve the inverse problem by estimating a direct mapping from data to model parameters (Puzyrev & Swidinsky, 2019; Moghadas, 2020; Bai et al., 2020). These methods estimate a single model, as the deterministic methods, and typically without accounting for uncertainty on geophysical data, and uncertainty on the predicted model parameters. Meier et al. (2007) suggest modeling the posterior distribution using a mixture of Gaussian models. Ardizzone et al. (2019) propose to make use of invertible neural networks, to estimate simultaneously both the forward and inverse mapping between data and model parameters, which allows generating multiple realizations of the posterior distribution, from which properties of the posterior distribution can be estimated. This method is applied to geophysical data by Zhang and Curtis (2021).

Here a novel method is proposed that allows direct estimation of statistical properties of the posterior distribution, and potentially any feature linked to the prior distribution (and hence also the posterior distribution), without the need to generate realizations of the posterior distribution. The method allows using arbitrarily complex prior information and forward models and naturally accounts for arbitrarily complex noise. It provides accurate properties of the posterior distribution at a fraction of the time used by sampling-based approaches.

In the following the method is outlined for inverse problems in general, using arbitrarily complex prior, forward, and noise models. Then, we demonstrate the method applied to the analysis of airborne electromagnetic data.

2 Method

Let $\mathbf{m} = [m_1, m_2, \ldots, m_M]$ represent 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 coordinate system. $\mathbf{m}$ can represent, for example, geophysical properties such as resistivity, velocity, or any geological/geophysical/geochemical parameter.

A key challenge in geoscience is that of inferring information about $\mathbf{m}$ from different types of available information, such as geological expert knowledge, geophysical data, well log data, etc. This challenge is generally referred to as an inverse problem. (Tarantola & Valette, 1982b; Tarantola, 2005) describe the inverse problem as problem of probabilistic integration of information. Available information about $\mathbf{m}$ is described in the form of a probability density and then combined using conjunction of information to obtain one probability density that describes all available information. Say a specific type of information about structural information is quantified by $\rho(\mathbf{m})$, information from seismic data and well logs by $L(\mathbf{m})$. 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

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

I.e. the conjunction of the independent information is proportional to the product of probability densities describing each independent set of information. The likelihood $L(\mathbf{m})$ quantifies a probability distribution that quantifies the difference between observed data $\mathbf{d}_{obs}$, and the noise-free data $\mathbf{d}$ computed by evaluating the forward model

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

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

The central problem in probabilistic inversion, the inverse problem, is inferring information about $\sigma(\mathbf{m})$, which in principle contain the combined information of, in this case, both structural prior information, through the prior $\rho(\mathbf{m})$, and information from geophysical data, through $L(\mathbf{m})$.

The most widely used method for solving probabilistic formulated inverse problems is by sampling the posterior distribution using variants of the Metropolis algorithm, Eqn. 1, (Metropolis et al., 1953; Hastings, 1970; Geman & Geman, 1984; Mosegaard & Tarantola, 1995; Laloy & Vrugt, 2012; T. M. Hansen et al., 2013). Such sampling methods can be extremely computational demanding to the point where they cannot be practically applied. These methods rely on solving the forward problem, Eqn 2, many (up to millions of) times. Solving the inverse problem by sampling $\sigma(\mathbf{m})$ is typically a much harder computational problem than solving the forward problem, Eqn 2.

T. M. Hansen (2021) demonstrate how one can sample the posterior distribution computationally efficiently using the extended rejection sampler, for a class of localized inverse problems, using informed prior models. This is achieved by constructing a lookup table consisting of models and corresponding data, where the models are generated as independent realizations of the prior $\rho(\mathbf{m})$, and the corresponding data as noise-free data obtained through Eqn. 2. The lookup table is constructed once and then used multiple times with the extended rejection sampler, to sample the posterior distribution for different observed data.

Here we consider the case when the parameter of interest may not be $\mathbf{m}$ itself, but instead a set of features/parameters $\mathbf{n}$ related to $\mathbf{m}$ through $\mathbf{n} = h(\mathbf{m})$. This is in practice often the case when $\mathbf{m}$ represents a geophysical parameter, such as resistivity or velocity, but where one is interested in the lithology or hydrological properties the geophysical parameter represents. A method is proposed that allows estimating posterior statistics of $\mathbf{m}$ and $\mathbf{n}$ by directly computing properties of $\sigma(\mathbf{m}, \mathbf{n})$, using a neural network trained on a data set representing a sample of all known information (including noise and modeling errors), in style with the lookup tables used by T. M. Hansen (2021).

The method is simple to apply and consists of two steps: A) construction of training set (A1) and construction and training of a neural network (A2). This is done once,

142 then in a second step B the trained machine learning algorithm is applied, very efficiently,
143 to potentially many sets of observed data (as demonstrated in the following example).

2.1 A1: Constructing training data

Eqn. 2 described the forward problem of computing noise free data. The forward problem of describing simulation of data with noise $\mathbf{d}_{obs}$ can be given by

$$\mathbf{d}_{sim} = g(\mathbf{m}) + n(\mathbf{m}), \tag{3}$$

145 where g represent a (possible) non-linear mapping typically describing some physical pro-
146 cess (e.g. solving the Maxwells equations), and n refer to a (possibly) non-linear noise
147 model.

Let $\mathbf{M}^* = [\mathbf{m}^{1*}, \mathbf{m}^{2*}, ...]$ represent N_r realizations of the prior, and let $\mathbf{D}^*_{sim} = [\mathbf{d}^{1*}_{sim}, \mathbf{d}^{2*}_{sim}, ...]$ represent N_r corresponding realizations of simulated noisy data, follow-ing Eqn. 3. Also, let $\mathbf{N}^* = [\mathbf{n}^{1*}, \mathbf{n}^{2*}, ...]$ represent N_r 'features' of each of the model realizations in $\mathbf{M}^*$. This constitutes a training data set

$$[\mathbf{N}^*; \mathbf{M}^*; \mathbf{D}^*; \mathbf{D}^*_{sim}], \tag{4}$$

148 that can be obtained simply by 1) sampling the prior, 2) solving the forward problem,
149 3) simulation of the noise, and 4) extracting/computing a feature of $\mathbf{n}$ from $\mathbf{m}$.

150 Usually, the forward mapping between $\mathbf{m}$ and noise-free data $\mathbf{d}$ is unique, and hence
151 if a large enough sample $[\mathbf{M}^*; \mathbf{D}^*]$ is available, one can in estimate the forward mapping
152 g arbitrarily precise using a neural network. The mapping between noise free $\mathbf{d}$ data and
153 $\mathbf{m}$ is though in general non-unique, as is the mapping from $\mathbf{d}_{sim}$ to $\mathbf{m}$ or $\mathbf{n}$

154 The sample in Eqn. 4 represents as much of the available information as can be
155 represented by a sample of size N_r. The larger the sample, the better representation of
156 the available information. If this sample is infinitely large, it will represent all available
157 information, and one could obtain a sample of the posterior of $\sigma(\mathbf{m}, \mathbf{n})$ simply by locat-
158 ing the entries in the training data set for which $\mathbf{d}_{obs} = \mathbf{d}^*_{sim}$. This is not viable, as the
159 probability of locating a match to the observed data in practice will be zero. Instead,
160 we propose to use machine learning to interpolate between the models in the lookup ta-
161 ble, to be able to estimate statistics of the posterior distribution.

162 Specifically, we consider a sample from $f(\mathbf{n}, \mathbf{d}_{sim})$ as $[\mathbf{N}^*; \mathbf{D}^*_{sim}]$, and how poste-
163 rior information on $\mathbf{n}$ can be obtained given some data with noise $\mathbf{d}_{obs}$.

2.2 A2: Train an algorithm to estimating properties of $\boldsymbol{\sigma}(\mathbf{m}, \mathbf{n})$

165 The idea is now to estimate statistical properties of $\sigma(\mathbf{n}|\mathbf{d}_{obs})$ by training a ma-
166 chine learning algorithm to estimate a mapping the observed data $\mathbf{d}_{obs}$ to the feature
167 $\mathbf{n}$ and/or the model parameters $\mathbf{m}, \mathbf{d}_{sim} \mapsto \mathbf{n}$. This mapping is for all practical pur-
168 poses non-unique, due to both the noise model and potentially the forward model. Here
169 we will use a neural network, but in principle, any machine learning method capable of
170 regression and classification, such as regression trees and support vector machines, can
171 be used(Bishop et al., 1995).

172 A neural network is constructed using a number of layers, that can be split into
173 an input layer, the central inner part of the neural network (which can consist of many
174 and different types of network layers), and an output layer.

175 The input layer refers here to the noisy data and consists of N_d nodes, organized
176 either as a 1D, 2D, or 3D array. The output layer has the number of nodes that refer
177 to the property of the posterior than one wishes to predict (see below). The inner part
178 of the network can be arbitrarily simple or complex, and it can consist of both (fully)

connected perceptrons, convolution operators, or combinations of these and other types of layers. In the present context, it is important that the complexity of the inner network is high enough that the desired mapping can be resolved and small enough such that overfitting will not be an issue.

When a neural network is trained using the training data set, its free parameters are adjusted to minimize a specific loss function, that measures the difference in the expected output from the training data, $\mathbf{n}^i$, set and the output of the neural network, $\hat{\mathbf{n}}$. For the methodology presented here, it is the choice of the loss function, and type of activation function for the output layer that is critical, to allow estimation of properties of the posterior distribution.

In general, a feature n can refer to a continuous parameter (such as velocity, resistivity, temperature), or a discrete parameter (such as lithology type, event type). Each type of parameter requires a specific choice of loss and activation function to estimate a specific property of the posterior distribution.

In the following $\hat{\mathbf{n}}^i$ will be the output of a neural network for the i'th entry in the training data set, $\mathbf{d}^i_{sim} \mapsto \hat{\mathbf{n}}^i$.

2.2.1 Regression type

Let $\mathbf{n}$ represent N_r continuous parameters, and that we wish to estimate the mean and covariance, $N(\widetilde{\mathbf{n}}, \widetilde{\mathbf{C}_n})$, of the posterior distribution $\sigma(\mathbf{n})$. Note the posterior distribution need not be distributed according to a Gaussian distribution. $N(\widetilde{\mathbf{n}}, \widetilde{\mathbf{C}_n})$ is simply a statistical property of the posterior distribution.

The probability that an estimated $\hat{\mathbf{n}}$ is a realization of $N(\widetilde{\mathbf{n}}, \widetilde{\mathbf{C}_n})$ is given by

$$f(\hat{\mathbf{n}}|N(\widetilde{\mathbf{n}}, \widetilde{\mathbf{C}_n})) = exp(-0.5 \ (\hat{\mathbf{n}} - \widetilde{\mathbf{n}}) \ \widetilde{\mathbf{C}_n}^{-1} \ (\hat{\mathbf{n}} - \widetilde{\mathbf{n}})^T). \tag{5}$$

The values of the mean and covariance that maximizes Eqn. 5, can be found by minimizing the loss function, in form og the negative log-likelihood loss function, that is

$$LOSS(\widetilde{\mathbf{n}}) = -\log(f(\mathbf{n_{NN}}). \tag{6}$$

Therefore, any neural network that uses the loss function in Eqn. 6, will lead to an estimate of the mean and covariance representing of the posterior distribution $\sigma(\mathbf{n}|\mathbf{d}_{obs})$. Typically no activation function is used for regression-type neural networks as the output could have any value.

Eqn. 6 is not widely used as a loss function, but is readily available using for example the tensorflow probability extension to tensorflow (Abadi et al., 2015; Dillon et al., 2017). The tensorflow-probability extension provides easy access to many statistical models, from which posterior statistics can be obtained, simply by using the log-likelihood for a specific distribution as the loss function.

To represent the posterior mean and full covariance, an output layer of $N_r + N_r^2$ nodes must be used. If only the posterior mean and variance is to be estimated, an output layer of $N_r + N_r$ nodes is needed. If only the posterior mean is of interest an output layer of N_r nodes is needed, and minimizing Eqn. 6 is similar to minimizing the widely used mean squared error loss function (Bishop et al., 1995).

2.2.2 Classification

When $\mathbf{n}$ represent N_r discrete parameters with nc classes, the goal can be to estimate the posterior probability of each of the nc classes given some data $\mathbf{d}_{obs}$.

Let $\hat{\mathbf{n}} = [\hat{n}_1, \hat{n}_2, ..., \hat{n}_{no}]$ represent the predicted probability of each possible outcome given some observed data $\mathbf{d}_{obs}$, as predicted by a neural network. Let $\mathbf{n^i} = [n_1^i, n_2^i, ..., n_{no}^i]$

represent the probability of the specific outcome i in the training image data set, where each entry is zero, except for the entropy that represent the specific outcome, which is 1. The likelihood that $\mathbf{n^i}$ is a realization of $\mathbf{\hat{n}^i}$ is then given by

$$f(\mathbf{n^i}|\mathbf{\hat{n}}\ ) = \prod_{j=1}^{no} \hat{n}_j^{n_j} \tag{7}$$

The choice of $\mathbf{\hat{n}^i}$ that maximizes Eqn. 7 can be found by minimizing the negative log-likelihood given by the loss function

$$LOSS(\mathbf{\hat{n}}) = -\log(f(\mathbf{n^i}|\mathbf{\hat{n}})) \tag{8}$$

$$= -\sum_{j=1}^{no} n_j \log(\hat{n}_j). \tag{9}$$

Eqn. 9 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, as it forces all probabilities to be in the range 0 to 1, and ensures that $\sum_{j=1}^{no} \hat{n}_j = 1$, such that the output parameters can be interpreted as a probability. A neural network minimizing Eqn. 9, using the softmax activation function in the output layer, therefore locate the maximum-likelihood of Eqn. 7, which represent directly $\sigma(\mathbf{n}|\mathbf{d}_{obs})$.

2.3 multiple data set

The methodology can be trivially extended to account for multiple data types. In case two types of data, A and B, are available (each with a specific forward and noise model), one can create training data sets for both types of data as

$$[\mathbf{N}^*; \mathbf{M}^*; \mathbf{D}^*; [\mathbf{D}_{sim}^{A*}, \mathbf{D}_{sim}^{B*}]], \tag{10}$$

and use the methodology described above to compute properties of the posterior distrubtion $\sigma((m,n))$.

2.4 Multiple prior models

If multiple prior models are available, one can consider using each prior in turn, or to mix the priors constructing $\mathbf{M}^*$ consisting of realizations of all prior models (proportional to one's prior belief in each prior). In the latter case, one can then trivially design a feature vector that contains the index of the prior, from which the posterior probability of the type of prior can be determined directly using the classification approach described above.

Results

As an example the methodology described above is applied to inversion of airborne electromagnetic (AEM) data from Morrill, Nebraska (Smith et al., 2010; Abraham et al., 2012).

These data have been analyzed previously using probabilistic inversion using a 1D transdimensional prior favoring fewer isotropic layers (Minsley, 2011; Minsley et al., 2021) and a 1D prior based on a geostatistical spatially correlated prior model (T. M. Hansen & Minsley, 2019). In these cases, inversion of a sounding took 5-10 minutes of a single CPU. (T. M. Hansen, 2021) used a localized rejection sampler to sample the posterior, using around 1 second per sounding, using the same information as in T. M. Hansen and

Minsley (2019). This work relied on the construction of a lookup table that is similar to the prior realizations corresponding to noise-free data $[\mathbf{M}^*, \mathbf{D}^*]$. Below, results obtained using the neural network approach will be compared to results obtained using this extended rejection sampler.

For this case, the same parameterization, noise model, and (initially) prior model is used as defined in T. M. Hansen and Minsley (2019). A 1D subsurface model, $\mathbf{m}$, is parameterized through $N_m = 125$ model parameters, representing the subsurface resistivity in 1m thick layers, from the surface down to 125 m depth. Initially, a Gaussian type prior with a trimodal 1-D marginal distribution is used, specifically $\rho_3(\mathbf{m}$ as defined in T. M. Hansen and Minsley (2019). The noise of the AEM data is assumed to be independent uncorrelated zero-mean Gaussian noise, with a standard deviation of 5 ppm (parts per million) plus 5 percent of the data value.

2.5 The posterior mean and standard deviation

Figure 1a shows the mean, (and the standard deviation as transparency) of the posterior distribution obtained using the extended rejection sampler with a lookup table of 100000 sets of models and noise-free data, $[\mathbf{M}^*, \mathbf{D}^*]$. See T. M. Hansen (2021) for details.

Using the same lookup table, a training data set of 100000 models and corresponding simulated data with noise is created as $[\mathbf{M}^*, \mathbf{D}^*_{sim}]$ and used as training data for regression type neural network, with noisy simulated data as input $N_d = 12$, and the mean and standard deviation of 125 resistivities, i.e. $n_{out} = 2 * Nm = 250$ parameters as output.

A fully connected multi-layer perceptron model, using 12 nodes in the input layers, 2 hidden layers with 40 nodes each, and 250 nodes in the output layer is constructed. The network is trained using 90% of the training data set, while 10% is used for validation. The loss function is based on Eqn. 6, using only the diagonal of $\widetilde{\mathbf{C}}_n$ (no correlation is estimated), which is minimized using Adam optimizer (Kingma & Ba, 2014).

Figure 1b shows the resulting estimated mean $\hat{n}$, with the standard deviation $\hat{\sigma}_n$ used for transparency (high standard deviation leads to high transparency). The results are strikingly similar to Figure 1a, with the results obtained using machine learning slightly more informed. One reason for the difference is that the neural network approach describes the posterior mean and standard deviation directly, while it is computed from 100 realizations of $\sigma(\mathbf{m})$ using the extended rejection sampling approach.

Training of the network takes around 35 minutes. Once trained, the prediction of the posterior mean and standard deviation for the 451 soundings takes around 4 ms. Around 100000 1D soundings can be analyzed per second. A similar analysis will take around 451 seconds using the extended rejection sampler (T. M. Hansen, 2021), and around 6 hours using the extended Metropolis algorithm (T. M. Hansen & Minsley, 2019).

2.6 Classification: probability of interfaces

For each of the N_r generated models in $\mathbf{M}^*$ a 'feature' n_{int} is estimated that defines whether the resistivity varies above 50% between neighboring model parameters. n_{int} thus represent a classification of 'interface' vs 'no interface'.

A fully connected multi-layer perceptron model, using 12 nodes in the input layers, 2 hidden layers with 40 nodes each, and 125 nodes in the output layer is constructed. The output layer represents the probability of having an interface at the location of the 125 model parameters. The network is trained using 90% of the training data set, while
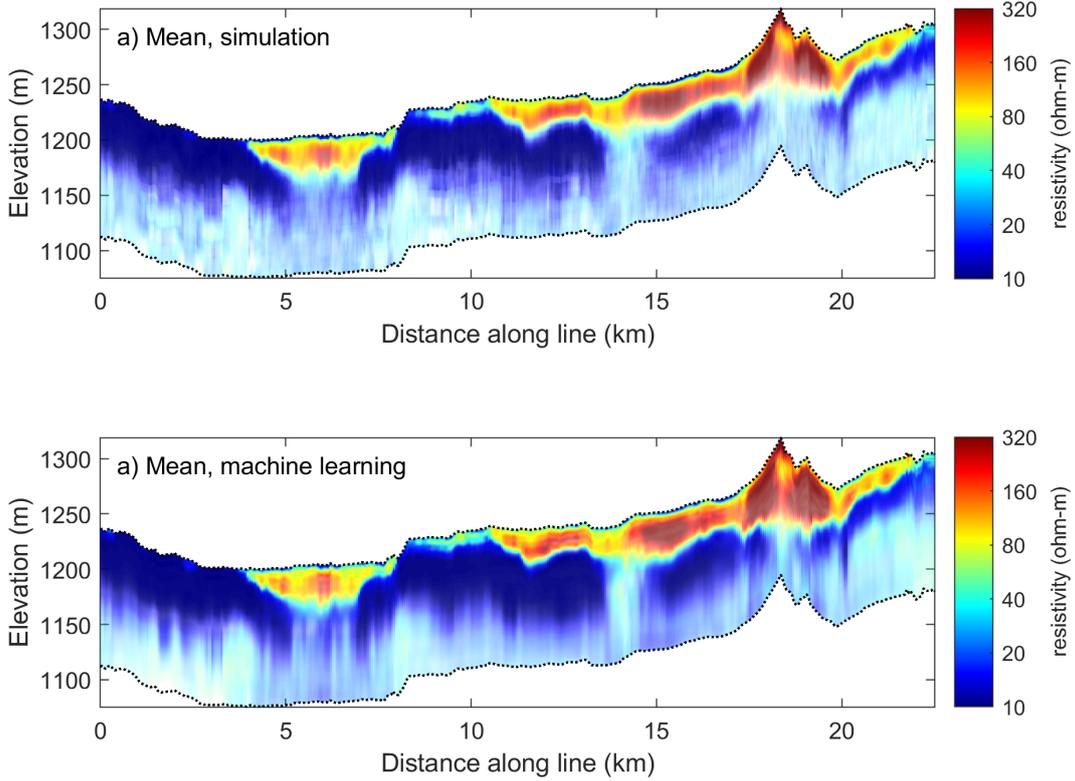
Figure 1. Pointwise mean obtained from $\sigma(m)$ obtained using McMC sampling, and directly using machine learning. Transparency is based on the pointwise estimated standard deviation.

10% is used for validation. The loss function is categorical cross-entropy, Eqn. 9, which is minimized using Adam optimizer (Kingma & Ba, 2014).

Figure 2a shows the posterior probability of layer interface, $\sigma(\mathbf{m}|\mathbf{d}_{obs})$ obtained using the extended rejection sampler (T. M. Hansen, 2021) using a lookup table of size 100000. Figure 2b shows the corresponding result obtained using the trained classification network described above. Again the results are similar, with a more pronounced localization of interfaces using machine learning as opposed to the sampling method.

Training of the network takes around 30 minutes. Once trained, the prediction of the probability of locating an interface for the 451 soundings takes around 1 ms. Around 400000 1D soundings can be analyzed to predict the probability of subsurface layer interfaces per second.

2.7 Facies classification

To illustrate a case of facies classification, a slight variation of the prior model considered above is used. It is constructed by simulation of the spatial distribution of three facies, after which the resistivities are simulated within each facies. This means, that for each realization of the prior, both the resistivity and the facies type are known. n refers to this facies type, that can have three outcomes, '1', '2', and '3'.

From this prior, 100000 models with lithology information are constructed as $\mathbf{N}^*$, that is converted into 100000 models of resistivity information $\mathbf{M}^*$, which is converted into the corresponding noise-free data $\mathbf{D}^*$, and data with simulated noise $\mathbf{D}^*_{sim}$.
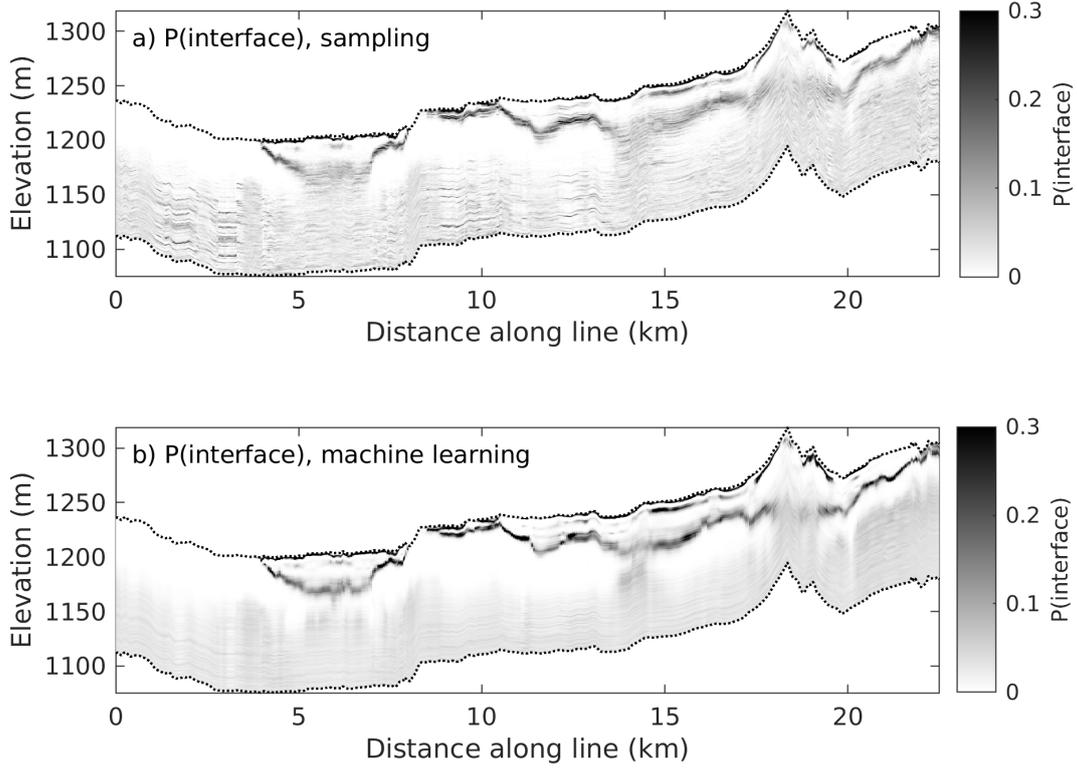
Figure 2. Pointwise probability of boundary (sharp interface) obtained using a) sampling methods, and b) using machine learning.

This means a lookup table $[\mathbf{N}^*, \mathbf{D}^*]$ can be constructed for use with the extended rejection sampler. The posterior probability of each lithology, obtained using the extended rejection sampler is shown in Figure 3a-c.

Similarly, a training data set $[\mathbf{N}^*, \mathbf{D}^*_{obs}]$ is available, and a neural network can be constructed similar to the one above, except that three outcomes are possible. The results, the probability of the lithology given data, is shown in Figure 3d-f.

Again, the results obtained by sampling $\sigma(\mathbf{n})$, Figures 3a-c, and direct computation using machine learning, Figures 3e are very similar.

Training of the network takes around 60 minutes. Once trained, the prediction of the posterior probability of lithology for the 125 model parameters in the 451 soundings takes around 13 ms.

3 Discussion

The machine learning methodology presented above provides a very fast approach for the computation of Gaussian statistics of the posterior distribution using continuous model parameters, and the direct computation of the full posterior distribution using categorical model parameters, of probabilistically formulated inverse problems.

The methodology is relatively easy to use and requires only that one can sample the prior, solve the forward problem, and evaluate the noise model. Then a machine learning algorithm, such as a multi-layer perceptron neural network, can be used to estimate properties of the posterior distribution.
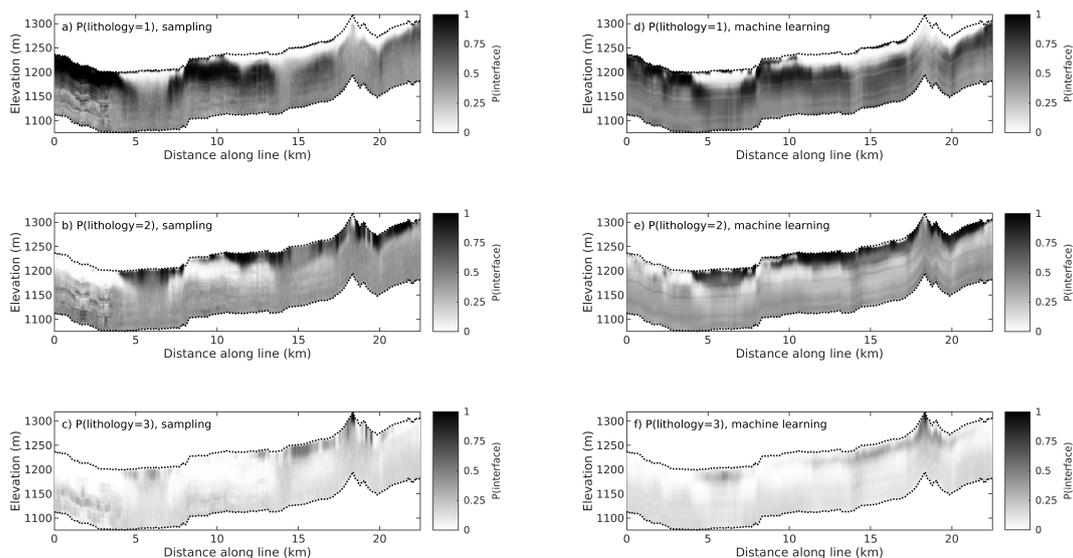
Figure 3. Posterior probability of lithology 1, 2, and 3, using a-c) sampling and d-f) machine learning.

The structure of the neural network structure considered here is quite simple and should be readily accessible with freely available open-source tools [1]

For the specific problem of inverting airborne EM data, the time it takes to construct the lookup table, and train the network, is not very significant, as the network will be applied to potentially many tens of thousands of individual sounding.

The real potential of access to such fast methods is that it allows end-users to try out multiple prior models/and noise models instantly, even on more realistic-sized data sets, with many thousands of soundings.

With this type of efficiency, it could be time to not disregard probabilistic inversion due to its high computational costs, but instead to embrace it, since it can account for much more complex, and realistic information, than deterministic and linearized inversion methods.

4 Conclusions

A simple, yet very powerful, approach to probabilistic inversion has been proposed. Its application requires that one can simulate sets of examples of what is known. 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 this set of models and data, a set of corresponding features of the model parameter and simulated noisy data can be obtained that represent, up to the limit of the finite size set of training data set, all known information.

From such features and data, the posterior statistics of the feature given the data can be obtained from, in this case, a neural network with an appropriately chosen output layer and activation function, whose free parameters have been obtained by mini-

---

[1] Pyhton notebooks will be made available at `http://github.com/cultpenguin/ip_and_ml/`.

mizing an appropriate loss function. This leads to fast and accurate estimation of posterior statistics.

A case study exemplified the methodology for inversion of AEM data and shows posterior statistics similar to those obtained using sampling methods, using a fraction of the computation time. This allows using and testing multiple prior models, for multiple features related to the prior distributions, in a fully probabilistic setting using only very limited computational resources.

Acknowledgments

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., . . . Kthe, U. (2019). *Analyzing inverse problems with invertible neural networks.*

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.

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.

Constable, S. C., Parker, R. L., & Constable, C. G. (1987). Occams 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.

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

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.

Hansen, P. C. (1992). Analysis of discrete ill-posed problems by means of the l-curve. *SIAM review*, *34*(4), 561–580.

Hansen, T. M. (2021). *Efficient probabilistic inversion using the rejection samplerexemplified on airborne em data* (Vol. 224) (No. 1). Oxford University Press.

Hansen, T. M., Cordua, K., Looms, M., & 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. (2017). Efficient monte carlo sampling of inverse problems using a neural network-based forwardapplied to gpr crosshole travel-time 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 problemsexemplified for crosshole tomography. *Geophysics*, *79*(3), H1–H21.

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.

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

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.

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. 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.

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*.

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).

Zhang, X., & Curtis, A. (2021). Bayesian geophysical inversion using invertible neural networks.
doi: 10.31223/X50K6J