

Parameterizing uncertainty by deep invertible networks, an application to reservoir characterization

Gabrio Rizzuti, Ali Siahkoohi, Philipp A. Witte, and Felix J. Herrmann

School of Computational Science and Engineering

Georgia Institute of Technology

ABSTRACT

Uncertainty quantification for full-waveform inversion provides a probabilistic characterization of the ill-conditioning of the problem, comprising the sensitivity of the solution with respect to the starting model and data noise. This analysis allows to assess the confidence in the candidate solution and how it is reflected in the tasks that are typically performed after imaging (e.g., stratigraphic segmentation following reservoir characterization). Classically, uncertainty comes in the form of a probability distribution formulated from Bayesian principles, from which we seek to obtain samples. A popular solution involves Monte Carlo sampling. Here, we propose instead an approach characterized by training a deep network that “pushes forward” Gaussian random inputs into the model space (representing, for example, density or velocity) as if they were sampled from the actual posterior distribution. Such network is designed to solve a variational optimization problem based on the Kullback-Leibler divergence between the posterior and the network output distributions. This work is fundamentally rooted in recent developments for invertible networks. Special invertible architectures, besides being computational advantageous with respect to traditional networks, do also enable analytic computation of the output density function. Therefore, after training, these networks can be readily used as a new prior for a related inversion problem. This stands in stark contrast with Monte-Carlo methods, which only produce samples. We validate these ideas with an application to angle-versus-ray parameter analysis for reservoir characterization.

INTRODUCTION

A classical probabilistic setup for full-waveform inversion (FWI, Tarantola, 2005) starts from the following assumptions on the data likelihood:

$$\mathbf{d} = F(\mathbf{m}) + \mathbf{n}. \quad (1)$$

Here, \mathbf{d} represents seismic data, F is the forward modeling map, \mathbf{m} collects the unknown medium parameters of interest, and \mathbf{n} is random noise. We are assuming that \mathbf{d} and \mathbf{n} are independent (knowing \mathbf{m}), and $\mathbf{n} \sim p_{\text{noise}}(\mathbf{n})$ is normally distributed according to $\mathcal{N}(0, I)$. Equation (1) then defines the likelihood density function

$$\mathbf{d}|\mathbf{m} \sim p_{\text{like}}(\mathbf{d}|\mathbf{m}) = p_{\text{noise}}(\mathbf{d} - F(\mathbf{m})).$$

In the Bayesian framework, the second ingredient is a prior distribution on the model space

$$\mathbf{m} \sim p_{\text{prior}}(\mathbf{m}). \quad (2)$$

Since for seismic experiments we do not have direct access to the Earth’s interior (except for localized information in the form

of well logs), the prior is typically hand crafted based on, e.g., Thikhonov or total variation regularization (see, for example, Esser et al., 2018; Peters et al., 2019). Alternatively, when this type of information is available, a prior can be implicitly encoded via a deep network (Mosser et al., 2019).

The posterior distribution, given data \mathbf{d} , is readily obtained from the Bayes’ rule:

$$p_{\text{post}}(\mathbf{m}|\mathbf{d}) = p_{\text{like}}(\mathbf{d}|\mathbf{m})p_{\text{prior}}(\mathbf{m})/p_{\text{data}}(\mathbf{d}), \quad (3)$$

where $p_{\text{data}}(\mathbf{d}) = \int p_{\text{like}}(\mathbf{d}|\mathbf{m})p_{\text{prior}}(\mathbf{m})d\mathbf{m}$. Sampling from the posterior probability (3) is the goal of uncertainty quantification (UQ). Besides the usual maximum-a-posteriori estimator (MAP)

$$\mathbf{m}_{\text{MAP}} = \arg \max_{\mathbf{m}} p_{\text{post}}(\mathbf{m}|\mathbf{d}), \quad (4)$$

being able to sample from the posterior allows to approximate the conditional mean and point-wise standard deviation:

$$\begin{aligned} \mu &= \mathbb{E}(\mathbf{m}|\mathbf{d}) = \int \mathbf{m} p_{\text{post}}(\mathbf{m}|\mathbf{d}) d\mathbf{m}, \\ \sigma^2 &= \mathbb{E}((\mathbf{m} - \mu)^2|\mathbf{d}) = \int (\mathbf{m} - \mu)^2 p_{\text{post}}(\mathbf{m}|\mathbf{d}) d\mathbf{m}. \end{aligned} \quad (5)$$

A great deal of research around UQ is devoted to Monte Carlo sampling (Robert and Casella, 2004), especially through Markov chains (MCMC) when dealing with high-dimensional problems like FWI (see, e.g., Malinverno and Briggs, 2004; Bennett and Malinverno, 2005; Malinverno and Parker, 2006; Martin et al., 2012; Sambridge et al., 2013; Ely et al., 2018; Parno and Marzouk, 2018; Zhu and Gibson, 2018; Izzatullah et al., 2019; Visser et al., 2019; Zhao and Sen, 2019). A Markov chain is a stochastic process that describes the evolution of the model distribution, and is designed to match the target distribution at steady state. A particularly egregious example is offered by the so-called Langevin dynamics (Neal et al., 2011):

$$\mathbf{m}_{n+1} = \mathbf{m}_n - \frac{\epsilon_n}{2} \nabla_{\mathbf{m}} \log p_{\text{post}}(\cdot|\mathbf{d})|_{\mathbf{m}_n} + \eta_n, \quad \eta_n \sim \mathcal{N}(0, \epsilon_n I), \quad (6)$$

which is akin to stochastic gradient descent where the update direction is perturbed by random noise η_n . Under some technical assumptions on the step-length decaying (Welling and Teh, 2011), collecting \mathbf{m}_n ’s as in Equation (6) is equivalent to sample from the posterior distribution (after an initial “burn-in” phase). While gradient-based sampling methods such as Langevin dynamics are gaining popularity in machine learning, we must face the computational challenge given by the evaluation and differentiation of the seismic modeling map involved in Equation (6), which needs be repeated for multiple sources thus resulting in an expensive scheme. Nevertheless, source encoding and network reparameterization could alleviate these issues (Herrmann et al., 2019; Siahkoohi et al., 2020a,b).

In this paper, we take an alternative route to UQ by employing model-space valued deep networks defined on a latent space

endowed with an easy-to-sample distribution (e.g. normal), as a way to mimic random sampling from the posterior (when properly trained). This is made possible by a special class of invertible networks for which the output density function can be computed analytically (Dinh et al., 2016; Kingma and Dhariwal, 2018; Lensink et al., 2019; Kruse et al., 2019). As such, these networks encapsulate a new prior, which can be used for subsequent tasks (as, for example, a similar imaging problem).

METHOD

In this section, we discuss the theoretical foundations of the proposed method for UQ via generative networks. We start by discussing a well-known notion of discrepancy between probabilities, the Kullback-Leibler divergence. Our program is suggested in Kruse et al. (2019), and entails the minimization of the posterior divergence with a known distribution pushed forward by an invertible mapping. The resulting optimization problem is further restricted over a special class of invertible networks, for which the log-determinant of the Jacobian can be computed.

The Kullback-Leibler (KL) divergence of two probabilities defined on the model space $\mathbf{m} \in M$, with density p and q respectively, is given by

$$\text{KL}(p||q) = \int p(\mathbf{m}) \log \frac{p(\mathbf{m})}{q(\mathbf{m})} d\mathbf{m}. \quad (7)$$

Let us fix, now, a normally distributed random variable

$$\mathbf{z} \sim p_Z(\mathbf{z}), \quad (8)$$

taking values on a latent space Z , and a deterministic map

$$T : Z \rightarrow M. \quad (9)$$

Selecting random samples \mathbf{z} and feeding it to T generates a random variable $T(\mathbf{z})$ with values in M , which is referred to as the push-forward of \mathbf{z} through T . The push-forward density will be denoted by

$$\mathbf{m} \sim T_{\#}p_Z(\mathbf{m}). \quad (10)$$

Our goal is the solution of the following minimization problem:

$$\min_T \text{KL}(T_{\#}p_Z||p_{\text{post}}). \quad (11)$$

Note that the KL divergence is notoriously not symmetric, and the order of the probabilities appearing in Equation (11) is designed. A more convenient form of Equation (11) is

$$\text{KL}(T_{\#}p_Z||p_{\text{post}}) = \mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} -\log p_{\text{post}}(T(\mathbf{z})|\mathbf{d}) - H(T_{\#}p_Z), \quad (12)$$

where

$$H(p) = - \int p(\mathbf{m}) \log p(\mathbf{m}) d\mathbf{m} \quad (13)$$

denotes the entropy of p . The loss function in Equation (12) ensures that the outputs of T adhere to the data without producing mode collapse (e.g. $T(\mathbf{z}) \equiv \mathbf{m}_{\text{MAP}}$, a low-entropy configuration).

While the first term in the KL divergence can be approximated by sample averaging, i.e. $\mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} f(\mathbf{z}) \approx \frac{1}{N} \sum_{i=1}^N f(\mathbf{z}_i)$, calculating the entropy of a push-forward distribution is not trivial (see, for instance, Kim and Bengio, 2016). In general, we typically can only evaluate $T(\mathbf{z})$ without a direct way to explicitly compute its density, as in Equation (10), for a given \mathbf{m} . A popular—but computationally expensive—workaround is based on generative adversarial networks (GAN, Goodfellow et al., 2014), where the discriminator is in fact related to the generator push-forward density (Che et al., 2020). Other approaches involve variational autoencoders (Kingma and Welling, 2013) and expectation-maximization techniques (Han et al., 2017). Nevertheless, if we focus on the specialized class of invertible mappings T (whose existence requires that $\dim Z = \dim M$ as manifolds), we obtain

$$H(T_{\#}p_Z) = H(p_Z) + \mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} \log |\det J_T(\mathbf{z})| \quad (14)$$

thanks to the change-of-variable formula. Here, $J_T(\mathbf{z})$ is the Jacobian of T with respect to the input \mathbf{z} .

Invertible networks are an enticing family of choice for T . This is especially true for sizable inverse problems, since they do not require storing the input before every activation functions (such as ReLU) and the memory overhead can be kept constant (Putzky and Welling, 2019). Furthermore, recent developments has lead to architectures that allows the computation (and differentiation) of the log-determinants that appears in Equation (14). As notable examples, we mention non-volume-preserving networks (NVP, Dinh et al., 2016), generative flows with invertible 1-by-1 convolutions (Glow, Kingma and Dhariwal, 2018), and hierarchically invertible neural transport (HINT, Kruse et al., 2019). These examples share the usage of special bijective layers whose Jacobian has a triangular structure (Knothe-Rosenblatt transformations, Parno and Marzouk, 2018). Hyperbolic networks (Lensink et al., 2019) are another instance of invertible maps, although volume preserving with null log-determinants. An invertible version of classical residual networks (resnets) is also offered in Behrmann et al. (2018), which however requires an implicit estimation of the log-determinant. Yet another example of invertible architecture, though not allowing log-determinant calculation, is discussed in Putzky and Welling (2019) for inverse problem applications.

By restricting the problem in Equation (11) to the family of suitable invertible networks just described, we end up with the following optimization:

$$\min_{\theta} \mathbb{E}_{\mathbf{z} \sim p_Z(\mathbf{z})} \frac{1}{2} \|\mathbf{d} - F(G_{\theta}(\mathbf{z}))\|^2 - \log p_{\text{prior}}(G_{\theta}(\mathbf{z})) - \log |\det J_{G_{\theta}}(\mathbf{z})|, \quad (15)$$

where the unknowns θ represent the parameters (such as weights and biases) of the invertible network

$$G_{\theta} : Z \rightarrow M, \quad G_{\theta}^{-1} : M \rightarrow Z. \quad (16)$$

We employ a stochastic gradient approach (ADAM, Kingma and Ba, 2014) with random selection of normally distributed mini-batches of \mathbf{z} 's. The expectation in Equation (15) is then replaced by Monte-Carlo averaging.

The solution of problem (15) comes in the form of a network G_θ , whose evaluation over random \mathbf{z} 's generates models \mathbf{m} 's as if they were sampled from the posterior distribution in Equation (3). After training, we can easily compute conditional mean and point-wise standard deviation as in Equation (5), or even higher-order statistics. Since G_θ possesses a suitable architecture that explicitly encodes its output density, we can reuse this result as a new prior for adjacent regions, time-lapse imaging, and so on.

UNCERTAINTY QUANTIFICATION FOR RESERVOIR CHARACTERIZATION

We now apply the setup described in the previous section, Equation (15), to full-waveform inversion for reservoir characterization under some simplifying assumptions. We postulate a purely acoustic model (constant density), with laterally homogeneous properties (e.g., 1D). Waves propagate in a 2D medium. The Radon transform yields the ‘‘1.5-D’’ wave equation

$$(\mathbf{m}(z) - p^2)\partial_{tt}\mathbf{u}(t, z, p) - \partial_{zz}\mathbf{u}(t, z, p) = \mathbf{q}(t)\delta(z), \quad (17)$$

where \mathbf{u} is a wavefield, \mathbf{q} a source wavelet, and p the ray parameter. Time-depth is indicated by (t, z) . Data is collected at the surface $z = 0$.

We synthesize a compressional velocity well-log from the open-source Sleipner dataset in Figure 1 (refer to Haffinger et al., 2018, for a previous AVP study). See the acknowledgements for more info. The true model is obtained by smoothing and subsampling the well-log slowness on a 3 m grid, and a background is obtained by severe smoothing. We select a source wavelet with corner frequencies (5, 30, 50, 80) Hz, and a angle range from $\theta = 12^\circ$ to $\theta = 30^\circ$, where $p = \sin\theta\sqrt{\mathbf{m}(z=0)}$. Synthetic data are generated by finite-differences, and random Gaussian noise is injected with SNR = 0 dB. Data are shown in Figure 2.

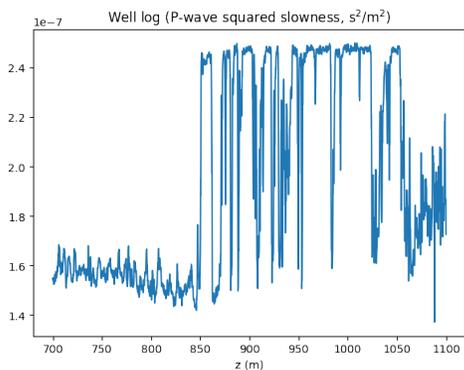


Figure 1: Compressional velocity well log for the Sleipner synthetic experiment

As of invertible architecture, we choose hyperbolic networks (Lensink et al., 2019), due to their inherent stability, augmented with one activation normalization layer (similarly to Kingma and Dhariwal, 2018) acting diagonally along the depth dimen-

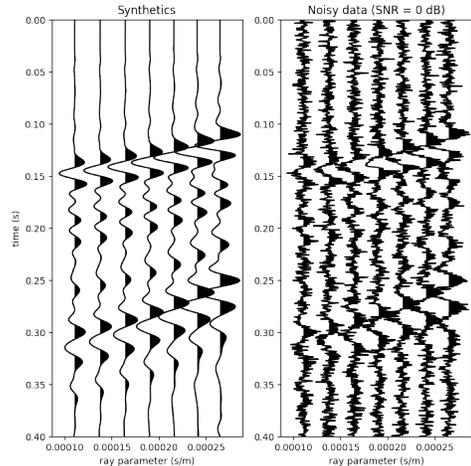


Figure 2: Angle-versus-ray parameter data

sion. This modification is necessary to obtain a non-volume preserving transformation, that is with a non-null log-determinant (otherwise the network entropy will remain constant, cf. Equation 14).

The inversion results are compared in Figure 3. We plot the conditional mean μ , as defined in Equation (5), with confidence intervals determined by the pointwise conditional standard deviation σ . The conditional μ and σ are calculated with sample average of the network outputs $G_\theta(\mathbf{z})$. Additionally we show multiple samples from the estimated posterior before and after training (Figure 4), which highlight the sharpening of the distribution around the mean once optimized. Note that G_θ was initialized as Gaussian perturbations of the MAP, cf. Figure 4(a). A structural comparison between true model and conditional statistics is made apparent in Figure 5. As expected, the standard deviation has generally higher values where the background model is further from the truth (the deeper portion), and displays peaks at the reflector positions. We also observe that uncertainties tend to grow with depth.

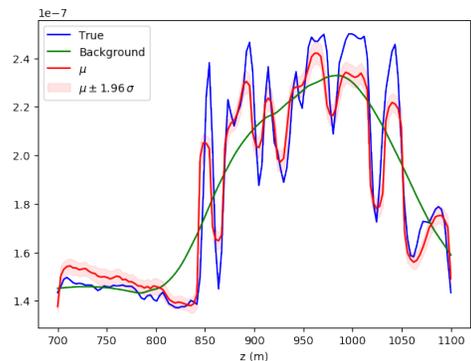


Figure 3: Inversion results with confidence intervals

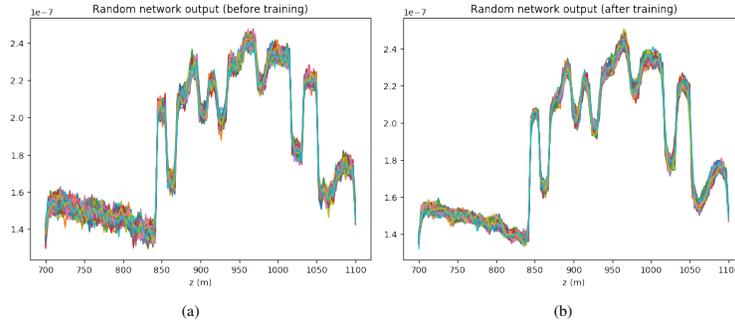


Figure 4: Network realizations: (a) before training (network initialized as Gaussian perturbation of the MAP); (b) after training

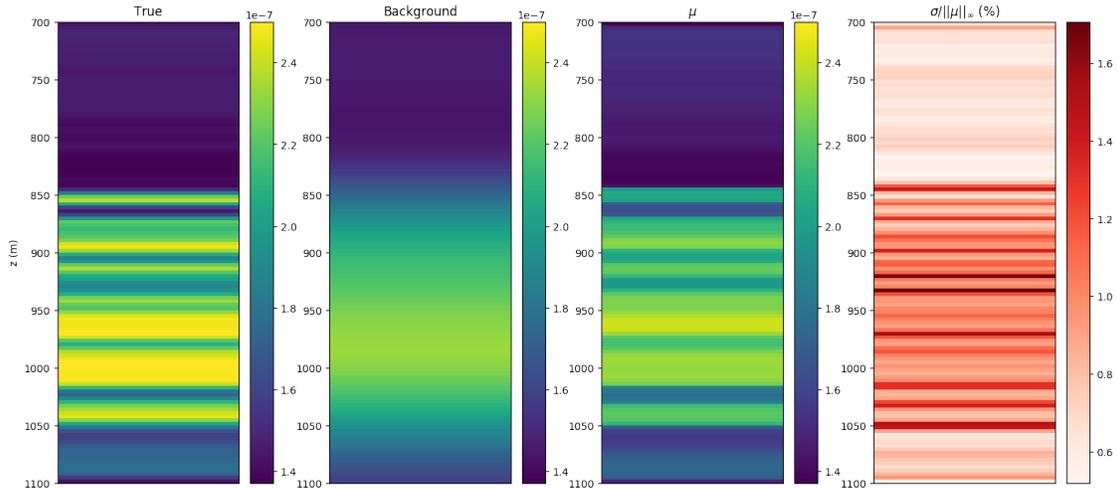


Figure 5: Uncertainty quantification, structural comparison

CONCLUSIONS

We presented an uncertainty quantification framework alternative to more classical MCMC sampling, where an invertible neural network is initialized according to a prior distribution and trained to approximate samples from the posterior distribution, given some data. This can be achieved by minimization of a probability discrepancy loss (the Kullback-Leibler divergence) between the posterior and the network output distribution. We discussed a full-waveform inversion application for reservoir characterization, which demonstrates the ability of this network to qualitatively capture the uncertainties of the problem. Contrary to MCMC methods, the result can be reused as a new prior for related inverse problems since, by construction, we are able to compute the density function of the output distribution. Despite the aim of this project being faithful posterior sampling, we remark that the network architecture does bias the uncertainty analysis (sometimes favorably as in deep prior regularization, Siahkoobi et al., 2020b), and should be carefully factored in. On the computational side, each iteration of stochastic gradient descent requires the calculation of a seismic gradient for the models of a given batch. Therefore, a feasible extension to 2D would certainly need randomized simultaneous sources. Future

work will involve elastics, incorporation of hard constraints (as opposed to priors used as a penalty term), and field data applications.

ACKNOWLEDGMENTS

We are grateful to Equinor and partners for providing open access to the Sleipner data through the CO₂ storage data consortium (<https://co2datashare.org/>).

RELATED MATERIALS

In order to facilitate the reproducibility of the results herein discussed, a Julia implementation of this work is made available on the SLIM github page:

github.com/slimgroup/Software.SEG2020

Several invertible network architectures are implemented as a Julia package in:

github.com/slimgroup/InvertibleNetworks.jl

REFERENCES

- Behrmann, J., W. Grathwohl, R. T. Q. Chen, D. Duvenaud, and J.-H. Jacobsen, 2018, Invertible Residual Networks.
- Bennett, N., and A. Malinverno, 2005, Fast Model Updates Using Wavelets: Multiscale Model. Simul., **3**, 106–130.
- Che, T., R. Zhang, J. Sohl-Dickstein, H. Larochelle, L. Paull, Y. Cao, and Y. Bengio, 2020, Your GAN is Secretly an Energy-based Model and You Should use Discriminator Driven Latent Sampling.
- Dinh, L., J. Sohl-Dickstein, and S. Bengio, 2016, Density estimation using Real NVP.
- Ely, G., A. Malcolm, and O. V. Poliannikov, 2018, Assessing uncertainties in velocity models and images with a fast nonlinear uncertainty quantification method: Geophysics, **83**, R63–R75.
- Esser, E., L. Guasch, T. van Leeuwen, A. Y. Aravkin, and F. J. Herrmann, 2018, Total-variation regularization strategies in full-waveform inversion: SIAM Journal on Imaging Sciences, **11**, 376–406.
- Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, 2014, Generative Adversarial Networks.
- Haffinger, P., F. J. Eyvazi, P. Steeghs, P. Doulgeris, D. Gisolf, and E. Verschuur, 2018, Quantitative Prediction of Injected CO₂ At Sleipner Using Wave-Equation Based AVO: **2018**, 1–5.
- Han, T., Y. Lu, S.-C. Zhu, and Y. N. Wu, 2017, Alternating back-propagation for generator network: Presented at the Thirty-First AAAI Conference on Artificial Intelligence.
- Herrmann, F. J., A. Siahkoochi, and G. Rizzuti, 2019, Learned imaging with constraints and uncertainty quantification: Presented at the Neural Information Processing Systems (NeurIPS) 2019 Deep Inverse Workshop.
- Izzatullah, M., T. van Leeuwen, and D. Peter, 2019, Bayesian uncertainty estimation for full waveform inversion: A numerical study: SEG Technical Program Expanded Abstracts 2019, 1685–1689.
- Kim, T., and Y. Bengio, 2016, Deep Directed Generative Models with Energy-Based Probability Estimation.
- Kingma, D. P., and J. Ba, 2014, Adam: A Method for Stochastic Optimization.
- Kingma, D. P., and P. Dhariwal, 2018, Glow: Generative Flow with Invertible 1x1 Convolutions.
- Kingma, D. P., and M. Welling, 2013, Auto-Encoding Variational Bayes.
- Kruse, J., G. Detommaso, R. Scheichl, and U. Köthe, 2019, HINT: Hierarchical Invertible Neural Transport for Density Estimation and Bayesian Inference.
- Lensink, K., E. Haber, and B. Peters, 2019, Fully Hyperbolic Convolutional Neural Networks.
- Malinverno, A., and V. A. Briggs, 2004, Expanded uncertainty quantification in inverse problems: Hierarchical Bayes and empirical Bayes: Geophysics, **69**, 1005–1016.
- Malinverno, A., and R. L. Parker, 2006, Two ways to quantify uncertainty in geophysical inverse problems: Geophysics, **71**, W15–W27.
- Martin, J., L. C. Wilcox, C. Burstedde, and O. Ghattas, 2012, A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion: SIAM Journal on Scientific Computing, **34**, A1460–A1487.
- Mosser, L., O. Dubrulle, and M. Blunt, 2019, Stochastic Seismic Waveform Inversion Using Generative Adversarial Networks as a Geological Prior: Mathematical Geosciences, **84**, 53–79.
- Neal, R. M., et al., 2011, MCMC using Hamiltonian dynamics: Handbook of Markov chain Monte Carlo, **2**, 2.
- Parno, M. D., and Y. M. Marzouk, 2018, Transport Map Accelerated Markov Chain Monte Carlo: SIAM/ASA Journal on Uncertainty Quantification, **6**, 645–682.
- Peters, B., B. R. Smithyman, and F. J. Herrmann, 2019, Projection methods and applications for seismic nonlinear inverse problems with multiple constraints: Geophysics, **84**, R251–R269.
- Putzky, P., and M. Welling, 2019, Invert to Learn to Invert.
- Robert, C., and G. Casella, 2004, Monte Carlo statistical methods: Springer-Verlag.
- Sambridge, M., T. Bodin, K. Gallagher, and H. Tkalčić, 2013, Transdimensional inference in the geosciences: Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, **371**, 20110547.
- Siahkoochi, A., G. Rizzuti, and F. J. Herrmann, 2020a, A deep-learning based bayesian approach to seismic imaging and uncertainty quantification: 82nd EAGE Conference and Exhibition 2020.
- , 2020b, Uncertainty quantification in imaging and automatic horizon tracking: a Bayesian deep-prior based approach.
- Tarantola, A., 2005, Inverse problem theory and methods for model parameter estimation: SIAM, **89**.
- Visser, G., P. Guo, and E. Saygin, 2019, Bayesian transdimensional seismic full-waveform inversion with a dipping layer parameterization: Geophysics, **84**, R845–R858.
- Welling, M., and Y. W. Teh, 2011, Bayesian Learning via Stochastic Gradient Langevin Dynamics: Proceedings of the 28th International Conference on International Conference on Machine Learning, 681–688.
- Zhao, Z., and M. K. Sen, 2019, A gradient based MCMC method for FWI and uncertainty analysis: SEG Technical Program Expanded Abstracts 2019, 1465–1469.
- Zhu, D., and R. Gibson, 2018, Seismic inversion and uncertainty quantification using transdimensional Markov chain Monte Carlo method: Geophysics, **83**, R321–R334.