
Wave-equation based inversion with amortized variational Bayesian inference

Ali Siahkoohi

School of Computational Science and Engineering,
Georgia Institute of Technology
alisk@gatech.edu

Rafael Orozco

School of Computational Science and Engineering,
Georgia Institute of Technology
rorozco@gatech.edu

Gabrio Rizzuti

Department of Mathematics,
Utrecht University
g.rizzuti@uu.nl

Felix J. Herrmann

School of Computational Science and Engineering,
Georgia Institute of Technology
felix.herrmann@gatech.edu

Abstract

Solving inverse problems involving measurement noise and modeling errors requires regularization in order to avoid data overfit. Geophysical inverse problems, in which the Earth’s highly heterogeneous structure is unknown, present a challenge in encoding prior knowledge through analytical expressions. Our main contribution is a generative-model-based regularization approach, robust to out-of-distribution data, which exploits the prior knowledge embedded in existing data and model pairs. Utilizing an amortized variational inference objective, a conditional normalizing flow (NF) is pretrained on pairs of low- and high-fidelity migrated images in order to achieve a low-fidelity approximation to the seismic imaging posterior distribution for previously unseen data. The NF is used after pretraining to reparameterize the unknown seismic image in an inversion scheme involving physics-guided data misfit and a Gaussian prior on the NF latent variable. Solving this optimization problem with respect to the latent variable enables us to leverage the benefits of data-driven conditional priors whilst being informed by physics and data. The numerical experiments demonstrate that the proposed inversion scheme produces seismic images with limited artifacts when dealing with noisy and out-of-distribution data.

1 Introduction

An inverse problem involves reliably estimating an unknown quantity from noisy indirect observations. This problem is commonly solved using optimization techniques to minimize the difference between predicted and observed data. Solely minimizing the data misfit negatively impacts the quality of the obtained solution due to noise in the data, modeling errors, and a nontrivial null-space of the forward operator [1]. To prevent this, it is crucial to capture and incorporate prior knowledge into the inverse problem [1], e.g., Gaussian or Laplace distribution priors [2–4]. While theoretically understood, these type of priors may lead to undesirable biases in the outcome of inversion.

The purpose of our contribution is to address this challenge by utilizing a formulation that exploits a data-driven conditional prior. To achieve this, following Orozco et al. [5], we train a conditional normalizing flow [NF, 6] to capture the conditional distribution of the unknown, given data, i.e., the posterior distribution. The training involves minimizing an amortized variational inference objective [6–10] using existing training pairs in the form of low-fidelity data and model pairs. After training, we are able to capture the low-fidelity posterior distribution for previously unseen seismic data. We use this network to reparameterize the unknown in an inversion scheme, involving physics-guided data misfit and a Gaussian prior on the NF latent variable. Due to the inherent invertibility of NFs, they can represent any model in the unknown space, which allows them to be used as priors when dealing with out-of-distribution data [5, 11].

There are three key advantages to our proposed method: (1) Data-driven priors make use of available data, such as high-resolution seismic images to capture prior knowledge about the Earth’s subsurface; (2) The use of a conditional prior favors solutions that are consistent with the data, which provides more specific knowledge about the unknown; (3) With the help of our formulation, we combine data-driven priors with conventional physics-based inversion methods, which offers the advantages of data-driven priors without relying solely on them as a black box.

In the following sections, we discuss conditional NFs, trained using an amortized variational inference procedure. Next, we present an inversion scheme for seismic imaging that incorporates conditional priors. We conclude by demonstrating this technique on a realistic seismic imaging problem involving noisy and out-of-distribution data.

2 Amortized variational inference

The problem setup entails applying variational inference [7] to approximate the posterior distribution [1] associated with the inverse problem $\mathbf{y} = F(\mathbf{x}) + \epsilon$, where $\mathbf{y} \in Y$ represents the observed data, $\mathbf{x} \in X$ unknown model, ϵ possibly non-Gaussian noise, and $F : X \rightarrow Y$ the possibly nonlinear forward operator. In the context of amortized variational inference [6], we wish to approximate the posterior distribution associated with this inverse problem for previously unseen data. This method has computational advantages as it does not require solving an additional instance of variational inference for new data. In this work, we choose NFs [12] that due to their invertibility (up to numerical precision) can be used to approximate a target distribution, of which we have only samples. NFs can be adapted to sample from the conditional distribution $p(\mathbf{x} | \mathbf{y})$ by using a block-triangular construction [13], $T_w(\mathbf{y}, \mathbf{x}) = (T_{w_1}(\mathbf{y}), T_{w_2}(\mathbf{y}, \mathbf{x}))$ with $\mathbf{w} = (\mathbf{w}_1, \mathbf{w}_2)$. The conditional NF $T_w : Y \times X \rightarrow Z \times Z$, which takes as input data and model pairs (\mathbf{y}, \mathbf{x}) , aims to yield two normally distributed outputs in the latent space $Z \times Z$. Training objective is based on minimizing the Kullback-Leibler divergence between the NF output distribution and the Gaussian latent distribution [6]:

$$\arg \min_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n \left[\frac{1}{2} \|T_w(\mathbf{y}^{(i)}, \mathbf{x}^{(i)})\|_2^2 - \log \left| \det \nabla_{(\mathbf{y}, \mathbf{x})} T_w(\mathbf{y}^{(i)}, \mathbf{x}^{(i)}) \right| \right]. \quad (1)$$

In the above objective, the ℓ_2 -norm follows from a Gaussian assumption on the latent variables and the second term is a regularization term that avoids T_w from converging to trivial solutions—e.g., $T_w := \mathbf{0}$. Computing $\det \nabla_{(\mathbf{y}, \mathbf{x})} T_w(\mathbf{y}, \mathbf{x})$ and its gradient adds almost no extra cost because of the particular design of invertible networks [14]. Following training, we can obtain samples from conditional distribution $p(\mathbf{x} | \mathbf{y})$ via $T_{w_2}^{-1}(T_{w_1}(\mathbf{y}), \mathbf{z}) \sim p(\mathbf{x} | \mathbf{y})$, $\mathbf{z} \sim p_z(\mathbf{z})$ [6, 13]. This amounts to feeding the latent code associated with observed data, i.e., $T_{w_1}(\mathbf{y})$, and Gaussian samples $\mathbf{z} \sim p_z(\mathbf{z})$ into the inverse network, $T_{w_2}^{-1}$. These samples may be used for Bayesian inference if we have an ideal training dataset [9, 15–17]. However, such an assumption is rarely correct in geophysical applications due to Earth’s strong heterogeneity [18–20], which highlights the importance of devising formulations that are robust to changes in data distribution during inference.

3 Seismic imaging with data-driven conditional priors

Using multiple processed shot records, $\{\mathbf{d}_i\}_{i=1}^{n_s}$, seismic imaging aims to estimate the short-wavelength component of the Earth’s subsurface squared-slowness model, denoted by $\delta \mathbf{m}$. The linearized Born scattering operator $\mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)$ relates the unknown seismic image $\delta \mathbf{m}^*$, to data, the i th source signature, \mathbf{q}_i , and the background squared-slowness model \mathbf{m}_0 . This relation can be

written as

$$\mathbf{d}_i = \mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)\delta\mathbf{m}^* + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I}). \quad (2)$$

Noise is denoted by ϵ_i and represents measurement noise and linearization errors, which for simplicity is assumed to be distributed according to a zero-centered Gaussian distribution with known covariance $\sigma^2\mathbf{I}$. We train a NF on pairs of low- and high-fidelity seismic images via the amortized variational inference objective by choosing \mathbf{x} to represent high-fidelity migrated images and $\mathbf{y} := \delta\mathbf{m}_{\text{RTM}}$ corresponding to low-fidelity reverse-time migrated images obtained by the process of demigration, followed by adding noise and migration. After training, the conditional NF captures the low-fidelity seismic imaging posterior distribution. In order to obtain a high-fidelity seismic image maximum a posteriori (MAP) estimate, we propose to reparameterize $\delta\mathbf{m}$ vis the pretrained NF and solve the optimization problem

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \frac{1}{2\sigma^2} \left[\sum_{i=1}^{n_s} \|\mathbf{d}_i - \mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)T_{w_2}^{-1}(T_{w_1}(\delta\mathbf{m}_{\text{RTM}}), \mathbf{z})\|_2^2 \right] + \frac{1}{2}\|\mathbf{z}\|_2^2, \quad (3)$$

followed by mapping $\delta\mathbf{m}_{\text{MAP}} := T_{w_2}^{-1}(T_{w_1}(\delta\mathbf{m}_{\text{RTM}}), \hat{\mathbf{z}})$. We initialize optimization problem 3 with $\mathbf{z}_0 = \mathbf{0}$. This initialization and a Gaussian prior on \mathbf{z} regularize the inversion by favoring solutions that are likely samples of the low-fidelity posterior distribution [11]. NFs’ inherent invertibility allows them to represent any image $\delta\mathbf{m}$ in the solution space. This limits the potentially negative bias of the conditional prior in domains where access to high-fidelity training data is limited. We demonstrate this through a numerical experiment in the next section.

4 Numerical experiments

We propose a realistic example in which we create 4750 2D training pairs of low- and high-fidelity seismic images, which the latter are 3075 m \times 5120 m sections extracted from the shallow part of [Parihaka](#) [21] pretrack Kirchhoff migration dataset. The low-fidelity images are obtain by migrating noisy synthetic data obtained from the high-fidelity images according to Equation 2. The acquisition geometry involves 102 shot records, 204 fixed receivers, Ricker wavelet with a central frequency of 30 Hz, and band-limited noise. We augment a 125 m water column on top of these models to limit the near source imaging artifacts. We train T_w according to the objective function in Equation 1 with the Adam optimization algorithm [22]. In order to evaluate the effectiveness of our inversion scheme when applied to out-of-distribution data, we select a 2D section from the deeper portions of the Parihaka dataset (see Figure 1a). As compared to training images, this image includes more noncontinuous reflectors, due to low signal-to-noise ratio in the deeper parts of the Parihaka dataset. We simulated linearized data with the same acquisition geometry described above to obtain a low-fidelity image (Figure 1b). We solve optimization problem 3 for 5 passes over the shot records, i.e., approximately the same cost as 5 reverse-time migrations. Figure 1c shows the initial guess of the optimization, i.e., $T_{w_2}^{-1}(T_{w_1}(\delta\mathbf{m}_{\text{RTM}}), \mathbf{z}_0)$. We can see that this image is not correctly recovering the reflectors, however, it can be considered a better starting guess than the reverse-time migrated image (Figure 1b) due to corrected amplitudes. Finally, Figure 1d shows the MAP estimate, obtained via solving the optimization problem 3, which successfully reconstructs most of the reflectors with limited artifacts.

Our example uses [JUDI](#) [23] to construct wave-equation solvers, which utilizes [Devito](#) [24, 25] as a highly optimized finite difference solver under the hood. The network architectures are implemented using [InvertibleNetworks.jl](#) [26], a memory-efficient framework for training invertible nets in Julia. Sample code to reproduce the results is provided on [GitHub](#).

5 Conclusions

Considering the Earth’s strong heterogeneity, designing regularization schemes to incorporate prior knowledge for solving ill-posed geophysical inverse problems is challenging. To address this challenge, we proposed a regularization scheme that takes advantage of existing data in the form of low- and high-fidelity seismic images to train a conditional normalizing flow (NF). This conditional NF approximates the imaging posterior distribution for previously unseen data. In order to minimize the impact of data distribution shifts during inference, we reparameterized the unknown image with the conditional NF and inverted for a Gaussian latent variable that fits the data. The resulting




Figure 1: Imaging with conditional NF priors. (a) High-fidelity image. (b) Reverse-time migrated image. (c) Initial guess, $T_{w_2}^{-1}(T_{w_1}(\delta\mathbf{m}_{\text{RTM}}), \mathbf{z}_0)$. (d) MAP estimation with conditional NF prior (Equation 3).

maximum a posteriori estimate takes advantage of the data-driven conditional prior while remaining bound to data and physics. Using numerical experiments, we demonstrated that this approach yields seismic images with limited imaging artifacts in the absence of high-fidelity training data. Further research on quantifying the uncertainty through this regularization technique is required.

6 Acknowledgment

This research was carried out with the support of Georgia Research Alliance and partners of the ML4Seismic Center.

References

- [1] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005. ISBN 978-0-89871-572-9. doi: 10.1137/1.9780898717921.
- [2] Alberto Malinverno and Victoria A Briggs. Expanded uncertainty quantification in inverse problems: Hierarchical bayes and empirical bayes. *GEOPHYSICS*, 69(4):1005–1016, 2004. doi: 10.1190/1.1778243.
- [3] James Martin, Lucas C. Wilcox, Carsten Burstedde, and OMAR Ghattas. A Stochastic Newton MCMC Method for Large-scale Statistical Inverse Problems with Application to Seismic Inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012. URL <http://epubs.siam.org/doi/abs/10.1137/110845598>.
- [4] Felix J. Herrmann and Xiang Li. Efficient least-squares imaging with sparsity promotion and compressive sensing. *Geophysical Prospecting*, 60(4):696–712, 07 2012. doi: 10.1111/j.1365-2478.2011.01041.x. URL <https://slim.gatech.edu/Publications/Public/Journals/GeophysicalProspecting/2012/herrmann11GPelsqIm/herrmann11GPelsqIm.pdf>.

- [5] Rafael Orozco, Ali Siahkoohi, Gabrio Rizzuti, Tristan van Leeuwen, and Felix Johan Herrmann. Photoacoustic imaging with conditional priors from normalizing flows. In *NeurIPS 2021 Workshop on Deep Learning and Inverse Problems*, 2021. URL <https://openreview.net/forum?id=woi1OTvROO1>.
- [6] Jakob Kruse, Gianluca Detommaso, Robert Scheichl, and Ullrich Köthe. HINT: Hierarchical Invertible Neural Transport for Density Estimation and Bayesian Inference. *Proceedings of AAAI-2021*, 2021. URL <https://arxiv.org/pdf/1905.10687.pdf>.
- [7] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An Introduction to Variational Methods for Graphical Models. *Machine Learning*, 37(2):183–233, 1999. doi: 10.1023/A:1007665907178.
- [8] Ali Siahkoohi, Gabrio Rizzuti, Mathias Louboutin, Philipp Witte, and Felix J. Herrmann. Preconditioned training of normalizing flows for variational inference in inverse problems. In *3rd Symposium on Advances in Approximate Bayesian Inference*, 1 2021. URL <https://openreview.net/pdf?id=P9m1sMaNQ8T>.
- [9] Ali Siahkoohi and Felix J Herrmann. Learning by example: fast reliability-aware seismic imaging with normalizing flows. In *First International Meeting for Applied Geoscience & Energy*, pages 1580–1585. Society of Exploration Geophysicists, 2021.
- [10] Nikola Kovachki, Ricardo Baptista, Bamdad Hosseini, and Youssef Marzouk. Conditional Sampling With Monotone GANs, 2021.
- [11] Muhammad Asim, Max Daniels, Oscar Leong, Ali Ahmed, and Paul Hand. Invertible generative models for inverse problems: mitigating representation error and dataset bias. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 399–409. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/asim20a.html>.
- [12] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. volume 37 of *Proceedings of Machine Learning Research*, pages 1530–1538. PMLR, 07–09 Jul 2015. URL <http://proceedings.mlr.press/v37/rezende15.html>.
- [13] Youssef Marzouk, Tarek Moselhy, Matthew Parno, and Alessio Spantini. Sampling via measure transport: An introduction. *Handbook of uncertainty quantification*, pages 1–41, 2016.
- [14] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. In *International Conference on Learning Representations, ICLR*, 2016. URL <http://arxiv.org/abs/1605.08803>.
- [15] Felix J. Herrmann, Ali Siahkoohi, and Gabrio Rizzuti. Learned imaging with constraints and uncertainty quantification. In *Neural Information Processing Systems (NeurIPS) 2019 Deep Inverse Workshop*, 12 2019. URL <https://arxiv.org/pdf/1909.06473.pdf>.
- [16] Rajiv Kumar, Maria Kotsi, Ali Siahkoohi, and Alison Malcolm. Enabling uncertainty quantification for seismic data preprocessing using normalizing flows (NF)—An interpolation example. In *First International Meeting for Applied Geoscience & Energy*, pages 1515–1519. Society of Exploration Geophysicists, 2021.
- [17] Konik Kothari, AmirEhsan Khorashadizadeh, Maarten de Hoop, and Ivan Dokmanić. Trumpets: Injective Flows for Inference and Inverse Problems. *arXiv preprint arXiv:2102.10461*, 2021. URL <https://arxiv.org/abs/2102.10461>.
- [18] Ali Siahkoohi, Mathias Louboutin, and Felix J. Herrmann. The importance of transfer learning in seismic modeling and imaging. *GEOPHYSICS*, 84(6):A47–A52, 11 2019. doi: 10.1190/geo2019-0056.1.
- [19] Hongyu Sun and Laurent Demanet. Elastic full-waveform inversion with extrapolated low-frequency data. In *SEG Technical Program Expanded Abstracts 2020*, pages 855–859. Society of Exploration Geophysicists, 2020. doi: 10.1190/segam2020-3428087.1.

- [20] Ali Siahkoohi, Gabrio Rizzuti, and Felix J Herrmann. Deep bayesian inference for seismic imaging with tasks. *arXiv preprint arXiv:2110.04825*, 2021.
- [21] WesternGeco. Parihaka 3D PSTM Final Processing Report. Technical Report New Zealand Petroleum Report 4582, New Zealand Petroleum & Minerals, Wellington, 2012.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. URL <https://arxiv.org/pdf/1412.6980.pdf>.
- [23] Philipp A. Witte, Mathias Louboutin, Navjot Kukreja, Fabio Luporini, Michael Lange, Gerard J. Gorman, and Felix J. Herrmann. A large-scale framework for symbolic implementations of seismic inversion algorithms in julia. *GEOPHYSICS*, 84(3):F57–F71, 2019. doi: 10.1190/geo2018-0174.1. URL <https://doi.org/10.1190/geo2018-0174.1>.
- [24] F. Luporini, M. Lange, M. Louboutin, N. Kukreja, J. Hüchelheim, C. Yount, P. Witte, P. H. J. Kelly, F. J. Herrmann, and G. J. Gorman. Architecture and performance of devito, a system for automated stencil computation. *CoRR*, abs/1807.03032, jul 2018. URL <http://arxiv.org/abs/1807.03032>.
- [25] M. Louboutin, M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, P. Velesko, and G. J. Gorman. Devito (v3.1.0): an embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development*, 12(3):1165–1187, 2019. doi: 10.5194/gmd-12-1165-2019. URL <https://www.geosci-model-dev.net/12/1165/2019/>.
- [26] Philipp Witte, Gabrio Rizzuti, Mathias Louboutin, Ali Siahkoohi, Felix Herrmann, and Bas Peters. InvertibleNetworks.jl: A Julia framework for invertible neural networks, March 2021. URL <https://doi.org/10.5281/zenodo.4610118>.