

Power-scaled Bayesian inference with score-based generative models

Huseyin Tuna Erdinc, Yunlin Zeng, Abhinav Prakash Gahlot, Felix J. Herrmann, Georgia Institute of Technology

SUMMARY

We propose a score-based generative algorithm for sampling from power-scaled priors and likelihoods within the Bayesian inference framework. Our algorithm enables flexible control over prior–likelihood influence without requiring retraining for different power scaling configurations. Specifically, we focus on generative models to synthesize seismic velocity models conditioned on imaged seismic. Our method enables sensitivity analysis by sampling from intermediate power posteriors, allowing us to assess the relative influence of the prior and likelihood on samples of the posterior distribution. Through a comprehensive set of experiments, we evaluate the effects of varying the power parameter in different settings: applying it solely to the prior, to the likelihood of a Bayesian formulation, and to both simultaneously. The results show that increasing the power of the likelihood up to a certain threshold improves the fidelity of posterior samples to the conditioning data (e.g., RTM), while decreasing the prior power promotes greater structural diversity among samples. Moreover, we find that moderate scaling of the likelihood leads to a reduced shot data residual, confirming its utility in posterior refinement.

INTRODUCTION

Subsurface velocity-model generation forms a critical component of hydrocarbon exploration (Zhang, Sengupta, Parekh, & Elliott-Lockhart, 2022), subsurface monitoring (Gahlot, Orozco, Yin, & Herrmann, 2024) and numerous other geophysical applications (Wagner & Uhlemann, 2021). Typically, subsurface characterization is achieved by analyzing the Earth’s response to physical stimuli, such as electrodynamics, gravity and acoustic wave propagation, to variations in subsurface properties. The resulting tomographic measurements are processed into images for downstream interpretation. In this work, we focus specifically on modeling acoustic properties by probing the Earth’s interior using acoustic waves. However, the proposed methodology is not limited to this particular application and can be extended to a broad class of inverse problems.

Among the various inversion techniques, Full-Waveform Inversion (FWI) has emerged as a leading method due to its ability to resolve high-resolution acoustic models in complex geological settings (Tarantola, 1984). Despite its strengths, FWI suffers from several practical limitations: it is computationally intensive, prone to convergence to local minima, and sensitive to initial models due to the problem’s inherently nonlinear and ill-posed nature (Virieux & Operto, 2009). Additionally, it requires solutions of wave-equation-based partial differential equations, which significantly increases computational cost.

To alleviate these challenges, recent research has explored generative models that leverage physics-informed summary

statistics—such as common-image gathers (CIGs) (Geng et al., 2021; Orozco, Erdinc, Zeng, Louboutin, & Herrmann, 2024; Yin, Orozco, & Herrmann, 2025; Yin, Orozco, Louboutin, & Herrmann, 2024) or reverse-time migration (RTM) images (Muller et al., 2023; Wang, Huang, & Alkhalifah, 2024)—to guide the inversion process. While promising, a common criticism of these generative approaches is that they often rely heavily on strong, structured priors, which can dominate the inference and limit the model’s responsiveness to observed data. In contrast, methods such as Zhao & Curtis (2024b) employ highly non-informative priors, placing more emphasis on the observed data. However, this comes at the cost of reduced regularization, which can affect robustness—especially in underdetermined or noisy regimes—and may increase susceptibility to local minima during inference.

Building on these insights, we propose a novel framework that combines classical Bayesian inference with score-based generative models trained on geologically structure-consistent priors. We introduce a modification to the score-based sampling process that enables sampling from the power-scaled versions of the prior and likelihood, allowing explicit control over their relative influence during the inference process. We validate our approach using RTM images as physics-based summary statistics (Deans, 2002; Orozco, Siahkoobi, Louboutin, & Herrmann, 2024) produced using a smoothed background model (Yin et al., 2024) and demonstrate its performance in generating diverse and data-consistent subsurface velocity models. It is important to note that, although the background models of RTMs are smoothed, they are not kinematically incorrect, so the RTM images retain essential information from the shot data. Unlike prior work that directly uses shot records for inference (Zhao & Curtis, 2024a, 2024b), our method achieves similar information usage through RTM representations.

THEORETICAL EXPLANATION

Seismic imaging and Bayesian inference

Estimation of the unknown subsurface property such as the acoustic wavespeed, \mathbf{x} , requires solving an inverse problem using observed data \mathbf{y} . In our context, we can define our forward problem as:

$$\mathbf{y} = \mathcal{F}(\mathbf{x}) + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim p(\boldsymbol{\varepsilon})$$

where \mathcal{F} represents the nonlinear forward operator and $\boldsymbol{\varepsilon}$ is bandlimited noise. The main complexity of inverting this problem stems from the nontrivial null-space of the forward operator and the compounding effect of the noise (Tarantola, 1984). As a result, multiple velocity models can explain the observed data equally well, which necessitates the use of a Bayesian framework to properly quantify uncertainties. Bayesian inference provides a probabilistic formulation for inverse problems

by computing the posterior probability density function (pdf) using Bayes' rule:

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})},$$

where $p(\mathbf{x})$ is the prior that describes available information about the velocity \mathbf{x} before the inference process, and $p(\mathbf{y}|\mathbf{x})$ is the likelihood function, which given any model value \mathbf{x} calculates the probability of observing imaged data denoted, with a slight abuse of notation, also by \mathbf{y} . The likelihood is used to describe how well \mathbf{y} matches synthetic image generated by a particular model \mathbf{x} . The denominator, $p(\mathbf{y})$, is the evidence or marginal likelihood, serving as a normalization constant to ensure that the posterior is a valid probability distribution.

Simulation-based inference via conditional score based networks

Simulation-based inference (SBI) is a framework that allows the training of surrogates for posterior pdf using neural estimators (Cranmer, Brehmer, & Louppe, 2020). The key idea is to use numerical simulators to generate training pairs $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, where each pair consists of a set of subsurface properties \mathbf{x}_i and the corresponding simulated observation \mathbf{y}_i derived using the forward simulation. In this study, rather than working directly with raw seismic data, we extract RTM images as summary statistics, which are used as \mathbf{y} during both training and inference. This approach retains key structural information while reducing the dimensionality and complexity of the data. The resulting pairs $(\mathbf{x}_i, \mathbf{y}_i)$ are then used to train a conditional generative network, which learns the posterior distribution of the velocities conditioned on RTM images. In this study, we will use conditional score-based generative models in a SBI setting.

Score-based models are density estimators based on learning of the annealed (time-dependent) score of the target distribution $\nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma(t))$, $\sigma(t)$ represents a noise schedule, by time reversing the gradual addition of noise to the target distribution, which corresponds to diffusion. After the completion of the score learning process, score-based models can sample from target distributions using the reverse process. The simplified score learning objective (when $\sigma(t) = t$) can be summarized as follows (Karras, Aittala, Aila, & Laine, 2022):

$$\hat{\theta} = \operatorname{argmin}_{\theta} E_{\mathbf{x} \sim p(\mathbf{x})} E_{\mathbf{n} \sim \mathcal{N}(0, t^2 \mathbf{I})} \|s_{\theta}(\mathbf{x} + \mathbf{n}; t) - \mathbf{x}\|_2^2,$$

where the score function at a specific time t is approximated by the trained network $s_{\theta}(\mathbf{x}; t)$. After completion of the training procedure, target distribution samples can be obtained by solving the following stochastic differential equation (SDE):

$$d\mathbf{x} = -(\dot{\sigma}(t) + \beta(t)\sigma(t))\sigma(t)\nabla_{\mathbf{x}} \log p(\mathbf{x}; \sigma(t)) dt + \sqrt{2\beta(t)\sigma(t)} d\omega_t,$$

by following existing strategies in (Karras et al., 2022). In this expression, ω_t is the reverse standard Wiener process and $\beta(t)$ is a function that describes the amount of stochastic noise during the sampling process. If $\beta(t) = 0$ for all t , then the process becomes a probabilistic ordinary differential equation (ODE), otherwise it represents time-varying Langevin diffusion SDE.

The previous discussion is solely based on the unconditional case and it can easily be extended to conditional distributions when we aim to find the conditional score $\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y}; \sigma(t))$. Then the new objective for the network can be written as follows:

$$\hat{\theta} = \operatorname{argmin}_{\theta} E_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{x})} E_{\mathbf{x} \sim p(\mathbf{x})} E_{\mathbf{n} \sim \mathcal{N}(0, t^2 \mathbf{I})} \|s_{\theta}(\mathbf{x} + \mathbf{n}, \mathbf{y}, t) - \mathbf{x}\|_2^2.$$

Power scaling in Bayesian inference

Previous works on Bayesian inference for seismic inversion often exhibit a trade-off—either overly relying on strong priors, which may decrease the influence of observed data, or placing too much emphasis on data, potentially at the cost of regularization. This motivates the need for a principled mechanism to adjust the relative influence of the prior and likelihood without incurring significant computational costs.

To address this, we introduce power scaling as a flexible tool to modulate this balance. Power scaling is a controlled, distribution-agnostic technique for modifying probability distributions by adjusting the relative influence of the likelihood and prior (Kallioinen, Paananen, Bürkner, & Vehtari, 2024). Intuitively, it can be understood as a mechanism to either amplify or attenuate the effect of each component in the posterior, without requiring any specific parametric assumptions.

The power-scaled posterior is defined as:

$$p^{\lambda, \alpha}(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})^{\lambda} p(\mathbf{x})^{\alpha}}{\int p(\mathbf{y}|\mathbf{x})^{\lambda} p(\mathbf{x})^{\alpha} d\mathbf{x}}.$$

Here, λ is the power scaling factor applied to the likelihood and α is the scaling factor for the prior.

From a statistical perspective, power-scaling the likelihood (i.e., increasing λ) mimics the effect of having more observations, thereby sharpening the posterior around high-likelihood regions. Conversely, reducing λ decreases the weight of the data and leads to more diffuse posteriors. Similarly, scaling the prior with exponent α adjusts its influence. Larger values of α (>1) concentrate the prior around its high-density regions, reinforcing prior assumptions. Smaller α values relax the prior, allowing the posterior to be more exploratory.

In this study, we propose to estimate power-scaled posterior using score-based generative modeling. The score of the power-scaled posterior can be written as:

$$\nabla_{\mathbf{x}} p^{\lambda, \alpha}(\mathbf{x}|\mathbf{y}) = \lambda \nabla_{\mathbf{x}} p(\mathbf{y}|\mathbf{x}) + \alpha \nabla_{\mathbf{x}} p(\mathbf{x}).$$

A primary challenge in this formulation lies in estimating the gradient of the likelihood term, which is often computationally intensive—especially in high-dimensional and physics-based models such as seismic inversion. To address this, we leverage a useful identity from Bayes' rule that expresses the likelihood score in terms of the posterior and prior scores: $\nabla_{\mathbf{x}} p(\mathbf{y}|\mathbf{x}) = \nabla_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) - \nabla_{\mathbf{x}} p(\mathbf{x})$. Substituting this into the power-scaled score expression yields the following formulation:

$$\nabla_{\mathbf{x}} p^{\lambda, \alpha}(\mathbf{x}|\mathbf{y}) = \lambda \nabla_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) + (\alpha - \lambda) \nabla_{\mathbf{x}} p(\mathbf{x}).$$

This expression corresponds to training surrogates for both the posterior score $\nabla_{\mathbf{x}}p(\mathbf{x}|\mathbf{y})$ and prior score $\nabla_{\mathbf{x}}p(\mathbf{x})$. Leveraging the approach introduced in classifier-free guidance (CFG) (Ho & Salimans, 2022), we can estimate both scores using a single conditional network by randomly masking the conditioning input during training. This eliminates the need to train two separate networks. Once the score function is learned, we can combine the posterior and prior scores at inference time using their respective power coefficients λ and α . To generate samples from the power-scaled posterior, we incorporate a Langevin-based correction step, such as predictor-corrector methods, into the sampling loop to improve accuracy and stability (Bradley & Nakkiran, 2024; Du et al., 2023).

The primary advantage of our approach is its flexibility: only a single amortized score network needs to be trained, yet it enables sampling from a wide range of power-scaled posterior distributions by adjusting power coefficients at inference time. This makes it ideal for sensitivity analysis, robustness testing or exploring different prior-data tradeoffs without retraining. A key limitation, however, is sampling efficiency. Unlike standard posterior sampling, our approach requires Langevin-type correction steps at each iteration, which can increase computational cost—particularly when a large number of steps are needed to ensure convergence. That said, the method still avoids the need for expensive likelihood evaluations and in practice, we achieve efficient inference with a per-sample generation time of approximately 4 seconds.

NUMERICAL CASE STUDY

To evaluate the proposed methodology, we conduct a numerical study using a synthetic 3D Earth model derived from the Compass model, which is representative of geological formations in the North Sea region (E. Jones, A. Edgar, I. Selvage, & Crook, 2012). The training dataset is constructed by slicing 2D velocity models from the 3D synthetic volume and pairing them with corresponding reverse-time migration (RTM) images. The dataset consists of 800 training samples, each defined on a 256×512 grid with a spatial resolution of 12.5 meters, covering an area of $3.2\text{km} \times 6.4\text{km}$. Seismic data are simulated using 16 sources and 256 receivers, with Ricker wavelets centered at 20Hz dominant frequency. To simulate realistic acquisition conditions, 10dB colored Gaussian noise is added to the shot records. RTM images are generated using a Gaussian-smoothed 2D background velocity model. Both wave simulation and migration are implemented using the open-source seismic inversion package JUDI (Louboutin et al., 2023).

To train the conditional score-based network for simultaneous posterior and prior estimation, we adopt a classifier-free guidance (CFG) strategy: the conditioning RTM input is randomly dropped with a probability of 0.2 and replaced with Gaussian noise during training (Ho & Salimans, 2022). For conditioning, RTM images are concatenated with the noisy velocity samples, and the network, is trained to denoise and recover the clean velocity samples. The model is trained for 12 GPU hours. For testing, we select an unseen RTM example and perform posterior sampling using the trained network.

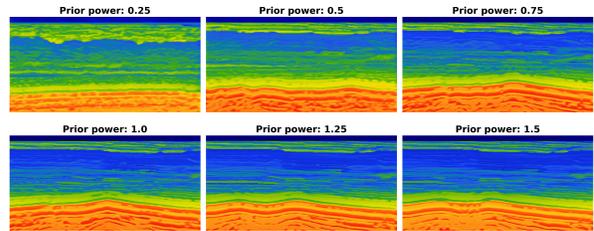


Figure 1: Prior samples generated with increasing power α from 0.25 to 1.5. As α increases, the generated velocity models become progressively sharper and more structurally consistent, especially in deeper layers. Lower powers result in higher variability and less coherent geologic features.

Prior scaling

In the first experiment, we investigate the effect of power scaling on the learned prior distribution in isolation. Specifically, we set the likelihood term aside and generate samples solely from the prior by varying the coefficient in the power-scaled prior $p(\mathbf{x})^\alpha$. This allows us to analyze how different levels of prior strength affect the structure and diversity of generated velocity models. Figure 1 shows the evolution of prior samples as α increases from 0.25 to 1.5. All samples are generated from the same random seed to ensure consistency in the comparison. At low prior powers (e.g., $\alpha = 0.25$), the generated samples exhibit high variability, particularly in the deeper regions. Layer boundaries are less coherent and geological structures appear diffuse and discontinuous. As α increases, the samples become increasingly regularized: layers become more consistent and sharper. This reflects the fact that the learning process concentrates probability mass around high-density regions of the prior, which tend to correspond to well-structured geological patterns seen in training data.

Likelihood scaling in Bayesian inference

In this experiment, we fix the prior power at its standard Bayesian setting, $\alpha = 1$ and vary the likelihood power λ from 0 to 16 for a specific RTM image. This allows us to investigate how scaling the influence of the likelihood affects the resulting posterior samples. Figure 2 shows velocity model samples generated with increasing values of λ , alongside the ground-truth velocity model and the RTM image used as the conditioning input. At low values of λ (e.g., $\lambda \leq 0.4$), the contribution of the likelihood is minimal, and the generated velocity models closely resemble unconditional prior samples. The influence of the RTM conditioning is barely noticeable. As λ increases, the posterior increasingly conforms to features present in the RTM image. For instance, at $\lambda = 1.0$, corresponding to standard Bayesian inference, the generated model shows moderate data fidelity. Interestingly, we observe that data fidelity continues to improve up to a certain point—most notably at $\lambda = 2.0$ —after which performance begins to degrade. This behavior is quantified in figure 3, which presents the data residual (i.e., difference between simulated shot data from the generated velocity models and ground-truth shot data). The residual is minimized at $\lambda = 2.0$, indicating optimal alignment with the ground truth at this power level. This result highlights

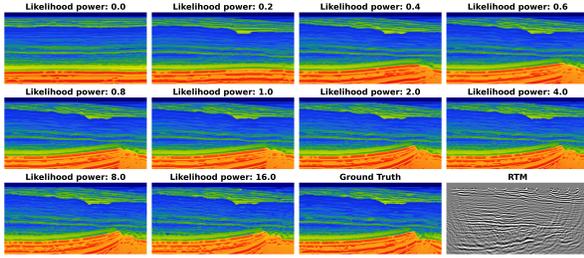


Figure 2: Posterior samples generated with varying likelihood power λ from 0.0 to 16.0, with fixed prior power $\alpha = 1$. As λ increases, the samples incorporate more structure from the conditioning RTM image. Maximum alignment with the ground truth occurs around $\lambda = 2.0$, beyond which overfitting and degradation in performance become apparent.

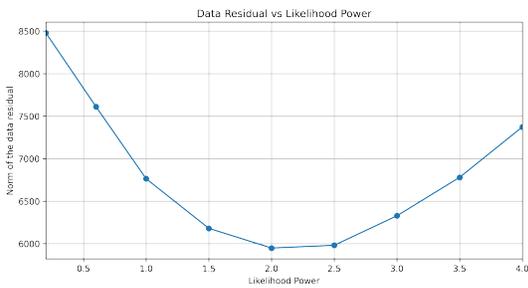


Figure 3: The ℓ_2 -norm of data residual as a function of likelihood power λ , with fixed prior power $\alpha = 1$. Residual decreases as λ increases, reaching a minimum around $\lambda = 2.0$, indicating optimal alignment with the ground truth. Beyond this point, performance begins to degrade, suggesting that excessive amplification of the likelihood leads to overfitting or reduced generalization.

a key insight: maximum posterior fidelity does not necessarily occur at the classical Bayesian setting $\lambda = 1.0$, but rather at an upweighted likelihood power. This demonstrates the practical value of power-scaling as a tool for tuning the trade-off between prior regularization and data conformity.

Power-scaling compass

Having independently explored the effects of prior and likelihood power scaling, we now investigate their joint influence by constructing a power-scaling compass. In this setup, we simultaneously vary the prior power α and the likelihood power λ to study their combined impact on posterior samples. Figure 4 presents posterior samples generated under various combinations of $\alpha \in \{0.5, 1.0, 2.0\}$ and $\lambda \in \{0.5, 1.0, 2.0\}$. This grid-like layout allows us to visually interpret how the balance between prior and likelihood power influences behavior of the samples.

We observe that when likelihood power exceeds prior power, the generated samples exhibit greater alignment with the conditioning RTM image—demonstrating stronger adherence to observed data. Conversely, when prior power dominates, the samples exhibit smoother layer transitions and more geologically consistent patterns, reflecting stronger regularization from

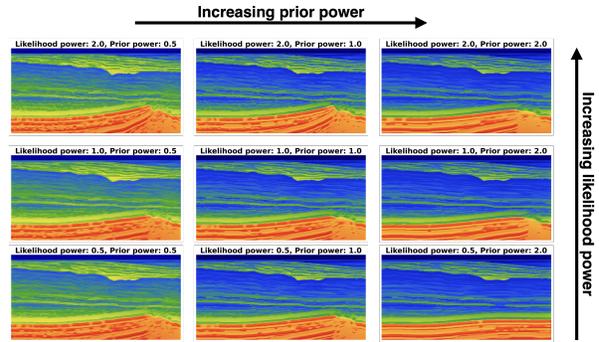


Figure 4: Posterior samples generated by jointly varying likelihood power λ and prior power α . Columns correspond to increasing prior power (left to right), while rows correspond to increasing likelihood power (top to bottom). Samples with larger λ better align with RTM conditioning data, while larger α values enforce structural consistency and sharper geological layering. Lower prior powers increase generative diversity and loosen geological constraints.

the prior distribution. Notably, reducing the prior power relaxes the structural constraints in the model and leads to increased variability and the emergence of less structured or more exploratory features. This visualization highlights how power-scaling enables fine-grained control over the trade-off between data fidelity and prior-driven structural regularity, offering a versatile tool for interpretability in seismic inversion tasks.

CONCLUSIONS

In this study, we proposed a method for applying power scaling to seismic velocity model generation within a Bayesian inference framework. Our approach is built on a single amortized score-based generative model trained using a classifier-free guidance strategy, enabling simultaneous generation of both prior and posterior samples. Crucially, this design allows flexible control over the relative influence of the prior and likelihood without requiring expensive likelihood evaluations, as all modifications are made at the inference stage. Experimental results demonstrate that power scaling offers a principled way to relax or constrain the prior, making it a valuable tool for generating diverse samples during training or for performing inference under less influential priors. Moreover, increasing the likelihood power was shown to improve data fidelity, enabling better alignment between posterior samples and observed seismic features. In future work, we aim to conduct a more rigorous investigation into the role of power scaling in uncertainty quantification and its impact on interpretability and robustness in seismic inversion tasks.

ACKNOWLEDGEMENT

This research was carried out with the support of Georgia Research Alliance, partners of the ML4Seismic Center. The overall readability is enhanced with ChatGPT.

References

10

Bradley, A., & Nakkiran, P. (2024). Classifier-free guidance is a predictor-corrector. *arXiv Preprint arXiv:2408.09000*.

Cranmer, K., Brehmer, J., & Louppe, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48), 30055–30062. doi:<https://doi.org/10.1073/pnas.2017141117>

Deans, M. C. (2002). Maximally informative statistics for localization and mapping. In *Proceedings 2002 IEEE international conference on robotics and automation (cat. No. 02CH37292)* (Vol. 2, pp. 1824–1829). IEEE.

Du, Y., Durkan, C., Strudel, R., Tenenbaum, J. B., Dieleman, S., Fergus, R., . . . Grathwohl, W. S. (2023). Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc. In *International conference on machine learning* (pp. 8489–8510). PMLR.

E. Jones, C., A. Edgar, J., I. Selvage, J., & Crook, H. (2012). Building complex synthetic models to evaluate acquisition geometries and velocity inversion technologies. In *74th EAGE Conference and Exhibition Incorporating EUROPEC 2012*, cp-293. doi:<https://doi.org/10.3997/2214-4609.20148575>

Gahlot, A. P., Orozco, R., Yin, Z., & Herrmann, F. J. (2024). An uncertainty-aware digital shadow for underground multi-modal CO₂ storage monitoring. *arXiv Preprint arXiv:2410.01218*.

Geng, Z., Zhao, Z., Shi, Y., Wu, X., Fomel, S., & Sen, M. (2021). Deep learning for velocity model building with common-image gather volumes. *Geophysical Journal International*, 228(2), 1054–1070. doi:<https://doi.org/10.1093/gji/ggab385>

Ho, J., & Salimans, T. (2022). Classifier-free diffusion guidance. *arXiv Preprint arXiv:2207.12598*.

Kallioinen, N., Paananen, T., Bürkner, P.-C., & Vehtari, A. (2024). Detecting and diagnosing prior and likelihood sensitivity with power-scaling. *Statistics and Computing*, 34(1), 57.

Karras, T., Aittala, M., Aila, T., & Laine, S. (2022). Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35, 26565–26577.

Louboutin, M., Witte, P., Yin, Z., Modzelewski, H., Kerim, Costa, C. da, & Nogueira, P. (2023, March). Slimgroup/JUDI.jl: v3.2.3. Zenodo. doi:<https://doi.org/10.5281/zenodo.7785440>

Muller, A. P., Bom, C. R., Costa, J. C., Klatt, M., Faria, E. L., Silva, B. dos S., . . . Albuquerque, M. P. de. (2023). Deep-tomography: Iterative velocity model building with deep learning. *Geophysical Journal International*, 232(2), 975–989.

Orozco, R., Erdinc, H. T., Zeng, Y., Louboutin, M., & Herrmann, F. J. (2024). Machine learning-enabled velocity model building with uncertainty quantification. Retrieved from <https://arxiv.org/abs/2411.06651>

Orozco, R., Siahkoohi, A., Louboutin, M., & Herrmann, F. (2024). ASPIRE: Iterative amortized posterior inference for

bayesian inverse problems. *Inverse Problems*.

Tarantola, A. (1984). Inversion of seismic reflection data in the acoustic approximation. *Geophysics*, 49(8), 1259–1266.

Virieux, J., & Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6), WCC1–WCC26.

Wagner, F. D., & Herrmann, F. J. (2021). A review of multi-method imaging approaches in environmental geophysics. *Advances in Geophysics*, 62, 1–72.

Wang, F., Huang, X., & Alkhalifah, T. (2024). Controllable seismic velocity synthesis using generative diffusion models. *Journal of Geophysical Research: Machine Learning and Computation*, 1(3), e2024JH000153.

Yin, Z., Orozco, R., & Herrmann, F. J. (2025). WISER: Multi-modal variational inference for full-waveform inversion without dimensionality reduction. *Geophysics*, 90(2), A1–A7. doi:<https://doi.org/10.1191/G1-2024-0483.1>

Yin, Z., Orozco, R., Louboutin, M., & Herrmann, F. J. (2024). WISE: Full-waveform variational inference via subsurface extensions. *Geophysics*, 89(4), A23–A28.

Zhang, L., Sengupta, S., Parekh, C., & Elliott-Lockhart, R. (2022). High-frequency FWI velocity model building for hydrocarbon delineation. In *SEG international exposition and annual meeting* (p. D011S043R004). SEG.

Zhao, X., & Curtis, A. (2024a). Physically structured variational inference for bayesian full waveform inversion. *Journal of Geophysical Research: Solid Earth*, 129(11), e2024JB029557.

Zhao, X., & Curtis, A. (2024b). Variational prior replacement in bayesian inference and inversion. *Geophysical Journal International*, 239(2), 1236–1256. doi:<https://doi.org/10.1093/gji/ggae334>