

Rock-Physics-Informed Diffusion Model for Direct Joint Seismic Inversion

Introduction

Predicting the spatial distribution of reservoir rock properties from seismic data is essential for quantitative reservoir characterization since facies, porosity, shale volume and fluid saturation directly control reservoir quality and flow behaviour. Sequential seismic reservoir characterization workflows usually recover these properties indirectly by first inverting seismic data for elastic attributes and then translating them into petrophysical variables through rock-physics models (Grana and Della Rossa, 2010; Azevedo et al., 2019). However, seismic inversion is non-unique and ill posed, and the sequential use of rock-physics interpretation can propagate elastic-inversion errors into the final rock-property model. This effect motivates direct seismic inversion of rock properties, where rock-physics relations are embedded in the inversion process rather than applied only as a post-inversion transform.

Deep learning methods, especially deep generative models, have recently shown strong ability to synthesize high-dimensional images and spatial models (Ho et al., 2020; Dhariwal and Nichol, 2021). These models have increasingly been explored for seismic inversion, including data-driven impedance inversion, velocity-model reconstruction and generative-prior-based waveform inversion (Alfarraj and AlRegib, 2019; Mosser et al., 2020). However, purely data-driven predictions may lack explicit physical consistency and can produce geologically implausible artefacts when they are not constrained by seismic forward modelling or rock-physics relations (Li et al., 2020). Here, we propose a rock-physics-informed diffusion model (RP-IDM) for simultaneous inversion of facies and rock properties from post-stack seismic reflection data.

Rather than first estimating elastic properties, RP-IDM directly samples facies, porosity, shale volume and water saturation. These variables are mapped to elastic parameters by a differentiable rock-physics model and then to reflection responses by seismic forward modelling. During diffusion posterior sampling (DPS; Chung et al., 2024), seismic and well-log misfits provide gradients that guide the denoising trajectory. The posterior ensemble therefore remains in the rock-property domain while being constrained by the learned geological prior, rock-physics consistency, seismic response and local well-log information.

Methodology

RP-IDM consists of three coupled components: a differentiable rock-physics model, rock-physics-embedded diffusion training, and physics-guided posterior diffusion sampling. The first component defines the physical mapping from rock-property variables to elastic properties (Figure 1). Specifically, porosity, shale volume and water saturation are transformed into elastic properties, including P-wave velocity, S-wave velocity and density, through explicit rock-physics relations such as mineral and fluid mixing, dry-frame effective-medium modelling and Gassmann fluid substitution. The elastic properties are then converted to acoustic impedance, from which normal-incidence reflection coefficients are computed and convolved with a seismic wavelet to generate synthetic post-stack seismic data. Because the entire chain is differentiable, the seismic misfit can provide gradients with respect to the rock-property variables, allowing RP-IDM to update reservoir properties directly rather than through a separate elastic-inversion step.




Figure 1 Differentiable rock-physics mapping from rock to elastic properties.

The second component trains the diffusion network in the rock-property domain (Figure 2). Multi-channel reservoir models, containing facies, porosity, shale volume and water saturation, are progressively perturbed by Gaussian noise, and the network learns the reverse denoising direction

through the standard denoising loss. In RP-IDM, this training objective is augmented by a rock-physics-based seismic loss: the denoised rock-property estimate is passed through the differentiable rock-physics model and seismic forward model to synthesize post-stack seismic data, which are compared with the corresponding target seismic response. The denoising and physical losses are back-propagated jointly, so the network learns not only the joint spatial prior of reservoir variables but also their consistency with the seismic response. As this prior is learned directly in the rock-property domain, the subsequent inversion can sample reservoir variables rather than intermediate elastic properties.




Figure 2 Schematic representation of rock-physics-embedded diffusion training in RP-IDM.




Figure 3 Representative joint inversion of facies, rock properties and seismic response.

In the inversion stage, RP-IDM combines the learned diffusion prior with seismic and well-log constraints through posterior diffusion sampling. At each reverse step, the current rock-property sample is passed through the differentiable rock-physics and seismic forward models, and the synthetic post-stack seismic trace is compared with the observed trace. Sparse well logs provide local conditioning constraints that encourage the sampled properties to match known absolute values at well locations, while the seismic misfit gradient guides the trajectory toward reduced seismic error. The final posterior samples therefore remain consistent with the learned geological prior, surface seismic response and local well information. Repeating the stochastic sampling process yields an ensemble for uncertainty assessment.




Figure 4 Well-log comparison of posterior realizations at conditioning and blind wells.

Results and Discussion

Synthetic application examples show that RP-IDM predicts spatially coherent joint inversions of facies and rock properties. Figure 3 compares the true model and observed seismic data with a single inversion result and the posterior mean from 50 independent realizations. The single realization preserves local geological variability, whereas the posterior mean emphasizes structures that are stable across the

ensemble. The recovered facies architecture and broad porosity, shale volume and saturation patterns agree with the true model, and the synthetic seismic response reproduces the main reflection events and amplitude trends of the observed post-stack seismic data. This demonstrates that the inversion is constrained by both rock-property structure and seismic forward-response consistency.

Quantitatively, RP-IDM achieved balanced performance across facies, rock-property and seismic-response metrics. The held-out validation results show a facies accuracy of 0.899, with RMSE values of 0.060, 0.225 and 0.208 for porosity, shale volume and water saturation, respectively. The seismic response was closely matched, with a normalized mean-squared seismic error of 0.0086 and an NCC loss of 0.0019. These results support the qualitative observation that RP-IDM can recover reservoir-property structure while maintaining seismic consistency.

Figure 4 compares posterior well-log realizations at three conditioning wells and one blind well. For facies, porosity, shale volume and water saturation, the conditioning wells show tightly concentrated posterior curves near the true logs, whereas the blind well exhibits broader uncertainty. This behaviour is consistent with the role of well information as a local conditioning constraint: uncertainty is reduced near known wells but remains broader away from direct control. The posterior ensemble therefore provides both a best-estimate inversion and a practical measure of spatially varying uncertainty. Results on real case applications are aligned with those obtained in the numerical examples.

Conclusions

RP-IDM enables direct joint inversion of facies and rock properties from post-stack seismic data by combining rock-physics-embedded diffusion training with physics-guided posterior sampling. The method recovered coherent facies, porosity, shale volume and water saturation models, reproduced the main seismic response, and provided ensemble-based uncertainty at both conditioning and blind wells. These results indicate that RP-IDM provides a physically constrained generative framework for joint seismic reservoir inversion in the rock-property domain.

Acknowledgements

The authors acknowledge the China-Portugal Joint Research Institute on AI for Climate and Energy and the National Key R&D Program of China (2024YFE0114000) and the support of Fundação para a Ciência e a Tecnologia via UID/04028/2025 (<https://doi.org/10.54499/UID/04028/2025>).

References

- Alfarraj, M., and AlRegib, G. [2019]. Semi-supervised learning for acoustic impedance inversion. SEG International Exposition and Annual Meeting.
- Azevedo, L., Grana, D., and Amaro, C. [2019]. Geostatistical rock physics AVA inversion. *Geophysical Journal International*, 216(3), 1728-1739.
- Chung, H., Kim, J., McCann, M. T., Klasky, M. L., and Ye, J. C. [2024]. Diffusion Posterior Sampling for General Noisy Inverse Problems (arXiv:2209.14687). arXiv.
- Dhariwal, P., and Nichol, A. [2021]. Diffusion Models Beat GANs on Image Synthesis. *Advances in Neural Information Processing Systems*, 34, 8780-8794.
- Grana, D., and Della Rossa, E. [2010]. Probabilistic petrophysical-properties estimation integrating statistical rock physics with seismic inversion. *Geophysics*, 75(3), O21-O37.
- Ho, J., Jain, A., and Abbeel, P. [2020]. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems*, 33, 6840-6851.
- Li, S., Liu, B., Ren, Y., Chen, Y., Yang, S., Wang, Y., and Jiang, P. [2020]. Deep-Learning Inversion of Seismic Data. *IEEE Transactions on Geoscience and Remote Sensing*, 58(3), 2135-2149.
- Mosser, L., Dubrule, O., and Blunt, M. J. [2020]. Stochastic Seismic Waveform Inversion Using Generative Adversarial Networks as a Geological Prior. *Mathematical Geosciences*, 52(1), 53-79.