



Research Paper

Physics-Guided Neural Inversion of 4D Seismic Data for CO₂ Plume Characterization in Heterogeneous Storage Reservoirs: An Original Three-Dimensional Simulation Study

J. Gopichand

Department of Informatics,

Universitas Muhammadiyah Yogyakarta, Yogyakarta, Indonesia

M. N. L. Sahity

Department of Informatics,

Universitas Muhammadiyah Yogyakarta, Yogyakarta, Indonesia

Abstract

Geological CO₂ storage is complicated due to the presence of reservoir heterogeneity, the limited resolution of the seismic data, the ambiguity of pressure-saturation, and model uncertainty in quantitative interpretation of time-lapse seismic data for geological CO₂ storage. This research designs and implements an innovative three-dimensional simulation benchmark for physics-guided neural characterization of CO₂ plumes. A $24 \times 24 \times 5$ grid was used to generate 100 reservoir realizations that were heterogeneous. The saturation and pressure snapshots were computed by a conservative finite-volume-based transport operator, and the noisy 4D acoustic-impedance changes were then computed using a band-limited rock-physics operator. 60 realizations were employed for training, 15 realizations for validation, 15 high connectivity vertical realizations for independent testing and 10 out-of-distribution for high vertical connectivity realizations. A 4-member multi-layer neural ensemble predicted 16 voxel and neighborhood features to CO₂ saturation and pressure change. This ensemble prediction was then corrected by a physics-constrained step using a physics-projection that implemented saturation bounds, exact mobile-volume conservation and weak smoothing in space. Local rock-physics inversion gave an RMS saturation error of 0.0148 on 60 independent test volumes, an intersection-over-union of 0.778 and a mean mobile-volume error of 24.31%. Neural ensemble reduced saturation error to 0.00574 with an increased plume intersection-over-union 0.897 and 3.26% mobile-volume error. Using physics-guided projection, we obtained a saturation error of 0.00642, a plume intersection-over-union of 0.884, and a mobile-volume error of less than 10^{-120} %. Validation-calibrated nominal 90% interval had a coverage rate of 90.1% on the independent test set, voxel-wise. Physics-guided saturation error rose to 0.00944 and the plume intersection-over-union fell to 0.768, with the broad training range resulting in 92.3% cover. The results demonstrate that the physics guidance can resolve a material inventory inconsistency with only a small loss in pointwise accuracy but that there is still a need to use independent geological stress testing for the purpose of establishing transferability.

Keywords: geological CO₂ storage; 4D seismic monitoring; neural inversion; physics-guided machine learning; plume characterization; uncertainty quantification

1. Introduction

One of the key monitoring challenges in geological storage is verification of the location and inventory of injected CO₂. Subsequent seismic surveys may be able to image changes in elastic properties due to fluid substitution and pressure changes; however, the inverse problem of estimating the saturation and pressure changes from the seismic images is ill-posed. The results of the Sleipner field monitoring showed the importance of integrating both time-lapse seismics and gravimetry to interpret the geometry of the plume and the mass of the free phase over a long time (Furre et al., 2017). Crosswell seismic tomography at Cranfield also identified local CO₂ invasion due to higher-permeability intervals, highlighting uncertainty due to rock-physics assumptions and sub-seismic mixing (Ajo-Franklin et al., 2013). The traditional seismic-reservoir inversion approach is based on rock-physics parameterization, regularization and probabilistic conditioning. Bayesian time-lapse inversion can simultaneously estimate static reservoir properties and dynamic pressure (or saturation) changes, and the computational burden increases rapidly as a function of model dimension and size of the ensemble (Grana & Mukerji, 2015). Rock-physics parameterization directly from the rock can decrease crosstalk when compared with sequential parameterization from elastic and petrophysical models, but it is still dependent on the forward-model assumptions and acquisition (Mardan et al., 2023). There is an alternative way using machine-learning methods. After training, neural networks that were trained using simulated seismic and saturation pairs are able to provide fast monitoring estimates and uncertainty estimates (Um et al., 2024). Geological-carbon-storage data assimilation and forecasting workflows have also been scaled up with deep surrogate models (Tang et al., 2021). Joint inversion of seismic, electrical and reservoir variables via differentiable physics (Liu et al., 2023) and encoding of geologically plausible priors and improved high-dimensional Bayesian inference with autoencoder and generative representation (Liu et al., 2024; Li et al., 2026). These advances cannot remove three practical "risks". First, voxel-level random splitting has the tendency to introduce almost identical spatial structures and time snapshots into both the training set and test set. Second, minimising the pointwise prediction error does not ensure that the reconstructed plume respects the injected inventory known. Third, the failure can be hidden by the performance within the training distribution if the connectivity of the reservoir is not the same as the simulated prior. Simulations of the heterogeneous Evetts depleted gas reservoir conducted on the field scale identified that dispersion of CO₂ over stacked formations and over multiple injection wells resulted in better pressure distribution and reduced plume expansion, and that fault transmissibility was a significant factor on containment (Williams et al., 2026). The present study is motivated by that field-scale complexity, and does not repeat an injection-optimization study, but instead focuses upon the monitoring inverse problem. The goal is to measure the accuracy and physical consistency, uncertainty calibration, and out-of-distribution behavior of a repeatable neural inversion pipeline. The contribution is an executed original simulation experiment using: (1) heterogeneous three-dimensional reservoir realizations, (2) physics guided projection onto saturation and mobile-volume constraints, (3) a high connectivity stress test, (4) a band-limited rock-physics observation model, (5) a conservative synthetic plume transport, and (6) a neural ensemble trained with realization-wise separation. The study does not include thermo-hydro-mechanical coupling, geomechanics, reactive transport, geothermal extraction or the use of the pressure-limited trapping simulator that has been used in a separate investigation.

Reproducible physics-guided 4D-seismic monitoring workflow




Figure 1. Reproducible physics-guided 4D-seismic monitoring workflow.

2. Methodology

2.1 Study design and reproducibility

The workflow was implemented in Python using NumPy, SciPy, and scikit-learn. The fixed study seed was 20260903. All reservoir properties and monitoring observations were synthetic. The reference domain measured $2400 \times 2400 \times 100$ m and was discretized into $24 \times 24 \times 5$ cells. The reduced resolution was selected so the complete experiment could be rerun on a CPU while preserving three-dimensional connectivity, buoyancy bias, multiple monitoring times, and independent geological realizations.

One hundred reservoir realizations were divided before voxel extraction. The training set contained 60 realizations, the validation set 15, the independent test set 15, and the out-of-distribution set 10. Each realization contributed monitoring snapshots at 25%, 50%, 75%, and 100% of the simulated injection period. No realization, spatial pattern, or snapshot from validation, testing, or out-of-distribution evaluation was used to fit network weights.

2.2 Heterogeneous reservoir realizations

A standard normal random field $g(\mathbf{x})$ was smoothed anisotropically and standardized to zero mean and unit variance. Horizontal permeability was assigned in millidarcies as

$$\log_{10} k_h(\mathbf{x}) = 2.0 + 0.58g(\mathbf{x}). \quad (1)$$

Porosity was correlated with permeability while retaining an independent geostatistical component ϵ_ϕ :

$$\phi(\mathbf{x}) = \text{clip}[0.18 + 0.025g(\mathbf{x}) + 0.012\epsilon_\phi(\mathbf{x}), 0.09, 0.30]. \quad (2)$$

Vertical permeability was

$$k_v(\mathbf{x}) = r_v k_h(\mathbf{x}) \exp[0.18g_v(\mathbf{x})], \quad (3)$$

where r_v was sampled uniformly between 0.04 and 0.14 for training, validation, and independent testing. Out-of-distribution cases used r_v between 0.22 and 0.30. This separation tests a physically interpretable change in vertical communication rather than random noise alone.

Baseline acoustic impedance Z_0 was correlated with porosity and a separate smooth perturbation ϵ_Z :

$$Z_0(\mathbf{x}) = 9.2 - 6.0[\phi(\mathbf{x}) - 0.18] + 0.22\epsilon_Z(\mathbf{x}). \quad (4)$$

The impedance values are dimensionless scaled units used consistently by the synthetic observation operator. They should not be interpreted as calibrated field measurements.

Synthetic reservoir and monitoring data, representative layer 5




Figure 2. Synthetic reservoir properties, true dynamic states, and observed 4D impedance change.

Table 1. Reference simulation and learning parameters.

Parameter	Value
Domain	2400 × 2400 × 100 m
Grid	24 × 24 × 5, 2880 cells
Training / validation / test / OOD realizations	60 / 15 / 15 / 10
Monitoring snapshots per realization	4
Training and test k_v/k_h	0.04-0.14
OOD k_v/k_h	0.22-0.30
Injected pore-volume fraction	0.035-0.065
Seismic noise fraction	0.025-0.075
Neural input features	16
Hidden layers	48 and 24 ReLU units
Ensemble members	4
Training voxels	95,000
Maximum epochs	55
Physics-guided iterations	8

2.3 Conservative CO₂ transport and pressure response

The synthetic transport target was based on the phase-occupancy balance

$$\frac{\partial(\phi S_g)}{\partial t} + \nabla \cdot (\mathbf{u}_g - \phi \mathbf{D} \nabla S_g) = q_g, \quad (5)$$

where S_g is mobile CO_2 saturation, $\mathbf{u}_g$ is an effective advective-buoyant flux, $\mathbf{D}$ is numerical dispersion, and q_g is injection. The executable implementation used a conservative cell-to-cell transition operator. For neighboring cells i and j , directional conductance was proportional to harmonic permeability:

$$T_{ij} = \frac{2k_i k_j}{k_i + k_j} b_{ij}, \quad (6)$$

where b_{ij} increased upward transport and reduced downward transport to represent buoyancy segregation. At each substep, a fraction γ of cell inventory $M_i = \phi_i S_{g,i}$ was distributed by normalized conductance:

$$M_i^{n+1} = M_i^n - \gamma M_i^n + \sum_{j \in N(i)} \gamma M_j^n \frac{T_{ji}}{\sum_{\ell \in N(j)} T_{j\ell}} + Q_i^n. \quad (7)$$

Reflecting boundaries returned unassigned flux to its source cell. Local saturation above 0.65 was conservatively redistributed to cells with available pore-volume capacity. Consequently, the discrete inventory satisfied

$$\sum_i M_i^{n+1} = \sum_i M_i^n + \sum_i Q_i^n \quad (8)$$

to floating-point precision.

Pressure change was generated by an iterated screened-diffusion response using the same directional permeability contrast. In continuous notation, the pattern approximates

$$-\nabla \cdot (k \nabla \psi) + \beta \psi = q_p, \quad (9)$$

and monitoring-time pressure change was scaled as

$$\Delta p(\mathbf{x}, t) = \Delta p_{\max} \sqrt{t/t_{\text{inj}}} \frac{\psi(\mathbf{x})}{\max_{\mathbf{x}} \psi(\mathbf{x})}, \quad (10)$$

with realization-specific $\Delta p_{\max}$ between 2.4 and 4.8 MPa. This pressure operator is a screening approximation used to introduce spatial pressure-saturation ambiguity. It is not a compositional pressure solve.

2.4 Synthetic 4D seismic observation model

Clean time-lapse acoustic-impedance change combined saturation and pressure effects:

$$\Delta Z_{\text{clean}} = G_{\sigma} \left[-a S_g + b \frac{\Delta p}{5 \text{ MPa}} \right], \quad (11)$$

where $a = 2.80$, $b = 0.36$, and G_{σ} is a three-dimensional Gaussian resolution operator with vertical and horizontal standard deviations of 0.35 and 0.90 cells. The monitored value was

$$\Delta Z_{\text{obs}} = \Delta Z_{\text{clean}} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_{\epsilon}^2). \quad (12)$$

The realization-specific noise standard deviation was 2.5% to 7.5% of clean-signal standard deviation, subject to a small lower bound. This operator captures band limitation, noise, and competing pressure and saturation contributions without claiming to reproduce elastic wave propagation.

2.5 Neural ensemble inversion

Each voxel was represented by 16 features: observed impedance change at three spatial scales; three directional gradients; a Laplacian; baseline impedance at two scales; normalized x , y , and z coordinates; normalized distance from the injector; monitoring time; injected pore-volume fraction; and k_v/k_h . A multilayer perceptron with hidden layers of 48 and 24 rectified-linear units predicted saturation and normalized pressure:

$$[\hat{S}_{g,i}, \hat{p}_i/5] = f_{\theta}(\mathbf{x}_i). \quad (13)$$

Network parameters minimized two-output mean-square error

$$L_{\text{data}}(\theta) = \frac{1}{N} \sum_{i=1}^N \|f_{\theta}(\mathbf{x}_i) - \mathbf{y}_i\|_2^2. \quad (14)$$

The training sample included 95,000 voxels and intentionally oversampled plume cells to prevent the zero-saturation background from dominating the loss. Four independently initialized networks used Adam optimization, batch size 512, L_2 coefficient 2×10^{-4} , and a maximum of 55 epochs.

Neural inversion and physics-guided correction




Figure 3. Neural inversion architecture and physics-guided correction.




Figure 4. Training-loss histories for four independently initialized neural networks.

2.6 Physics-guided projection

The neural mean was corrected by a projected-gradient procedure. The unconstrained objective was

$$J(S_g) = \frac{1}{2} \| S_g - \hat{S}_{g,N} \|_2^2 + \frac{\lambda_d}{2} \| G_\sigma[-aS_g + b\hat{p}/5] - \Delta Z_{\text{obs}} \|_2^2 + \frac{\lambda_s}{2} \| S_g - G_s(S_g) \|_2^2, \quad (15)$$

subject to

$$0 \leq S_{g,i} \leq 0.65, \quad \sum_i \phi_i V_i S_{g,i} = V_{\text{CO}_2}. \quad (16)$$

Here, $\hat{S}_{g,N}$ and $\hat{p}$ are neural-ensemble means and V_{CO_2} is known mobile CO_2 pore volume in the synthetic experiment. Eight weak smoothing iterations were followed by an exact bounded inventory projection. This design tests whether a physically important global constraint can be restored without materially degrading plume geometry.

2.7 Ensemble uncertainty and calibration

For $K = 4$ networks, ensemble mean and standard deviation were

$$\bar{S}_g = \frac{1}{K} \sum_{k=1}^K \hat{S}_g^{(k)}, \quad s_S = \sqrt{\frac{1}{K-1} \sum_{k=1}^K \left(\hat{S}_g^{(k)} - \bar{S}_g \right)^2}. \quad (17)$$

A validation multiplier $q_{0.90}$ was calculated from the 90th percentile of normalized absolute error:

$$q_{0.90} = Q_{0.90} \left(\frac{|S_g - \bar{S}_g|}{s_S + 0.005} \right). \quad (18)$$

The reported nominal 90% interval was $\bar{S}_g \pm q_{0.90}(s_S + 0.005)$. The same multiplier was frozen and applied to independent and out-of-distribution sets.

2.8 Evaluation metrics

Saturation root-mean-square error was

$$\text{RMSE}_S = \sqrt{\frac{1}{N} \sum_i (S_{g,i} - \hat{S}_{g,i})^2}. \quad (19)$$

Plume cells were defined by $S_g > 0.02$, and intersection-over-union was

$$\text{IoU} = \frac{|\Omega_{\text{true}} \cap \Omega_{\text{pred}}|}{|\Omega_{\text{true}} \cup \Omega_{\text{pred}}|}. \quad (20)$$

The mobile-volume error and seismic consistency error were

$$e_M = 100 \frac{|\sum_i \phi_i V_i \hat{S}_{g,i} - V_{\text{CO}_2}|}{V_{\text{CO}_2}}, \quad (21)$$

$$e_Z = \frac{\|G_\sigma[-a\hat{S}_g + b\hat{p}/5] - \Delta Z_{\text{obs}}\|_2}{\|\Delta Z_{\text{obs}}\|_2}. \quad (22)$$

Pressure accuracy was reported as RMSE in MPa. The local baseline used $\hat{S}_g = \text{clip}(-\Delta Z_{\text{obs}}/a, 0, 0.65)$ and zero pressure change. It represents a simple local conversion, not a full conventional seismic inversion.

3. Results

3.1 Numerical verification and training behavior

Maximum relative transport-balance error across all 100 realizations was 3.00×10^{-16} . This confirms conservation by the implemented transition operator, but it does not verify physical completeness of the simplified flow model. Data generation required 0.39 s, neural training 8.13 s, and the complete study 10.07 s in the reported x86-64 Linux CPU environment using Python 3.12.13.

All ensemble members reduced training loss throughout the fixed 55-epoch budget. Independent initialization produced nonzero ensemble spread and distinct model trajectories.

3.2 Independent-test inversion accuracy

Figure 5 and Table 2 compare three inversion strategies across 60 independent monitoring volumes. Local inversion achieved mean saturation RMSE of 0.0148 and plume IoU of 0.778. The neural ensemble reduced RMSE to 0.00574 and raised IoU to 0.897. Its pressure RMSE was 0.123 MPa.

Physics-guided projection produced saturation RMSE of 0.00642, which was 56.6% lower than local inversion but 11.9% higher than the unconstrained ensemble. The key change was inventory consistency: mean mobile-volume error declined from 3.26% to numerical precision. Seismic NRMSE increased from 0.136 to 0.165, showing that exact inventory consistency imposed a measurable data-fit tradeoff.




Figure 5. Independent-test performance for local, neural, and physics-guided inversion.

Table 2. Mean performance across 60 independent test volumes.

Method	Saturation RMSE	Plume RMSE	Plume IoU	Mobile-volume error (%)	Pressure RMSE (MPa)	Seismic NRMSE
Local rock-physics inversion	0.0148	0.0221	0.778	24.310	0.558	0.131
Neural ensemble	0.00574	0.00719	0.897	3.256	0.123	0.136
Physics-guided neural inversion	0.00642	0.00781	0.884	<0.000000000001	0.123	0.165

3.3 Spatial plume reconstruction

The representative test example in Figure 6 shows that local inversion broadened the plume boundary and produced diffuse saturation around the high-saturation core. The neural ensemble recovered the compact front more closely. Physics-guided projection preserved the main neural plume outline while shifting saturation to satisfy the global mobile-volume constraint. Ensemble standard deviation was concentrated near the advancing front and injector, which are locations where small seismic-response changes corresponded to larger saturation ambiguity.




Figure 6. Representative independent-test saturation predictions, error, and ensemble uncertainty.

3.4 Uncertainty calibration

The validation procedure produced $q_{0.90} = 0.773$. The resulting nominal 90% interval achieved 90.0% coverage on validation and 90.1% on the independent test set. Mean test half-width was 0.00929 saturation units. Because the multiplier was chosen on held-out realizations rather than training voxels, the result measures transfer to unseen geological patterns inside the specified prior range.

Coverage should not be interpreted as a universal probability guarantee. Voxel errors are spatially correlated, the calibration set is finite, and the interval is calibrated marginally rather than simultaneously over an entire plume.

3.5 High-connectivity out-of-distribution performance

Increasing k_v/k_h from 0.04-0.14 to 0.22-0.30 produced more vertically distributed plume geometries. Physics-guided saturation RMSE increased from 0.00642 to 0.00944, plume IoU declined from 0.884 to 0.768, and pressure RMSE increased from 0.123 to 0.173 MPa. Mean interval half-width more than doubled from 0.00929 to 0.01974, and coverage increased to 92.3%. The ensemble therefore detected part of the distribution shift through wider intervals, but the substantial geometric-accuracy loss still demonstrates limited transfer.




Figure 7. Representative physics-guided inversion under high vertical connectivity outside the training range.




Figure 8. Comparison of independent-test and high-connectivity out-of-distribution error and coverage.

Table 3. Physics-guided performance inside and outside the training connectivity range.

Evaluation set	k_v/k_h range	Saturation RMSE	Plume IoU	Pressure RMSE (MPa)	90% interval coverage	Mean half-width
Independent test	0.04-0.14	0.00642	0.884	0.123	0.901	0.00929
High-connectivity OOD	0.22-0.30	0.00944	0.768	0.173	0.923	0.01974

4. Discussion

The saturation RMSE was smallest with the unconstrained ensemble, and its predicted mobile inventory was on average 3.26% off from synthetic truth. That inventory error was eliminated by physics-guided projection without a drop in plume IoU (0.884). This result provides a good example of why choosing the inversion for voxelwise error is not a good choice. A small spatial error can exist with material integrated-volume bias for monitoring and accounting. The physics-guided solution was not the best at all of the metrics. It had a higher saturation RMSE and seismic NRMSE compared to the neural ensemble. The correction is thus a deliberate compromise and not the assertion that the introduction of constraints will necessarily enhance all predictions. The constraint weights and inventory tolerance should account for uncertainty in application to the field in the following: injection accounting; rock physics; acquisition repeatability; baseline reservoir characterization. 4.3 Measurement of the effects of natural resource management practices. 4.4 Application of a more comprehensive monitoring framework. The results validate previous studies on the rapid saturation inference from simulated seismic responses using neural methods (Um et al., 2024) and the stabilization effect of flow or geostatistical information on inversion (Liu et al., 2023, 2024; Li et al., 2026). The present contribution is a more limited approach than full-waveform or multiphysics inversion. It illustrates in a practical experiment using reduced resolution, how realization-wise validation, mobile-volume projection and an explicit connectivity shift affect standard

performance drawing conclusions. This model is easier to interpret than deep surrogate data-assimilation workflows (Tang et al., 2021) because it directly predicts the monitoring states as opposed to updating an entire reservoir-parameter ensemble. Ensemble spread is used as a computationally efficient uncertainty indicator and is calibrated from independent ensemble realisations compared to the Bayesian time-lapse inversion (Grana & Mukerji, 2015). The above simplifications are necessary to minimize the computational burden, and limit the interpretation of uncertainty. 4.3 The importance of shift in geographical distribution The high-connectivity test caused an increase of 47.0% in saturation RMSE and a decrease in plume IoU. The ensemble intervals had nominal marginal coverage however this does not represent the preservation of plume geometry. This limitation cannot be identified by a model evaluated solely by random voxel splitting as there is strong structural dependence between neighboring voxels and repeated monitoring times. Realization-wise separation and physically defined out-of-distribution scenarios should be considered as minimum criteria for claims of geological generalization. 4.4 Limitations This study is not a field demonstration, but is an original, repeatable screening experiment. The transport operator saves mobile pore volume and accounts for the anisotropic and heterogeneous conductance, but is not a solution of a complete compositional CO₂-brine system. These are eliminated: Capillary pressure, residual trapping, dissolution, geochemical reactions, fault leakage, and geomechanical deformation. The pressure is created using a screened-diffusion approximation. The seismic operator transforms the saturation and pressure to change in acoustic-impedance in the bands. It does not include the effect of elastic wave propagation, acquisition geometry, multiples, anisotropy, attenuation or nonrepeatability between surveys. Therefore, the low errors indicate that the synthetic problem was controlled, and cannot be taken as expected field accuracy. This is a neural model, not a three-dimensional convolutional network, provided with multiscale neighborhood features. The ensemble is of four members and includes only variability associated with initialization. Validation-calibrated intervals are at the edges and don't account for structural uncertainty in transport or rock-physics equations. When it comes to field application, compositional simulation, calibrated petroelastic modeling, realistic simulation of surveys, independent monitoring data and prospective validation are required.

5. Conclusions

A physics-guided neural inversion of synthetic 4D seismic data to characterize CO₂ plume was evaluated using an original 3D simulation study. According to the results, data separation in realisation allows to evaluate heterogeneous reservoirs reliably without losing information within the process of the monitoring snapshot and without leaking the information between different snapshots. The neural ensemble also achieved a lower saturation RMSE of 0.00574 and higher IoU for plume of 0.897 than local rock-physics inversion, which had a RMSE of 0.0148 and an IoU of 0.778. The Physics-guided projection reduced the test saturation RMSE to 0.00642 and mean mobile volume error from 3.26% to numerical precision. Within the training connectivity range, the nominal 90% uncertainty interval covered 90.1% of the data, while outside of training connectivity range, increasing the vertical connectivity resulted in increasing RMSE to 0.00944 and decreasing plume IoU to 0.768. The results show that, in order to build physics-guided monitoring systems, the accuracy of the neural model, consistency of inventory, and reliability for out-of-distribution data should be considered simultaneously..

References

1. Ajo-Franklin, J. B., Peterson, J., Doetsch, J., & Daley, T. M. (2013). High-resolution characterization of a CO₂ plume using crosswell seismic tomography: Cranfield, MS, USA. *International Journal of Greenhouse Gas Control*, 18, 497-509. <https://doi.org/10.1016/j.ijggc.2012.12.018>

2. Furre, A.-K., Eiken, O., Alnes, H., Vevatne, J. N., & Kiaer, A. F. (2017). 20 years of monitoring CO₂-injection at Sleipner. *Energy Procedia*, 114, 3916-3926. <https://doi.org/10.1016/j.egypro.2017.03.1523>
3. Grana, D., & Mukerji, T. (2015). Bayesian inversion of time-lapse seismic data for the estimation of static reservoir properties and dynamic property changes. *Geophysical Prospecting*, 63(3), 637-655. <https://doi.org/10.1111/1365-2478.12203>
4. Li, H., Wang, N., Durlowsky, L. J., & Biondi, B. L. (2026). Bayesian full-waveform monitoring of CO₂ storage with fluid-flow priors via generative modeling. *Journal of Geophysical Research: Machine Learning and Computation*, 3(3), e2025JH001190. <https://doi.org/10.1029/2025JH001190>
5. Liu, M., Grana, D., & Mukerji, T. (2024). Geostatistical inversion for subsurface characterization using Stein variational gradient descent with autoencoder neural network: An application to geologic carbon sequestration. *Journal of Geophysical Research: Solid Earth*, 129(7), e2024JB029073. <https://doi.org/10.1029/2024JB029073>
6. Liu, M., Vashisth, D., Grana, D., & Mukerji, T. (2023). Joint inversion of geophysical data for geologic carbon sequestration monitoring: A differentiable physics-informed neural network model. *Journal of Geophysical Research: Solid Earth*, 128(3), e2022JB025372. <https://doi.org/10.1029/2022JB025372>
7. Mardan, A., Giroux, B., Fabien-Ouellet, G., & Saberi, M. R. (2023). Monitoring fluid saturation in reservoirs using time-lapse full-waveform inversion. *Geophysical Prospecting*, 71(6), 1012-1029. <https://doi.org/10.1111/1365-2478.13363>
8. Tang, H., Fu, P., Sherman, C. S., Zhang, J., Ju, X., Hamon, F., Azzolina, N. A., Burton-Kelly, M. E., & Morris, J. P. (2021). A deep learning-accelerated data assimilation and forecasting workflow for commercial-scale geologic carbon storage. *International Journal of Greenhouse Gas Control*, 112, 103488. <https://doi.org/10.1016/j.ijggc.2021.103488>
9. Um, E. S., Alumbaugh, D., Lin, Y., & Feng, S. (2024). Real-time deep-learning inversion of seismic full waveform data for CO₂ saturation and uncertainty in geological carbon storage monitoring. *Geophysical Prospecting*, 72(1), 199-212. <https://doi.org/10.1111/1365-2478.13197>
10. Williams, M. O., Ampomah, W., Sibaweihi, N., Amosu, A., & Dupree, R. (2026). Optimizing CO₂ storage in a heterogeneous depleted gas reservoir: A case study from the Evetts Site, Permian Basin. *Carbon Capture, Utilization, and Storage Conference* (Paper CCUS-4436159). <https://doi.org/10.15530/ccus-2026-4436159>