CCUS: 4443708

Optimizing Long-Term Geological Carbon Storage Stability via Deep Learning-Assisted Particle Swarm Optimization

Guodong Ren[1], Oluwatobi Akinyede*[1], Siddharth Misra[1], 1. Texas A&M University, College Station, TX, United States.

Abstract

This study optimizes CO_2 injection strategies for geological carbon storage by integrating deep-learning surrogates with Particle Swarm Optimization (PSO) to maximize long-term stable trapping of CO_2 in the form of residual and solubility trapped CO_2. This optimization workflow designs injection schedules to maximize residually and solubility trapped CO_2 while minimizing movable trapped CO_2 under diverse geological settings; thereby promoting long-term storage stability. The injection schedules can have variable rates and intermittent shut-ins to better replicate realistic field injection operations. A trained long short-term memory (LSTM) sequence-to-sequence surrogate model is employed, using nine geological properties and dynamically varying injection profiles as inputs. PSO with 4,000 particles is coupled to the surrogate to efficiently explore high-dimensional control spaces under operational constraints. The objective function combines trapping metrics with penalty–reward terms to encourage near-maximum allowable injection volumes while maintaining feasibility. Optimized injection strategies significantly outperform constant-rate schemes, increasing combined residual and solubility trapping by up to 50% under favorable geological settings and by more than 15% over a 30-year injection period. The surrogate-based optimization evaluates thousands of candidate schedules within minutes and consistently converges within ten iterations, demonstrating computational tractability and robustness for large-scale storage design.

Introduction

Geological carbon storage (GCS) in saline aquifers requires injection strategies that promote long-term immobilization of CO_2 through residual and solubility trapping while limiting the persistence of mobile plumes. Injection rate scheduling directly influences pressure evolution and plume migration, thereby affecting post-injection trapping efficiency and storage security (Huber et al., 2016). Despite its importance, systematic optimization of time-varying injection strategies remains computationally challenging. High-fidelity compositional reservoir simulators accurately resolve coupled multiphase and multicomponent flow processes but require substantial computational effort over multi-decadal horizons. A single full-physics saline-aquifer simulation spanning several decades may require tens of minutes of wall-clock time, rendering optimization workflows that require thousands of forward evaluations impractical (Jin and Durlofsky, 2018). Recent advances in deep-learning-based surrogate modeling have enabled rapid forecasting of CO_2 plume migration and trapping behavior across heterogeneous geological settings (Falola et al., 2025; Nunez et al., 2025). Sequence-based surrogate models have demonstrated the ability to reproduce the temporal evolution of plume geometry and trapping metrics under dynamically varying injection conditions with orders-of-magnitude reduction in computational cost relative to full-physics simulation (Talabi et al., 2026; Ren et al., 2024). While prior studies have established the predictive capability of data-driven surrogates, most applications focus on forward forecasting under prescribed injection schedules. Fewer efforts have addressed direct optimization of time-varying injection strategies within a surrogate framework, particularly for enhancing residual and solubility trapping. Analytical and numerical studies indicate that injection rate modulation can accelerate residual trapping and alter plume evolution (Huber et al., 2016), motivating optimization-based approaches that systematically exploit this control. This work addresses this gap by coupling a pre-trained long short-term memory (LSTM) sequence-to-sequence surrogate with Particle Swarm Optimization (PSO) to directly optimize dynamic CO_2 injection schedules. The objective is to identify feasible, time-varying injection strategies that maximize residual and solubility trapping under realistic operational constraints without repeated invocation of computationally expensive reservoir simulators**.**

Method

Figure 1 summarizes the surrogate-based optimization framework. A pre-trained sequence-to-sequence long short-term memory (LSTM) surrogate serves as a fast forward model that emulates a full-physics compositional reservoir simulator.

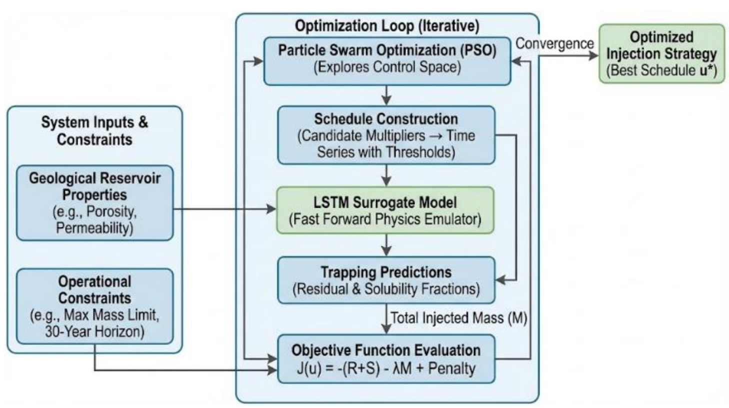

Figure 1. Conceptual workflow for surrogate-based optimization of dynamic CO_2 injection strategies. The workflow integrates a pre-trained LSTM surrogate model with Particle Swarm Optimization to maximize CO_2 trapping efficiency. The surrogate serves as a fast-forward emulator, rapidly mapping static geological properties and time-varying injection schedules to end-state residual and solubility trapping fractions. The PSO algorithm iteratively refines the 30-year injection strategy by evaluating the objective function which incentivizes storage performance and mass utilization while enforcing operational constraints via a penalty mechanism.

The surrogate maps static geological properties together with a time-varying CO_2 injection rate schedule to end-of-injection trapping outcomes. The geological inputs comprise nine normalized parameters describing reservoir and fluid conditions, while the injection input is a 30-year time series of monthly CO_2 injection rates. The surrogate outputs are the predicted fractions of CO_2 trapped residually and dissolved in brine at the end of the injection period. The surrogate architecture and training procedure follows Talabi et al. (2026) and Ren et al. (2024), where similar models were shown to accurately reproduce plume migration and trapping trends under dynamic injection conditions. The injection horizon is discretized into uniform six-month control blocks, yielding 60 segments over 30 years. Injection within each block is defined by a dimensionless rate multiplier. The first multiplier is fixed to establish a physically reasonable initial injection condition, while the remaining 59 multipliers constitute the optimization variables. Multipliers below a prescribed threshold are treated as effective shut-ins, enabling intermittent injection patterns consistent with operational practice. Particle Swarm Optimization (PSO) is used to search for the resulting high-dimensional control space in a derivative-free manner. PSO is selected for its robustness in nonlinear, nonconvex problems and its ability to efficiently explore large decision spaces without gradient information. At each iteration, PSO proposes candidate injection schedules, which are expanded into monthly rate profiles and evaluated using the surrogate model. The optimization objective is defined as:

$$\boldsymbol{J}(\boldsymbol{u}) = -\big(\boldsymbol{R}(\boldsymbol{u}) + \boldsymbol{S}(\boldsymbol{u})\big) - \boldsymbol{\lambda M}(\boldsymbol{u}) + \boldsymbol{P}(\boldsymbol{M}(\boldsymbol{u}))$$

where $u = \{u_2, \ldots, u_{60}\}$ denotes the optimized injection multipliers, $R(u)$ and $S(u)$ are the residual and solubility trapping fractions at the end of injection, and $M(u)$ is the total injected CO_2 mass. The objective is formulated as a cost minimization. The term $-\big(R(u) + S(u)\big)$ promotes trapping efficiency, while the small reward coefficient λ discourages trivial under-injection by favoring schedules that utilize available injection capacity. Operational feasibility is enforced through the penalty term P, which activates only when the total injected mass exceeds a prescribed upper limit; otherwise, P=0. This separation of physical performance objectives from feasibility constraints yields stable optimization behavior and rapid convergence. Coupling PSO with the trained surrogate enables thousands of candidate injection schedules to be evaluated within minutes, making systematic optimization of time-varying injection strategies computationally tractable without repeated high-fidelity reservoir simulation.

Results

Figure 2 illustrates an example optimized injection schedule obtained using the surrogate-assisted Particle Swarm Optimization framework for a fixed geological setting, objective definition, and operational constraints.

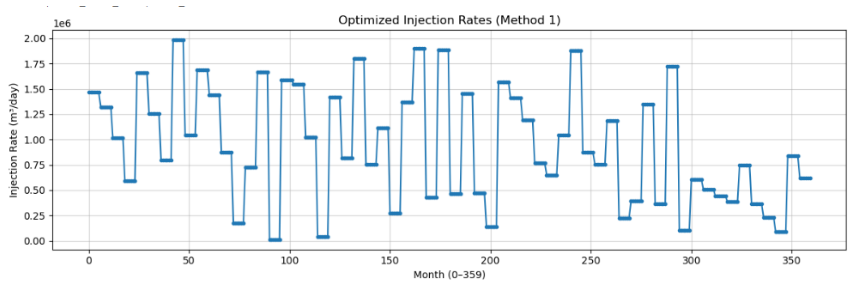

Figure 2. Optimized time-varying CO_2 injection schedule for a single geological and operational scenario. Injection rates are defined over six-month control blocks across a 30-year injection period and are determined by surrogate-assisted Particle Swarm Optimization to maximize combined residual and solubility trapping subject to a constraint on total injected CO_2 volume. The optimized schedule exhibits non-uniform rate modulation and intermittent low-rate periods resulting from the optimization of block-wise injection multipliers.

The optimization targets maximization of combined residual and solubility trapping at the end of a 30-year injection period, subject to a constraint on total injected CO_2 volume. The resulting injection schedule departs from constant-rate injection by exhibiting time-varying rate modulation and intermittent low-rate or shut-in periods over the injection horizon. The optimized injection schedule is not rule-based or pre-defined but is an emergent outcome of optimizing time-varying injection rates against a trapping-based objective under operational constraints. For the optimized solution shown in Figure 2, the framework achieves a combined residual and solubility trapping fraction of 0.48 at a total injected CO_2 volume of 27 MMT, corresponding to residual trapping of 0.34 and solubility trapping of 0.14 at the end of injection. This outcome demonstrates the ability of the surrogate-based optimization framework to identify feasible, non-uniform injection schedules that enhance trapping metrics relative to unconstrained or manually selected rate profiles. The optimization converges within ten PSO iterations for this case, despite the 59-dimensional control space defined by six-month injection blocks. The use of a pre-trained surrogate model enables rapid evaluation of thousands of candidate schedules per iteration without repeated execution of a full-physics reservoir simulator, making such high-dimensional optimization computationally tractable

Discussion

Sensitivity of the optimized injection schedule to operational deviations is examined by applying stochastic perturbations of ±5–25% to the block-wise injection multipliers (Figure 3). The optimal schedule achieves a combined residual and solubility trapping fraction of 0.48, whereas perturbed schedules yield lower combined trapping values, ranging from approximately 0.41 to 0.43 across all perturbation levels tested. Moderate perturbations (±5–15%) result in limited degradation of trapping performance, while larger perturbations (±20–25%) lead to greater variations in trapping outcomes.

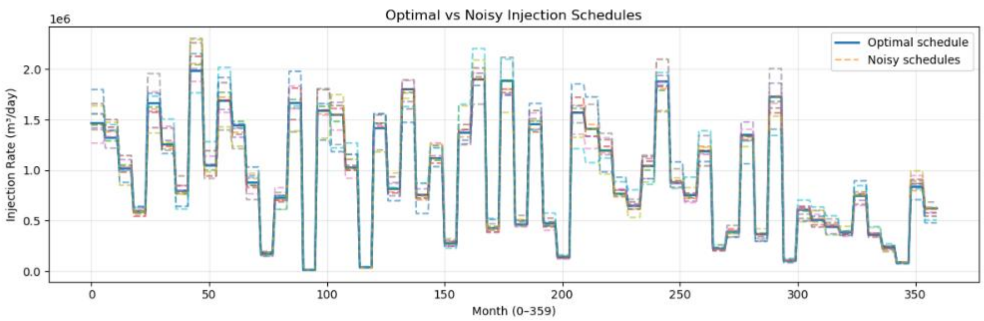

Figure 3. Optimized injection schedule compared with stochastically perturbed schedules obtained by applying ±5–25% random variations to block-wise injection multipliers. Perturbed schedules exhibit similar temporal structure but yield reduced combined residual and solubility trapping relative to the optimal case

These results indicate that trapping efficiency is sensitive to deviations from the optimized temporal injection pattern. Although the perturbations do not enforce constant injected volume or provide a comparison against constant-rate injections, they demonstrate that departures from the optimized schedule systematically reduce trapping performance. This finding highlights the importance of maintaining the optimized time-varying injection structure when translating surrogate-optimized schedules into practical injection designs.

Conclusions

This study demonstrates that coupling a sequence-to-sequence LSTM surrogate with Particle Swarm Optimization enables efficient exploration of time-varying CO_2 injection schedules to optimize geological carbon storage for stable, long-term storage. By replacing repeated computationally expensive, full-physics simulation with a fast deep-learning-based surrogate, the framework evaluates thousands of candidate injection schedules within minutes to boost residual and solubility trapping. Relative to non-optimized schedules under the same operational and geological constraints, the deep-learning boosted optimization workflow identifies dynamic strategies that increase combined residual and solubility trapping, representing stable, long-term storage, while decreasing the movable CO_2 structurally trapped under the seal. For the cases investigated, optimized injection schedules achieved up to 0.48 combined residual and solubility trapping at a total injected volume of 27 MMT over a 30-year period. Perturbation tests indicate that deviations from the optimized schedule reduce residual and solubility trapping performance by approximately 10-15%, highlighting sensitivity to injection control. Overall, the results illustrate the potential of surrogate-assisted optimization as a practical tool for injection strategy design within computational and operational limit to find injection scenarios that can lead to enhanced geological carbon storage.

References

Falola, Y., Misra, S., and Nunez, A. 2025. Geological and engineering insights from transfer learning with Fourier neural operators: A case study of CO_2 storage forecasting in disparate saline aquifers. Applied Soft Computing 178: 113272. https://doi.org/10.1016/j.asoc.2025.113272.

Huber, E.J., Stroock, A.D., and Koch, D.L. 2016. Analysis of a time-dependent injection strategy to accelerate the residual trapping of sequestered CO_2 in the geologic subsurface. International Journal of Greenhouse Gas Control 44: 185–198. https://doi.org/10.1016/j.ijggc.2015.11.024.

Jin, Z.L. and Durlofsky, L.J. 2018. Reduced-order modeling of CO_2 storage operations. International Journal of Greenhouse Gas Control 68: 49–67. https://doi.org/10.1016/j.ijggc.2017.08.017.

Nunez, A., Misra, S., and Falola, Y. 2025. Transfer learning for geological carbon storage forecasting using neural operator. Advances in Water Resources 199: 104948. https://doi.org/10.1016/j.advwatres.2025.104948.

Talabi, O., Ren, G., and Misra, S. 2026. AI-driven rapid forecasting of CO_2 plume migration and trapping efficiency in geological carbon storage. International Journal of Greenhouse Gas Control 149: 104535. https://doi.org/10.1016/j.ijggc.2025.104535.

Tang, M., Ju, X., and Durlofsky, L.J. 2022. Deep-learning-based coupled flow–geomechanics surrogate model for CO_2 sequestration. International Journal of Greenhouse Gas Control 118: 103692. **https://doi.org/10.1016/j.ijggc.2022.103692.**