

Brownian Bridge Augmented Surrogate Simulation and Injection Planning for Geological CO₂ Storage

Haoyue Bai¹, Guodong Chen^{2,3}, Wangyang Ying¹, Xinyuan Wang¹, Nanxu Gong¹, Sixun Dong¹, Giulia Pedrielli¹, Haoyu Wang⁴, Haifeng Chen⁴, Yanjie Fu¹

¹Arizona State University

²University of California Berkeley

³Lawrence Berkeley National Laboratory

⁴NEC Labs America

{haoyueba, wying4, xwang735, ngong6, sdong46, giulia.pedrielli, yanjie.fu}@asu.edu, gchen6@lbl.gov, {haoyu, haifeng}@nec-labs.com

Abstract

Geological CO₂ storage (GCS) involves injecting captured CO₂ into deep subsurface formations to support climate goals. The effective management of GCS relies on adaptive injection planning to dynamically control injection rates and well pressures to balance both storage safety and efficiency. Prior literature, including numerical optimization methods and surrogate-optimization methods, is limited by real-world GCS requirements of smooth state transitions and goal-directed planning within limited time. To address these limitations, we propose a Brownian Bridge–augmented framework for surrogate simulation and injection planning in GCS and develop two insights (i) Brownian bridge as smooth state regularizer for better surrogate simulator; (ii) Brownian bridge as goal-time-conditioned planning guidance for better injection planning. Our method has three stages: (i) learning deep Brownian bridge representations with contrastive and reconstructive losses from historical reservoir and utility trajectories, (ii) incorporating Brownian bridge-based next state interpolation for simulator regularization (iii) guiding injection planning with Brownian utility-conditioned trajectories to generate high-quality injection plans. Experimental results across multiple datasets collected from diverse GCS settings demonstrate that our framework consistently improves simulation fidelity and planning effectiveness while maintaining low computational overhead.

Code — https://github.com/HaoyueBai98/Brownian_Bridge_Augmented_CO2_Storage

Introduction

Geological CO₂ storage (GCS) is a technique that involves injecting captured carbon dioxide into deep subsurface formations, such as saline aquifers and depleted reservoirs, for long-term containment (Mao and Ghahfarokhi 2024; Ismail and Gaganis 2023). GCS plays a role in climate change mitigation, energy system resilience, and the transition to carbon neutrality. GCS operates by capturing, compressing, and injecting CO₂ into the deep subsurface through a sequence of injection events while controlling injection rates and well

pressures over time. Effective GCS management needs to ensure: i) safety: preventing excessive reservoir pressures, leakage, and induced seismicity, and ii) efficiency: optimizing the reservoir’s storage capacity and improving injectivity. One of the essential tasks in GCS is adaptive injection planning that dynamically controls CO₂ injection rates and well pressures over time to ensure safety and efficiency.

Solving the adaptive injection planning requires adaptive simulation of evolving reservoir states over time and optimization of the injection plan (Chen, Huan, and Ma 2006). In prior literature, numerical simulation and optimization based methods can provide accurate results but are computationally intensive and prohibitively slow for real-time or large-scale applications (Zhang and Agarwal 2013; Liu et al. 2021; Witte, Hewett, and Chandra 2022; Song et al. 2023). To alleviate these issues, integrated surrogate modeling with optimization algorithms has emerged as an effective alternative (Liu et al. 2024a; Sun et al. 2021). Surrogate models can provide fast and accurate approximations of high-fidelity simulations, reduce computational costs, and enable iterative optimization in complex storage scenarios.

Real world GCS practices impose two new requirements of adaptive injection planning for deployments: i) In physical systems like CO₂ injection and geochemical reactions in GCS, state transitions are typically smooth and gradual; Without mechanisms for temporal continuity and smoothing, surrogate models are sensitive to noise and variance in data, often producing abrupt state changes that distort subsurface physics and lead to unreliable predictions; ii) GCS planning is inherently goal-oriented and time-constrained: injection planning aims to achieve a target utility (e.g., sufficient CO₂ storage, pressure stability) within a limited time. Thus, injection planning requires the ability to guide decisions toward a target goal within the limited time and to maintain alignment with a consistent, goal-directed trajectory over time. The two practical requirements call for a new surrogate-optimization framework.

Our Perspective: deep Brownian bridges as smooth state regularizer and goal-time conditioned planning guidance. As a stochastic process, Brownian bridge (Revuz and Yor 2013) provides unique mathematical properties: i) *smooth and gradual transitions*: the Brownian bridge mod-

els trajectories that are continuous and smooth between a defined start and end point, which aligns with GCS. ii) *goal and time-constrained trajectory planning*: the Brownian bridge is conditioned to reach a specific target state at a designated future time and can ensure decision-making to stay on a projected, goal-aligned path. Along these lines, we identify three insights centered on Brownian bridge to advance surrogate simulation and injection planning: i) a deep version of Brownian bridge can learn a Brownian embedding space in which we can incorporate smooth and gradual state transitions into the surrogate simulator, and incorporate goal- and time-conditioned planning into adaptive injection. ii) the next state interpolation in Brownian embedding space can serve as an auxiliary supervision signal to regularize the surrogate simulator to learn smooth and physically consistent transitions. iii) the Brownian bridge projected trajectory can be seen as goal- and time-conditioned guidance in injection planning toward a high storage utility goal yet completed at a specific time.

Contributions. 1) Problem: We tackle the AI for science problem: adaptive injection in GCS as a simulation to optimization task. 2) Framework: We propose a Brownian Bridge-augmented surrogate simulation and injection planning framework, where the simulator predicts both storage utility and future reservoir states given the current condition and injection plan, while the planner generates injection plans over time. 3) Techniques: we leverage contrastive and reconstructive losses to learn two deep Brownian bridges, each structured by an encoder, a generator, and a decoder, for reservoir states and storage utilities from observed data; we leverage Brownian bridge-interpolated next state as auxiliary supervision to regularize simulator learning; we leverage Brownian bridge-interpolated trajectory as goal-time conditioned guidance for injection planning model learning. 4) Validations: extensive experiments on multiple CO₂ injection datasets demonstrate that our approach consistently achieves superior operational efficiency, improved predictive accuracy, and enhanced strategic robustness compared to state-of-the-art methods.

Preliminaries and Problem Statement

Geological CO₂ Storage. Geological CO₂ storage (GCS) is a cornerstone technology in climate change mitigation and energy system decarbonization. GCS prevents CO₂ release into the atmosphere by injecting it through multiple wells into deep geological formations, such as, saline aquifers and unmineable coal seams (Figure 1(left)). Because geological formations exhibit heterogeneous properties and slow diffusion-driven dynamics, which means the full impact of an injection decision may take months or even years to manifest (Juanes et al. 2006), effective CO₂ sequestration requires multiple injection steps, continuous monitoring, feedback, and adaptive control, often over a long period (e.g., decades). The entire storage process is structured into a sequence of injection stages, collectively referred to as the *GCS lifecycle*. Figure 1(middle) shows that each lifecycle consists of multiple discrete time steps. At each step, the system can be characterized by a *reservoir state* and requires the design of an *injection plan*. The utility of this action is

observed through a measurable *storage utility*.

Reservoir State, Injection Plan, and Storage Utility in GCS Systems. In GCS, the reservoir state of the storage system includes reservoir pressure, gas saturation, oil saturation, and water saturation at various spatial grid points. These features describe the reservoir’s condition and its response to different operational strategies. We define this reservoir state as a vector $\mathbf{o} \in \mathbb{R}^{|\mathbf{o}|}$, where \mathbf{o}_t represents the reservoir state at the t^{th} time step. The system’s operation relies on strategic decisions, such as adjusting CO₂ injection rates and well pressure. We represent these strategies as a vector $\mathbf{s} \in \mathbb{R}^{|\mathbf{s}|}$, where \mathbf{s}_t is the injection plan at the t^{th} time step. Each dimension of \mathbf{s}_t is a decision variable of the t -th well, for example, the CO₂ injection rate of the t -th well. The effectiveness of these strategies is assessed through performance metrics, which capture utilities like injection efficiency, storage integrity, and pressure stability. These metrics form a storage utility vector $\mathbf{r} \in \mathbb{R}^{|\mathbf{r}|}$, with \mathbf{r}_t denoting the storage utility at the t^{th} time step. Over a complete lifecycle, these storage utilities provide a comprehensive view of storage performance, revealing whether the strategies achieve long-term stability and effectiveness.

The AI for Science Problem: Surrogate Simulation and Injection Plan Optimization. Given a dataset $\mathbb{D} = \{\mathbf{o}, \mathbf{s}, \mathbf{r}\}^{N \times T}$, comprising N GCS trajectories, each with a lifecycle length of T time steps, our task is twofold: i) we aim to construct an accurate surrogate simulator \mathcal{S} that maps the reservoir states \mathbf{o} and the strategies \mathbf{s} to corresponding storage utility \mathbf{r} . ii) we aim to develop a decision-making model \mathcal{D} that generates optimal strategies \mathbf{s} based on the current reservoir state \mathbf{o} . The ultimate objective is to optimize the decision model to produce strategies maximizing the average storage utility over the entire lifecycle of the system.

Technical Background: The Brownian Bridge. A Brownian bridge is a stochastic process characterized by Brownian motion conditioned on fixed start and end points (Revuz and Yor 2013) (Figure 1(right)). It models a continuous trajectory that begins at a specified starting value and ends at a designated target value, while evolving stochastically between these two points. Formally, given the starting point z_0 at the time $t = 0$ and the ending point z_T at the time $t = T$, the Brownian bridge at an intermediate time step $t \in [0, T]$ is defined by a Gaussian distribution:

$$z_t \sim \mathcal{N} \left(\left(1 - \frac{t}{T} \right) z_0 + \frac{t}{T} z_T, \frac{t(T-t)}{T} \right). \quad (1)$$

This formulation represents a probabilistic interpolation between the initial and final points, with uncertainty maximized midway and minimized at the endpoints.

Proposed Methodology

Framework Overview

To advance GCS management and operations, we propose a Brownian Bridge-enhanced surrogate simulation and injection planning model framework (Figure 2), where the surrogate simulator estimates the storage utility and next reservoir state given a reservoir state and an injection plan; the injection planning model generates adaptive injection plans.

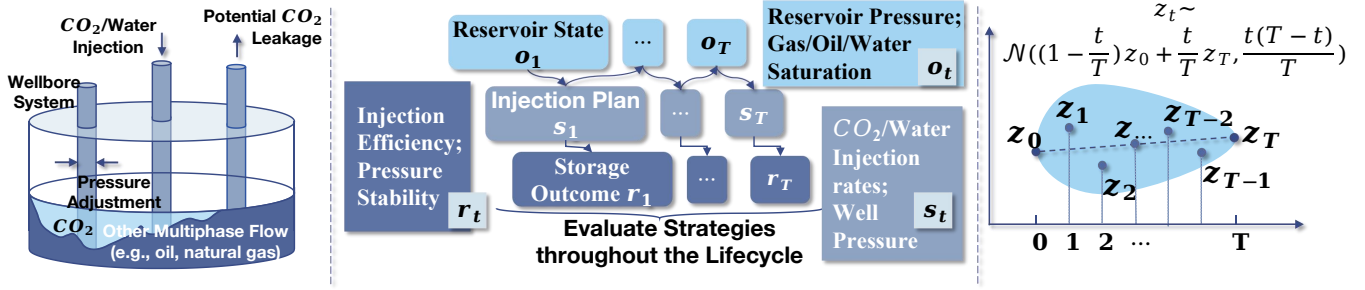


Figure 1: Problem and Technique Background

We found that Brownian bridge provides two opportunities: i) the ability to enforce smooth state transitions and interpolate the next state can serve as auxiliary supervision and regularization signals to advance next reservoir state estimation; ii) the ability to condition a model to reach a goal can improve knowledge-guided injection planning. Our method includes three steps: *Step 1* integrates contrastive and reconstruction losses to learn two deep Brownian bridges to embed desired reservoir state and storage utility trajectories into a Brownian latent space. *Step 2* incorporates the smooth interpolation of the next reservoir state in Brownian space to regularize surrogate simulator learning. *Step 3* leverages the goal-conditioned pursuing ability of Brownian bridges to guide long-term and forward-thinking injection planning.

Deep Brownian Bridge: GCS Lifecycle State and Utility Dynamics Modeling as Interpolation in Brownian Embedding Space

Given a start point and an end point, the classic Brownian bridge is a Gaussian process that models the most possible trajectory between two points. Inspired by (Wang, Lin, and Li 2023; Wang et al. 2022), we extend the Brownian bridge to a latent representation space, where a trajectory is represented as a sequence of latent embedding vectors, instead of a point sequence in an explicit Gaussian space. There are two benefits of using the latent Brownian bridge: i) Less susceptible to non-Gaussian scenarios, and more robust and generalized. ii) After mapping real system reservoir state or storage utility dynamics to a Brownian space, predicting the next reservoir state or inferring a path toward optimal storage utility is as simple as direct proportional interpolation, instead of using complex parameterized deep recurrent network families. In our AI for science problem, given the GCS lifecycle training data of many CO₂ injection event sequences, the latent Brownian bridge learning has two tasks: i) Learn a Brownian encoder to encode the start and end points to two embedding vectors in a latent space; ii) learn a Brownian generator to interpolate a latent trajectory between the start and end points. We propose to integrate data augmentation, Brownian embedding encoder and interpolation generator, contrastive and reconstruction losses to solve the two tasks.

Data Augmentation to Overcome Science Data Scarcity.

We need large training data, for example, the injection event sequences (e.g., injection plan, state changes, utility) of

many GCS lifecycles, to train a Brownian bridge to understand, represent, and infer the reservoir state and storage utility dynamics of a GCS system. However, GCS lifecycle data are scarce because i) real-world earth science data are limited and ii) numerical simulation is costly (Witte, Hewett, and Chandra 2022; Song et al. 2023). With limited GCS data available, we develop a two-step data augmentation method to generate more training data from existing GCS lifecycle event sequences and enable self-supervised learning to represent and reason Brownian bridges. The first step is sampling: given a GCS lifecycle (i.e., a complete state-injection-utility injection event sequence), we randomly sample diverse fixed-length subsequences. The second step is diversification: we add controlled Gaussian noises to each sampled subsequence to create more diversified subsequences to enrich training data. Here, we can control how much noise to add (stochastic variation) by a scaling factor α .

Learning Reservoir State and Storage Utility Related Deep Brownian Bridges. We propose to learn a deep encoder-generator-decoder model to map a sequence of dynamic reservoir state vectors or storage utility vectors into a latent embedding space, where the path between two points is defined and regularized by the Brownian bridge.

State Bridge: Brownian Encoder, Generator and Decoder for Reservoir State Sequences. We aim to construct a state-related deep Brownian bridge for reservoir state sequences. The model includes three parts: i) The state-related Brownian encoder is to map an original state vector at a time step into a latent embedding, given by: $\mathbf{z}_t = \mathcal{B}_e^o(\mathbf{o}_t)$, where \mathcal{B}_e^o is the Brownian encoder, \mathbf{o}_t is the reservoir state at time step t , \mathbf{z}_t is the latent embedding. ii) the state-related Brownian generator is to interpolate a smooth trajectory between the embedding of the start point and the embedding of the end point under the definition of Brownian bridge, given by:

$$\hat{\mathbf{z}}_t = \mathcal{B}_g^o(\mathbf{z}_0, \mathbf{z}_T, t) = \left(1 - \frac{t}{T}\right) \mathbf{z}_0 + \frac{t}{T} \mathbf{z}_T, \quad (2)$$

where \mathcal{B}_g^o is the Brownian generator and $\hat{\mathbf{z}}_t$ is generated latent embedding at time t ; iii) the state-related Brownian decoder is to reconstruct original reservoir state vectors over time from embeddings over time, given by: $\hat{\mathbf{o}}_t = \mathcal{B}_d^o(\hat{\mathbf{z}}_t)$, where \mathcal{B}_d^o is the Brownian decoder and $\hat{\mathbf{o}}_t$ is reconstruct original reservoir state vector. The objective function includes: i) a reconstruction loss that measures the gap between the Brownian decoder outputs and the original reservoir state in-

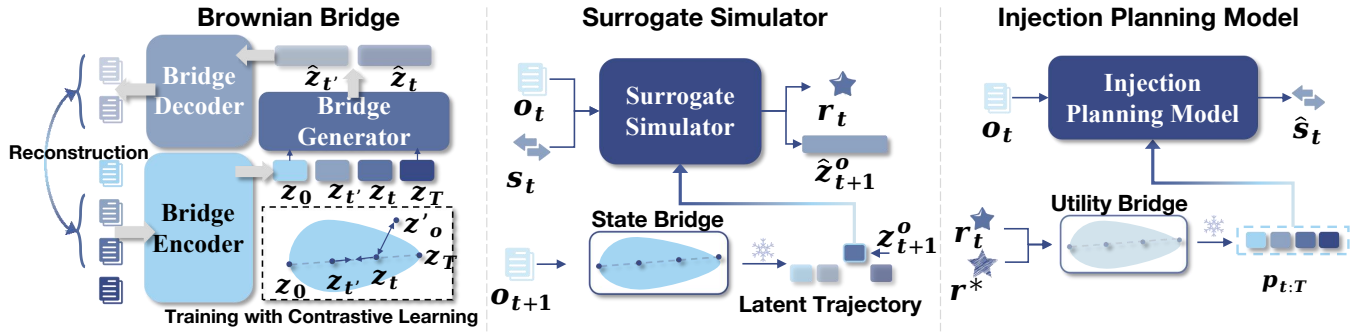


Figure 2: Framework

put, and ii) a contrastive loss that pulls positive pairs closer in latent space and pushes negative pairs apart, where positive pairs are latent embeddings from the same reservoir state sequence, and negative pairs are sampled from different reservoir state sequences, by minimizing the following:

$$\begin{aligned} & \min \|\mathcal{B}_d(\hat{\mathbf{z}}_t) - \mathbf{o}_t\|_2^2 \\ & - \log \sum_{i,j,k} \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)}{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau) + \sum_k \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k)/\tau)} \end{aligned} \quad (3)$$

where \mathcal{B}_d^o is the Brownian decoder implemented by a MLP, $\text{sim}(\cdot, \cdot)$ are the cosine similarity between two vectors, and τ is a temperature hyperparameter that controls the sharpness of the softmax distribution, \mathbf{z}_i is the anchor embedding, \mathbf{z}_j is a positive sample from the same trajectory as \mathbf{z}_i , and \mathbf{z}_k are negative samples drawn from other trajectories in the batch. **Utility Bridge: Brownian Encoder, Generator, and Decoder for Utility Sequences.** Similarly, the Brownian bridge model of storage utility sequences, including the utility-related Brownian encoder \mathcal{B}_e^r , generator \mathcal{B}_g^r , and decoder \mathcal{B}_d^r , has the same structure and training method using storage utility vector sequences.

Learning Surrogate Simulator with Smooth Interpolation of Brownian Next Reservoir States as Auxiliary Supervision

In GCS management and operations, numerical optimization based simulations, like physical solvers, incur substantial computational overhead and are high-cost (Witte, Hewett, and Chandra 2022; Song et al. 2023). A data-driven learnable surrogate simulator serves as a low-cost alternative to evaluate the impact of a CO₂ injection plan given the current reservoir state. Classic GCS simulators estimate the immediate storage utility and next reservoir state, given the reservoir state and injection plan. Such a simulation method can be enhanced by incorporating an auxiliary task: one-step ahead prediction of the next reservoir state’s Brownian embedding, benchmarked by Brownian interpolation.

Modeling Intuitions. Enforcing the simulator to predict the next state in a Brownian latent space, instead of the next state observed in the physical world, can ensure a smooth, gradual state transition path, rather than abrupt jumps. This is

particularly useful for environments involving physical systems. In GCS, many physical processes (e.g., CO₂ injection and subsequent movement, geochemical reactions, and geomechanical effects) involve gradual changes in pressure, saturation, and fluid composition. In contrast, abrupt transitions in state transition would not reflect the actual physics and chemistry occurring in the subsurface, thus leading to unreliable predictions. Therefore, we propose two modeling intuitions (i.e., tasks) for learning the surrogate simulator: i) a good simulator should accurately estimate storage utility given reservoir state and injection plan; ii) a good simulator should accurately infer the embedding of next reservoir state in the Brownian latent space.

Simulator Design. To bring in Brownian supervision to guide the surrogate simulator learning, we propose a different design: given the current reservoir state and injection plan, the surrogate simulator should estimate storage utility and the embedding of the next reservoir state in the Brownian latent space. In other words, we introduce a secondary auxiliary objective that requires the simulator to infer the latent embedding of the next reservoir state in a Brownian latent space. Formally, let t be the current time, \mathbf{o}_t is the current reservoir state, \mathbf{s}_t is the injection plan, $\hat{\mathbf{r}}_t$ is the predicted storage utility, \mathbf{o}_{t+1} is the next reservoir state observed in GCS, and $\hat{\mathbf{z}}_{t+1}^o$ is the Brownian latent embedding of the next reservoir state observed in GCS. The simulator simulates the storage utility and the Brownian latent next reservoir state embedding, given by:

$$\hat{\mathbf{r}}_t, \hat{\mathbf{z}}_{t+1}^o = \mathcal{S}(\mathbf{o}_t, \mathbf{s}_t), \quad (4)$$

where \mathcal{S} is a generic estimation function, such as, MLP, graph neural networks, or convolutional neural networks.

Objective Function. There are two objectives of learning the surrogate simulator: 1) minimizing the gap between the simulator-estimated storage utility and the real storage utility; 2) minimizing the gap between the simulator-estimated Brownian embedding of the next reservoir state and the actual latent embedding output by the state-related Brownian encoder of encoding the actual next reservoir state, which is given by: $\mathbf{z}_{t+1}^o = \mathcal{B}_e^o(\mathbf{o}_{t+1})$. Therefore, the ultimate optimization objective is:

$$\mathcal{L}_S = \|\hat{\mathbf{r}}_t - \mathbf{r}_t\|_2^2 + \eta \|\hat{\mathbf{z}}_{t+1}^o - \mathbf{z}_{t+1}^o\|_2^2, \quad (5)$$

where $\|\hat{\mathbf{r}}_t - \mathbf{r}_t\|_2^2$ is a storage utility estimation loss,

$\|\hat{\mathbf{z}}_{t+1}^o - \mathbf{z}_{t+1}^o\|_2^2$ is a next reservoir state embedding inference loss, and η is a hyperparameter balancing two losses.

Solving the Optimization. Reservoir state, injection plan, and storage utility of each time step is used as an instance to provide the necessary input and supervision signal for the surrogate simulator \mathcal{S} and the stochastic gradient descent is performed to update parameters of \mathcal{S} by minimizing $\mathcal{L}_{\mathcal{S}}$.

Learning Brownian Goal Conditioned Injection Plans

After learning the surrogate simulator, we aim to learn the injection planning model to decide an injection plan including water injection, CO₂ injection, and pressure control of injection wells, at each time step of the GCS lifecycle. Moreover, we want an instant decision not just to optimize instant utility but also future, long-term, total utility with respect to storage efficiency and safety.

Modeling Intuitions. Deep Brownian Bridge is a neural process that is conditioned to reach a specific final state (e.g., the upper bound of storage utility) at a fixed time, and, moreover, can interpolate the most likely trajectory towards the final state. This ability provides two benefits: i) Goal guidance: It can guide the decision model toward high-reward terminal states that are guaranteed to reach a goal (i.e., a high utility) at a target time, which is useful for goal-conditioned planning. ii) Exploration efficiency: It can enable more effective transitions of utility-relevant trajectories rather than wasting effort on aimless explorations and drifts.

Incorporating Brownian Goal Guidance into Injection Planning. Classic injection planning models are used to perceive the current reservoir state and project an injection plan, such as decisions on water injection, CO₂ injection, and pressure control. Unlike the traditional solution, we propose to leverage the utility-related deep Brownian bridge to inform the injection planning model to make decisions by conditioning on not just the current reservoir state but also a forward-looking trajectory about how to pursue a target utility from the current utility. The idea is to leverage the utility-related Brownian generator to generate a prospective latent storage utility trajectory guided by the goal of moving from the recent storage utility to the target utility, then integrate both such utility trajectory and current state as inputs of the injection planning model. Formally, let \mathbf{r}_{t-1} be the recent storage utility, \mathbf{r}^* is a predefined storage utility target (e.g., the empirical maximum or a theoretical optimum), the latent utility embedding trajectory in the utility-related Brownian space is given by: $\{\mathcal{B}_g^r(\mathbf{r}_{t-1}, \mathbf{r}^*, t') \mid t' \in [t, T]\}$, where \mathcal{B}_g^r is the utility-related Brownian generator, t' indexes each time step from t to T , and $\mathcal{B}_g^r(\mathbf{r}_{t-1}, \mathbf{r}^*, t')$ is the desired storage utility at the time t' . The injection planning model produces an injection plan by conditioning on the current reservoir state and the planning trajectory:

$$\hat{\mathbf{s}}_t = \mathcal{D}(\mathbf{o}_t, \{\mathcal{B}_g^r(\mathbf{r}_{t-1}, \mathbf{r}^*, t') \mid t' \in [t, T]\}). \quad (6)$$

where \mathcal{D} is the injection planning model, $\hat{\mathbf{s}}_t$ is an injection plan, and \mathbf{o}_t is the current reservoir state.

Objective Function. After the injection planning model generates an injection plan at the t -th time step, we exploit the surrogate simulator to estimate the storage utility of

the generated injection plan. Meanwhile, the utility-related Brownian decoder can decode the embedding of the desired utility at the t -th timestep in the latent utility embedding trajectory, into a desired storage utility. The objective is to learn the injection planning model by minimizing the gap between the estimated storage utility and the desired storage utility over time. Formally, let \mathcal{S} be a simulator, \mathbf{o}_t is current reservoir state, $\hat{\mathbf{s}}_t$ is the proposed injection plan at t , the estimated utility at the t -th timestep is: $\mathcal{S}(\mathbf{o}_t, \hat{\mathbf{s}}_t)$. The injection planning model is optimized by minimizing the gap between the estimated storage utility and the desired storage utility, given by:

$$\mathcal{L}_{\mathcal{D}} = \|\mathcal{S}(\mathbf{o}_t, \hat{\mathbf{s}}_t) - \mathcal{B}_g^r(\mathbf{r}_{t-1}, \mathbf{r}^*, t')\|_2^2, \quad (7)$$

where the desired storage utility signal $\mathcal{B}_g^r(\mathbf{r}_{t-1}, \mathbf{r}^*, t')$ at the t -th time step is interpolated from the planning trajectory by the utility Brownian generator.

Experiment

Experiment Setting

Datasets. The datasets employed in this study were constructed from the high-fidelity numerical simulator ECLIPSE 2016 (Schlumberger 2016), focusing on two distinct CO₂ storage scenarios that reflect different geological settings and operational complexities. The first scenario, denoted as **Homogeneous-WAG (H-WAG)**, represents a small-scale, homogeneous sandstone reservoir configured with a five-spot injection-production pattern. The geological model is discretized into a Cartesian grid of $60 \times 60 \times 7$ cells, with a uniform cell size and relatively thin layering to approximate near two-dimensional flow behavior. A water-alternating-gas (WAG) injection scheme was employed to control CO₂ mobility and improve storage utility. The lifecycle in this scenario consists of 120 time steps. We constructed six datasets with increasing scales, containing 50, 100, 150, 250, 350, and 450 full lifecycle trajectories, denoted as **H-WAG_1** through **H-WAG_6**. The second scenario, denoted as **Heterogeneous-Complex (H-COM)**, depicts a larger and geologically more heterogeneous system, featuring a $60 \times 60 \times 9$ grid with denser well placement and more complex stratigraphy. Initial reservoir pressures were maintained at comparable levels to the H-WAG case. The lifecycle in this scenario extends to 240 time steps. Four datasets were constructed at different scales, containing 50, 100, 200, and 300 full lifecycle trajectories, and are referred to as **H-COM_1** through **H-COM_4**. The training, validation, and testing sets are divided into 8:1:1 parts according to the number of trajectories.

Evaluation. We evaluate both surrogate modeling accuracy and injection plan optimization performance. *For surrogate simulation evaluation*, we measure the accuracy by computing the mean squared error (MSE) between the predicted storage utility and the ground-truth storage utility included in datasets. This provides a direct assessment of how well the surrogate simulator captures the system dynamics across different operational conditions. *For injection plan optimization evaluation*, we define a Storage Performance Index (SPI) based on simulator outputs to assess CO₂ storage performance of a whole lifecycle. Specifically, the SPI for each

Method	H-WAG_1	H-WAG_2	H-WAG_3	H-WAG_4	H-WAG_5	H-WAG_6	H-COM_1	H-COM_2	H-COM_3	H-COM_4
CNN	0.00224	0.00272	0.00248	0.00214	0.00163	0.00176	0.00079	0.00073	0.00075	0.00073
CNN-Ours	0.00211	0.00265	0.00235	0.00197	0.00159	0.00140	0.00075	0.00069	0.00071	0.00070
Improvement	5.87%	2.49%	5.46%	8.20%	2.65%	20.17%	4.99%	6.29%	4.54%	3.95%
AE-CNN	0.00230	0.00313	0.00317	0.00246	0.00206	0.00197	0.00134	0.00107	0.00089	0.00081
AE-CNN-Ours	0.00222	0.00290	0.00276	0.00237	0.00189	0.00185	0.00123	0.00096	0.00085	0.00076
Improvement	3.53%	7.24%	12.97%	3.97%	8.14%	6.38%	7.78%	9.72%	4.18%	6.68%
ConvLSTM	0.00193	0.00259	0.00404	0.00302	0.00243	0.00198	0.00139	0.00081	0.00077	0.00069
ConvLSTM-Ours	0.00165	0.00230	0.00374	0.00252	0.00188	0.00153	0.00104	0.00072	0.00074	0.00067
Improvement	14.59%	11.46%	7.64%	16.47%	22.40%	22.74%	25.67%	10.52%	3.77%	2.90%
GNSM	0.01208	0.00688	0.00734	0.02203	0.01550	0.01281	0.00629	0.00375	0.00312	0.00241
GNSM-Ours	0.01145	0.00559	0.00458	0.02055	0.01450	0.01161	0.00533	0.00318	0.00289	0.00227
Improvement	5.18%	18.79%	37.62%	6.73%	6.46%	9.38%	15.23%	15.26%	7.40%	5.93%
ANN	0.01085	0.01958	0.01610	0.00679	0.00822	0.00543	0.03822	0.02142	0.02064	0.01243
ANN-Ours	0.00947	0.01846	0.01508	0.00611	0.00678	0.00520	0.01416	0.00990	0.01240	0.00576
Improvement	12.70%	5.71%	6.30%	10.08%	17.46%	4.40%	62.96%	53.78%	39.94%	53.71%

Table 1: Performance Comparison Across Different Simulation Scenarios




Figure 3: Pressure changes and underground CO₂ storage distribution.

lifecycle is given by:

$$SPI = \overline{FGIR} - \overline{FGPR} + ((FGIT - FGPT)/FGIT) - \sigma_{FPR} \quad (8)$$

where \overline{FGIR} and \overline{FGPR} denote the average CO₂ injection and production rates, respectively, $FGIT$ and $FGPT$ represent the cumulative injected and produced CO₂ volumes, and σ_{FPR} is the standard deviation of the reservoir pressure. These parameters can be obtained from the storage utility. A higher SPI value indicates better CO₂ storage performance, characterized by enhanced injection capacity, minimized leakage, and stabilized reservoir pressure. We compute the mean SPI across all lifecycles in the test dataset as the final metric. We compare the SPI achieved by the injection plan when executed in both the surrogate simulator and the numerical simulator. For the baseline and our model, we repeat the experiment 5 times and report the average value.

Baseline Algorithms. *For comparisons of surrogate simulator methods*, the baseline algorithms include: (1) **ANN** (Pan et al. 2014): a standard feedforward neural network used as a surrogate to approximate CO₂ storage responses; (2) **CNN** (Wang et al. 2024): a surrogate model that predicts reservoir saturation using convolutional neural networks; (3) **AE-CNN** (Mo et al. 2019): a fully convolutional encoder-decoder network with dense blocks; (4) **ConvLSTM** (Feng et al. 2024): combining convolutional layers and ConvLSTM cells to model evolution of reservoir states; (5) **GNSM** (Tang and Durlofsky 2024): a graph neural network surrogate model. *For comparisons of injection plan optimization methods*, the baseline methods include:

- (1) **random policy (RAND)**: generating injection plan randomly and averages the results over 10 repetitions;
- (2) **SAC** (Zhang et al. 2022): a reinforcement learning method that learns a control policy through environment interactions;
- (3) **POMDP** (Corso et al. 2022): a method that optimizes strategies under uncertainty by modeling the process as a partially observable Markov decision process; and
- (4) **NSGA-II** (Liu et al. 2024b): a population-based evolutionary algorithm that performs optimization via non-dominated sorting and elite preservation.

Effectiveness of Bridge-Enhanced Surrogate

This experiment aims to evaluate how well the idea of Brownian bridge interpolated next state as simulator regularization can improve the performance of baseline methods. By incorporating bridge-enhanced objective into different surrogate simulator, we assess whether our method leads to more accurate and temporally consistent predictions over the full CO₂ injection lifecycle. Table 1 shows our method consistently improves the performance of all baseline models across different scenarios. For instance, in the H-WAG_6 scenario, the prediction error of the CNN model decreases from 0.00176 to 0.00140 (a 20.17% improvement), while in the more complex H-COM_1 scenario, the error of the ANN model is reduced by 62.96%, from 0.03822 to 0.01416. These results validate our hypothesis that modeling temporal continuity is critical for accurate surrogate simulation.

Dataset	H-WAG_1	H-WAG_2	H-WAG_3	H-WAG_4	H-WAG_5	H-WAG_6	H-COM_1	H-COM_2	H-COM_3	H-COM_4
RAND	0.01224	-0.87519	-0.32341	0.20419	-0.18921	0.04251	0.08970	-1.05086	0.32720	-0.20124
SAC	0.74475	0.89808	0.86930	0.88703	<u>1.20986</u>	1.08951	0.88418	1.00218	0.81055	<u>1.28069</u>
POMDP	<u>0.85544</u>	<u>1.07623</u>	0.97523	<u>1.04814</u>	1.15012	<u>1.18885</u>	0.97278	<u>1.29326</u>	<u>0.85680</u>	1.26131
NSGA-II	0.72146	0.92391	<u>1.15187</u>	0.92426	1.13187	0.31259	<u>1.20690</u>	0.90107	0.81491	0.97578
Ours	1.35140	1.35223	1.36021	1.36291	1.35927	1.35646	1.33317	1.39324	1.39760	1.39605

Table 2: Comparison of Decision Strategies Across Scenarios

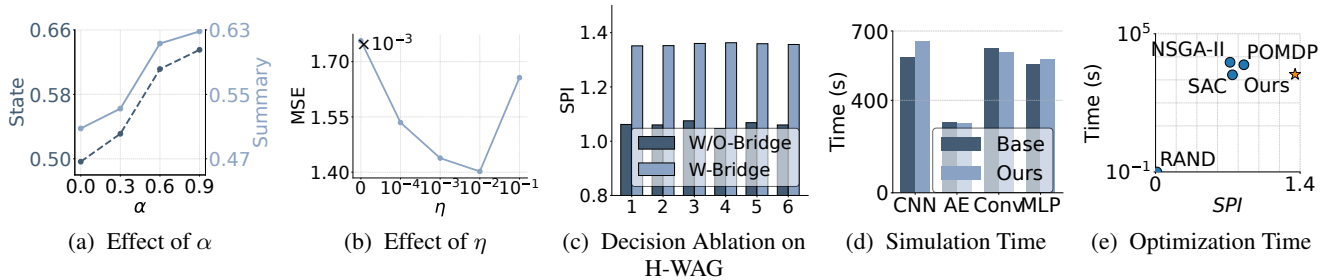


Figure 4: Investigation of Proposed Method

Effectiveness of Injection Plan Optimization

This experiment evaluates the effectiveness of our injection planning model in improving CO₂ storage performance. Table 2 shows our method consistently outperforms all baseline approaches. For instance, in the H-COM_3 case, our injection plan achieves an SPI of 1.3976, compared to 0.8568 by the strongest baseline (POMDP). To further verify the physical plausibility and effectiveness of our method, we selected the strongest baseline (POMDP) and our method to generate multiple injection plans to control the whole lifecycle, which were then evaluated using the ECLIPSE 2016 numerical simulator on the H-COM setting. Figure 3a shows the cumulative distribution function (CDF) of CO₂ saturation (SGAS) reveals that our method consistently achieves higher saturation levels across multiple time steps. This indicates that our strategy enables more effective CO₂ storage. Figures 3b and 3c depict the pressure evolution throughout the lifecycle. Compared to the baseline, our method results in more regular and stable pressure dynamics. This reflects greater control and operational safety, as sudden pressure fluctuations are known to increase the risk of reservoir damage or leakage. Figures 3d and ?? provide a 3D visualization of the underground CO₂ distribution at the end of the lifecycle. Due to the lower density of CO₂, injected CO₂ tends to accumulate in the upper layers. The resulting distribution from our strategy aligns well with this physical expectation, suggesting that our decisions are geophysically consistent. Furthermore, the CO₂ plume generated by our method is more uniformly distributed across the reservoir, indicating more balanced injection and less local saturation bias.

Component Analysis and Ablation Study

Figures 4a and 4b analyze the sensitivity of trajectory construction and surrogate learning. Moderate to high values of the noise parameter α lead to higher similarity between gen-

erated and ground-truth trajectories, while the auxiliary loss weight η exhibits a U-shaped effect on H-WAG_6, with optimal performance around $\eta = 10^{-3}$. Figure 4c shows that bridge-guided trajectory planning consistently improves SPI across all H-WAG scenarios, highlighting the importance of trajectory-aware decision guidance. As shown in Figures 4d and 4e, our method maintains competitive training and optimization efficiency despite the auxiliary task, benefiting from a compact latent representation. Overall, the results demonstrate that Brownian bridge-based trajectory modeling improves both prediction accuracy and decision quality without incurring significant computational overhead.

Conclusion Remarks and Limitations

We present a Brownian Bridge-augmented framework for CO₂ injection strategy optimization in GCS, addressing temporal continuity and goal-conditioned planning. Our method integrates Brownian bridge representations into surrogate simulation and injection planning by (i) learning smooth latent trajectories of reservoir states and utilities, (ii) improving simulation fidelity through Brownian interpolation, and (iii) guiding adaptive planning via utility-conditioned trajectories. Experiments across multiple GCS datasets show improved simulation accuracy and injection plan quality with low computational cost. A remaining limitation is the multi-stage design, which may accumulate errors; future work will explore tighter integration.

Acknowledgments

Dr. Yanjie Fu is supported by the National Science Foundation (NSF) via the grant numbers: 2426340, 2416727, 2421865, 2421803.

References

- Chen, Z.; Huan, G.; and Ma, Y. 2006. *Computational methods for multiphase flows in porous media*. SIAM.
- Corso, A.; Wang, Y.; Zechner, M.; Caers, J. K.; and Kochenderfer, M. J. 2022. A POMDP Model for Safe Geological Carbon Sequestration. *ArXiv*, abs/2212.00669.
- Feng, Z.; Tariq, Z.; Shen, X.; Yan, B.; Tang, X.; and Zhang, F. 2024. An encoder-decoder ConvLSTM surrogate model for simulating geological CO₂ sequestration with dynamic well controls. *Gas Science and Engineering*, 125: 205314.
- Ismail, I.; and Gaganis, V. 2023. Carbon capture, utilization, and storage in saline aquifers: subsurface policies, development plans, well control strategies and optimization approaches—a review. *Clean Technologies*, 5(2): 609–637.
- Juanes, R.; Spiteri, E.; Orr Jr, F.; and Blunt, M. 2006. Impact of relative permeability hysteresis on geological CO₂ storage. *Water resources research*, 42(12).
- Liu, G.; Pu, L.; Zhao, H.; Chen, Z.; and Li, G. 2024a. Multi-objective optimization of CO₂ ejector by combined significant variables recognition, ANN surrogate model and multi-objective genetic algorithm. *Energy*, 295: 131010.
- Liu, G.; Pu, L.; Zhao, H.; Chen, Z.; and Li, G. 2024b. Multi-objective optimization of CO₂ ejector by combined significant variables recognition, ANN surrogate model and multi-objective genetic algorithm. *Energy*, 295: 131010.
- Liu, S.; Agarwal, R.; Sun, B.; Wang, B.; Li, H.; Xu, J.; and Fu, G. 2021. Numerical simulation and optimization of injection rates and wells placement for carbon dioxide enhanced gas recovery using a genetic algorithm. *Journal of Cleaner Production*, 280: 124512.
- Mao, J.; and Ghahfarokhi, A. J. 2024. A review of intelligent decision-making strategy for geological CO₂ storage: Insights from reservoir engineering. *Geoenergy Science and Engineering*, 240: 212951.
- Mo, S.; Zhu, Y.; Zabarar, N.; Shi, X.; and Wu, J. 2019. Deep convolutional encoder-decoder networks for uncertainty quantification of dynamic multiphase flow in heterogeneous media. *Water Resources Research*, 55(1): 703–728.
- Pan, I.; Babaei, M.; Korre, A.; and Durucan, S. 2014. Artificial neural network based surrogate modelling for multi-objective optimisation of geological CO₂ storage operations. *Energy Procedia*, 63: 3483–3491.
- Revuz, D.; and Yor, M. 2013. *Continuous martingales and Brownian motion*, volume 293. Springer Science & Business Media.
- Schlumberger. 2016. *ECLIPSE Reservoir Simulator*. <https://www.software.slb.com/products/eclipse>.
- Song, Y.; Jun, S.; Na, Y.; Kim, K.; Jang, Y.; and Wang, J. 2023. Geomechanical challenges during geological CO₂ storage: A review. *Chemical Engineering Journal*, 456: 140968.
- Sun, Z.; Xu, J.; Espinoza, D. N.; and Balhoff, M. T. 2021. Optimization of subsurface CO₂ injection based on neural network surrogate modeling. *Computational Geosciences*, 25(6): 1887–1898.
- Tang, H.; and Durlofsky, L. J. 2024. Graph network surrogate model for subsurface flow optimization. *Journal of Computational Physics*, 512: 113132.
- Wang, J.; Lin, D.; and Li, W. 2023. Dialogue Planning via Brownian Bridge Stochastic Process for Goal-directed Proactive Dialogue. In *Findings of the Association for Computational Linguistics: ACL 2023*, 370–387. Toronto, Canada: Association for Computational Linguistics.
- Wang, R. E.; Durmus, E.; Goodman, N. D.; and Hashimoto, T. 2022. Language modeling via stochastic processes. *ArXiv*, abs/2203.11370.
- Wang, S.; Xiang, J.; Wang, X.; Feng, Q.; Yang, Y.; Cao, X.; and Hou, L. 2024. A deep learning based surrogate model for reservoir dynamic performance prediction. *Geoenergy Science and Engineering*, 233: 212516.
- Witte, P. A.; Hewett, R.; and Chandra, R. 2022. Industry-scale CO₂ flow simulations with model-parallel fourier neural operators. In *NeurIPS 2022 workshop tackling climate change with machine learning*.
- Zhang, K.; Wang, Z.; Chen, G.; Zhang, L.; Yang, Y.; Yao, C.; Wang, J.; and Yao, J. 2022. Training effective deep reinforcement learning agents for real-time life-cycle production optimization. *Journal of Petroleum Science and Engineering*, 208: 109766.
- Zhang, Z.; and Agarwal, R. 2013. Numerical simulation and optimization of CO₂ sequestration in saline aquifers. *Computers & Fluids*, 80: 79–87.