

Multi-Objective Bayesian Optimization Framework for CO_2 Injection Strategy Design: A Sleipner-Inspired Study

Muhammad Amir Saeed^{1,2,3*} Jo Eidsvik^{2,3†} Antonio Candelieri^{1‡}

May 28, 2026

¹ Department of Economics, Management and Statistics, University of Milano-Bicocca, Milan, Italy.

² Department of Mathematical Sciences, Norwegian University of Science and Technology (NTNU), Trondheim, Norway.

³ Center for Geophysical Forecasting (CGF), Norwegian University of Science and Technology (NTNU), Trondheim, Norway.

Abstract

This study develops a multi-objective Bayesian optimization framework for the design of geological CO_2 injection strategies. The problem is formulated over a mixed decision space that includes well configuration, injection rate, and staged allocation of injection across reservoir layers. Evaluating candidate strategies requires repeated simulation of subsurface migration processes, which motivates the use of data-efficient optimization techniques.

The optimization is posed with three competing objectives: maximizing retained CO_2 storage, minimizing upward migration through a leakage proxy, and minimizing total project cost. To handle the mixed categorical–continuous decision variables and non-smooth response behavior, a multinomial logit surrogate model is employed. Within this framework, several acquisition strategies are investigated, including expected hypervolume improvement, scalarized upper confidence bounds, Thompson sampling, and expected preference improvement. Optimization performance is assessed using hypervolume-based metrics, convergence behavior, and characteristics of the resulting Pareto sets.

Testing on the Sleipner in the North Sea, the proposed framework is useful for trade-offs involved in designing CO_2 storage. Here, we rely on using a simulator for fluid migration in multi-layer storage system running through an integrated R-Julia workflow. Results show that the suggested method finds a variety of high-quality trade-off solutions within a small evaluation budget. The Scalarized upper confidence bounds method performs the best of all the tested strategies, creating the largest and most evenly spaced Pareto front. The analysis shows that vertical injection allocation is the most important factor in controlling how well storage works and how it behaves when it is full.

Keywords: CO_2 storage; Bayesian optimization; Multi-objective optimization; Pareto front; Hypervolume; Multinomial logit surrogate; Sleipner.

*Corresponding author. Email: m.saeed@campus.unimib.it. ORCID: 0000-0003-1650-8194.

†Corresponding author. Email: jo.eidsvik@ntnu.no. ORCID: 0000-0002-9757-9252.

‡Corresponding author. Email: antonio.candelieri@unimib.it. ORCID: 0000-0003-1431-576X.

1 Introduction

High levels of CO_2 in the atmosphere is a major cause of anthropogenic climate change. Geological carbon capture and storage (CCS) has become a promising approach for cutting CO_2 emissions. In CCS operations, CO_2 is taken from industrial activities and pumped into deep geological formations like saline aquifers, depleted hydrocarbon reservoirs, or coal seams. Injected CO_2 can stay in storage formations for a long time, and when trapped by capillary forces, it will dissolve in the formation brine, or slowly change into stable minerals through geochemical reactions over time (Ringrose, 2020; Alqahtani et al., 2023).

Despite the large impact that CCS can have on mitigating climate change, there are also risks involved. The main risk is associated with migration from the storage units which could eventually lead to leakage of CO_2 to the surface and into the atmosphere. This risk is mainly driven by the pressure build-up in the storage units which is again influenced by the fluid flow in the subsurface formations. The pressure build-up is partly governed by the decision maker via injection control rates in one or more wells. It is also influenced by geological formation geometries, cap rock properties, faulting, and other factors. These factors are important to consider when designing a large-scale CCS project. Moreover, the operational feasibility and costs will of course vary with different injection strategies (injection rates, locations, number of wells, injection layers, etc.).

Numerical reservoir simulation have been used for the assessment and design of CO_2 storage systems. Relying on foundational differential equations and decades of knowledge from the oil and gas industry (Lie, 2019), they can simulate the CO_2 plume over time at high levels of spatial detail. They are however quite costly, and in complex optimization systems one must often rely on simpler simulators that capture the main mechanisms. One such model is that of invasion percolation (IP) which relies on vertical migration and buoyancy forces (Carruthers, 2003; Callioli Santi et al., 2025).

Bayesian optimization (BO) has recently gained attention as a practical approach for optimizing expensive black-box models when only a limited number of evaluations are feasible (Garnett, 2023). BO uses statistical surrogate models with uncertainty quantification to guide the selection of new simulation runs, balancing exploration of uncertain regions with exploitation of promising decision designs. In CCS injection planning, the problem naturally leads to a multi-objective formulation, since objectives such as maximizing stored CO_2 , minimizing leakage risk, and controlling operational cost are often in conflict. Rather than searching for a single optimal solution, the goal is to identify a set of optimal solutions that illustrate the trade-offs between competing objectives. Multi-objective Bayesian optimization (MOBO) incorporates concepts of Pareto dominance directly into the acquisition process. A commonly used approach is then to select new decision designs for a simulator based on their expected contribution to the Pareto front. Recent work in MOBO has proposed numerous variants of this type of optimization (Candelieri et al., 2024). Despite these advances, many existing methods primarily target continuous design spaces. In practical CCS applications, however, the decision space often includes discrete or categorical variables, such as the number of wells, selection of injection layers, staged injection schedules, etc. These types of variables are not always easily handled by standard models.

In this paper we develop a three-objective MOBO framework for CO_2 injection planning. The three objectives considered are (i) maximizing the amount of CO_2 retained within the target storage layer, (ii) minimizing cross-layer migration, represented through a leakage proxy associated with vertical CO_2 movement, and (iii) minimizing overall project cost. The decision space includes both continuous variables, such as injection rate, and categorical variables, such as well configuration and injection-layer allocation. To handle categorical decisions, we employ a multinomial logit-based surrogate model and examine several acquisition strategies. Optimization performance is also evaluated using multiple criteria. The proposed method is tested on a model inspired by the

Sleipner CO_2 storage project. The IP-based simulator facilitate MOBO.

The remainder of the paper is organized as follows. Section 2 discusses the literature and related work. Section 3 introduces the case study and formulates the multi-objective optimization problem. Section 4 describes the Bayesian optimization framework, including the surrogate model and acquisition strategies used in this study. Section 5 presents and discusses the optimization results. Finally, Section 6 summarizes the main findings and outlines directions for future work.

2 Related Work

2.1 Geological CO_2 Storage and reservoir simulation

Geological CCS has attracted growing research interest as a large-scale mitigation pathway for industrial emissions (Ringrose, 2020). Successful underground storage entails that the injected CO_2 builds up over time to get trapped in structural accumulation beneath low-permeability caprocks. In layered formations such as the Utsira Sand at Sleipner, buoyancy forces drive CO_2 upward through inter-layer connections, leading to a characteristic stacked plume structure that is sensitive to both permeability contrasts and the integrity of thin shale barriers. Studies on Sleipner specifically have shown that vertical-equilibrium (VE) models can be sufficient to capture CO_2 migration in high-permeability layers where capillary transition zones are thin (Bandilla et al., 2014). These physical complexities make CO_2 storage an inherently multi-physics problem in which containment performance, plume migration, and pressure evolution must all be considered together when designing injection strategies.

High-fidelity reservoir simulators are the standard tool for assessing CO_2 storage behavior, as they can capture multiphase flow, buoyancy-driven transport, and pressure dynamics at field scale. However, repeated simulation of realistic storage formations is computationally expensive, which limits the number of injection scenarios that can be run on the computer. VE and hybrid-VE approximations have therefore become widely used in optimization and uncertainty workflows, providing large reductions in computational cost while retaining the essential physics of buoyancy-driven spreading. The Sleipner 2019 benchmark dataset, released by Equinor as an openly accessible reference model with plume outlines, well logs, and velocity maps, has become a standard testbed for methods that combine simulation and monitoring data (Equinor / CO2DataShare, 2019).

To further reduce the cost of repeated forward evaluations, machine learning-based models have also been widely adopted in CCS workflows. A rigorous machine learning-assisted workflow for uncertainty and sensitivity analysis of CO_2 storage prediction in deep saline aquifers combines Latin hypercube sampling for training data generation with data-driven models that map the non-linear relationship between input parameters and storage outcomes, enabling uncertainty propagation via Monte Carlo simulation and global sensitivity analysis through Sobol indices (Alqahtani et al., 2023). Beyond statistical approaches, neural network architectures—including convolutional networks, U-Net variants, recurrent networks, and Fourier neural operators—have been applied to predict the spatial and temporal evolution of pressure and CO_2 saturation fields in heterogeneous reservoirs (Liu et al., 2025b). Even though such methods require much training, the developments have made it feasible to build computationally tractable proxies that may support optimization at field scale. Nevertheless, most existing studies focus on continuous input spaces and single-objective formulations.

2.2 Bayesian Optimization for CO_2 Injection Design

BO provides a principled framework for optimization. By maintaining a probabilistic surrogate model—typically a Gaussian process (GP)—and selecting new candidate points through an acquisition function that balances exploration with exploitation, BO can identify good solutions with far fewer simulator calls than gradient-free or evolutionary methods. An early demonstration of BO for geological carbon sequestration coupled a high-fidelity compositional reservoir simulator with a GP-based Bayesian optimization accelerator for parallel optimization of injection schedules, achieving a competitive objective value with over 60% fewer forward model evaluations compared to a genetic algorithm and a covariance matrix adaptation evolution strategy (Lu et al., 2022). More recent work has extended BO to optimize CO_2 injection rates alongside well placement, demonstrating that relaxing fixed group-control constraints on injection rates allows the optimization to explore alternative solutions and achieve improved storage performance over 40-year injection horizons (Fotias et al., 2024).

In practice, CCS injection design involves competing objectives: maximizing stored volumes, minimizing leakage risk, and controlling cost that cannot be reduced to a single criterion without losing important trade-off information. Multi-objective optimization approaches aim to identify the Pareto front: the set of solutions for which no objective can be improved without degrading another. Evolutionary algorithms such as NSGA-II and MOPSO have been widely used in this context. A hybrid framework combining long-short-term memory proxy models with NSGA-II has been applied to simultaneously maximize oil recovery and CO_2 storage while minimizing fault slip displacement, demonstrating that geomechanical risk can be incorporated as a third objective in a single optimization workflow (Liu et al., 2023). Similarly, deep learning-assisted approaches have been applied to jointly maximize CO_2 storage and minimize geomechanical risks by adjusting injection schedules within bounded constraints (Liu et al., 2025b), with proxy models being essential for making such coupled workflows computationally tractable.

MOBO extends BO to the Pareto-aware setting by incorporating concepts of Pareto dominance directly into the acquisition process (Candelieri et al., 2024). The most widely used acquisition function in this setting is Expected Hypervolume Improvement (EHVI), which selects new evaluations based on their expected contribution to the volume of objective space dominated by the current Pareto front. The differentiable q EHVI formulation extends EHVI to parallel and constrained evaluation settings, enabling gradient-based acquisition optimization via automatic differentiation and demonstrating computational tractability across many practical scenarios (Daulton et al., 2020). Further extensions address noisy observations: the NEHVI acquisition function applies a Bayesian treatment to the EHVI criterion by integrating over uncertainty in the Pareto frontier itself, and its parallel variant q NEHVI reduces computational complexity from exponential to polynomial with respect to batch size while remaining one-step Bayes-optimal (Daulton et al., 2021). Recent work has also addressed numerical issues in gradient-based optimization of EHVI, showing that log-space reformulations can resolve vanishing gradient problems and significantly improve performance across a wide range of benchmark problems (Ament et al., 2023). Thompson sampling (TS) has been proposed as an alternative acquisition strategy for MOBO, combining the accuracy of evolutionary multi-objective search with the sample efficiency of surrogate-based optimization by solving a cheap multi-objective problem on GP posterior samples (Daulton et al., 2020).

2.3 Handling Categorical and Mixed Decision Spaces

A recurring challenge in practical CCS optimization is that the decision space contains both continuous variables (such as injection rate and well depth) and discrete or categorical variables (such as the

selection of injection layers, well configurations, or staged schedules). Standard GP-based MOBO methods assume a continuous input domain and do not directly accommodate non-ordinal categorical choices. Several approaches have been proposed to address this. Bartoli et al. (2025) developed a multi-objective surrogate method developed for continuous and categorical design variables using mixtures of experts, allowing the construction of accurate Pareto fronts with a few function evaluations. A mixed-categorical correlation kernel for Gaussian processes has also been developed to preserve the structure of categorical inputs within GP surrogate models (Saves et al., 2023). Piecewise affine surrogates have been proposed for linearly constrained mixed-variable problems, relying on mixed-integer linear programming for acquisition optimization (Zhu and Bemporad, 2025).

An alternative family of models—discrete-choice surrogates, including multinomial logit (MNL) models provides a natural framework for modeling non-ordinal categorical alternatives by expressing category probabilities as functions of covariates through a softmax transformation (Candelieri et al., 2024). High-dimensional extensions of the MNL model have been developed to handle large choice sets with many categorical explanatory variables by clustering parameters across outcome and explanatory categories using a Bayesian Dirichlet process prior, reducing the effective number of model parameters and providing interpretable structure (Nibbering, 2024). Recent work has further demonstrated the use of multinomial logit surrogates within Bayesian optimization frameworks for categorical and mixed-variable optimization problems (Saeed and Candelieri, 2026). The incorporation of such discrete-choice surrogates into sequential MOBO loops has received limited attention in the CCS literature, and their relative performance compared to GP and tree-based alternatives has not been systematically evaluated on realistic subsurface benchmarks.

2.4 Economic Objectives in CO_2 Storage Planning

In addition to geophysical containment objectives, economic feasibility is a central concern in CO_2 storage projects. Project costs span capital expenditure for drilling and infrastructure, operating expenses for compression and monitoring, and potential penalty exposure if injection targets contracted with emitters are not met (Alqahtani et al., 2023).

Integrated multi-objective optimization frameworks for offshore CCS system design have demonstrated that balancing storage capacity, safety, and economic feasibility yields Pareto-optimal solutions across key transport and injection parameters (Loke et al., 2025). The inclusion of explicit cost and contractual penalty terms in a multi-objective framework forces decision-makers to confront trade-offs between storage performance, containment security, and project economics in a transparent way. However, few existing MOBO studies for CCS injection planning directly include economic objectives alongside physical storage metrics within a unified Pareto framework. Liu et al. (2025a) examined a related class of multi-objective problems for CO_2 storage in saline aquifers, but focused primarily on physical objectives.

2.5 Research Gaps and Contributions

The body of work reviewed above reveals strong methodological foundations in VE-based simulation, BO for injection scheduling, and MOBO for competing physical objectives. Nevertheless, several gaps remain. First, existing MOBO applications in the CCS domain have largely addressed two-objective formulations, with limited attention to three-objective settings that simultaneously capture containment, migration, and project economics. Second, standard BO frameworks are not well suited to the mixed categorical–continuous decision spaces that arise naturally in injection design, and discrete-choice surrogates such as the MNL model have not been systematically evaluated in this context. Third, empirical comparisons of acquisition strategies—including EHVI,

TS, scalarized upper confidence bounds (ScalUCB), and preference-based alternatives—on realistic CCS benchmarks are scarce.

The present work addresses all these three gaps by developing a three-objective MOBO framework with a MNL surrogate for mixed decision spaces. We evaluate multiple acquisition strategies, and benchmarking performance on a simulation-based model inspired by the Sleipner CO_2 storage system.

3 Problem Formulation

3.1 Sleipner Case

3.1.1 Background

Sleipner was the first offshore CCS project. Since 1996, the Sleipner field has captured and stored around one million tons CO_2 each year. The site is situated in the North Sea, off the Western coast of Norway (Figure 1a). It was initially prompted by the Norwegian offshore CO_2 tax, encouraging CCS.

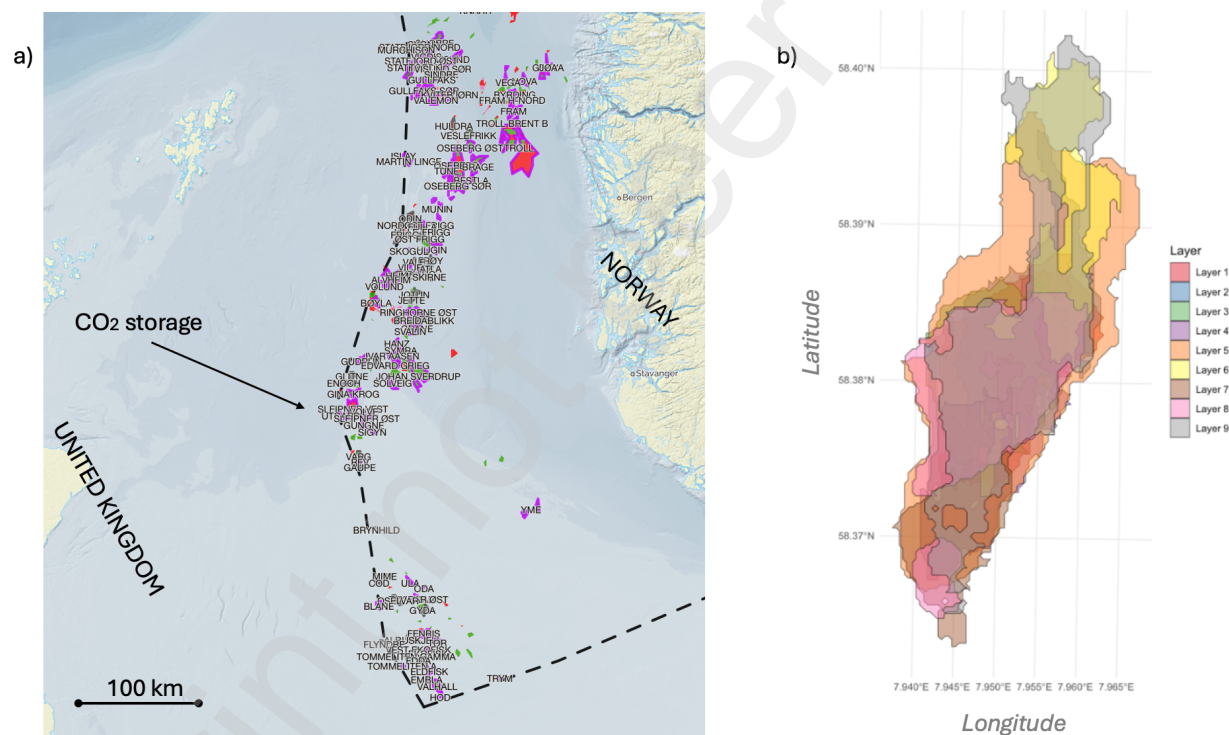


Figure 1: a) The Sleipner field in the North Sea where CO_2 has been injected since 1996. b) Polygon representation of the plumes in the nine storage layers identified at the Sleipner site.

The Sleipner 2019 Benchmark Model (Equinor / CO2DataShare, 2019) provides a widely used reference dataset for CCS research. The dataset includes grid geometry, petrophysical properties, and well information describing the storage formation. It was developed to support research and development activities related to CO_2 storage, and it builds upon the earlier 2011 Sleipner Benchmark Model, which has been used by more than 40 research institutions and has served as the basis for numerous studies.

At Sleipner, the storage formation consists of nine sandstone layers separated by thin shale and mudstone barriers. Time-lapse seismic monitoring has shown that injected CO_2 has migrated into these layers, forming nine distinct plume accumulations (Figure 1b). While the earlier 2011 benchmark model focused only on the uppermost storage layer (Layer 9), the 2019 benchmark model represents all nine layers and therefore provides a more complete three-dimensional description of the storage system.

At Sleipner, the single injection well has historically injected between 0.8 and 1 million tonnes of CO_2 per year. Although its location is known, in the optimization experiments we allow variation in injection rate and also consider alternative well configurations. In addition, scenarios with multiple wells are examined, where both the injection rates and well placements can vary.

3.1.2 Objective

We use Sleipner as an inspiration to use, develop and study properties of multi-objective optimization problems. The optimization problem in geological CO_2 storage involves balancing several competing objectives. In this study we consider three main objectives:

- *Maximize* the net volume of CO_2 retained within the storage formation,
- *Minimize* the cumulative CO_2 volume migrating out of the storage layers (leakage proxy),
- *Minimize* the total project cost associated with injection and monitoring.

Evaluating these objectives requires simulation of CO_2 migration for different choices of injection strategies. We conduct MOBO by playing with a variety of decision variables.

3.2 Decision Variables and Simulator

The reservoir simulator is at the core of evaluating the trade-offs among project duration, cost, and storage performance. This simulator relies on several inputs. Decision variables refer to input operational parameters that are chosen by the operator and can be changed during optimization. Fixed variables describe elements that a decision maker cannot control, e.g. the geological structure, cap rock properties, etc., which are assumed known in this work.

3.2.1 Decision Variables

- **Number of injection wells:** The default Sleipner configuration consists of a single injection well. In this study we also consider scenarios with three wells at known locations. Increasing the number of wells may improve storage efficiency and reduce pressure buildup per well, but it also increases operational costs.
- **Annual CO_2 injection rate:** The injection rate is varied between 0.8 and 1.0 million metric tons per year, consistent with historical Sleipner operations. This range allows investigation of how injection intensity affects plume development, pressure buildup, and storage efficiency.
- **Injection schedule across layers:** The total injection period is fixed to 14 years, consistent with the Sleipner reference scenario. However, the allocation of injection across different reservoir layers during this period is treated as a decision variable. Injection is implemented through a staged schedule in which injection initially targets lower reservoir layers and may later shift toward intermediate or upper layers. This strategy enables investigation of vertical plume migration driven by pressure buildup and caprock constraints.

The Sleipner benchmark model provides estimates of the reference depth and geometry of each layer. These properties vary substantially across the nine layers, resulting in different storage capacities and plume distributions. However, they cannot be controlled in the optimization process. In this work we treat the depths and geometry of the nine stratigraphic layers as known. Capillary pressure threshold is also assumed to be known. This guarantees that discrepancies in optimization outcomes stem solely from alterations in operational decisions. The simulator must further make assumptions about the residual oil trapping in each layer, which implicitly makes assumptions about the porosity in each sand unit. We must further employ closed lateral boundary conditions.

3.2.2 Simulator

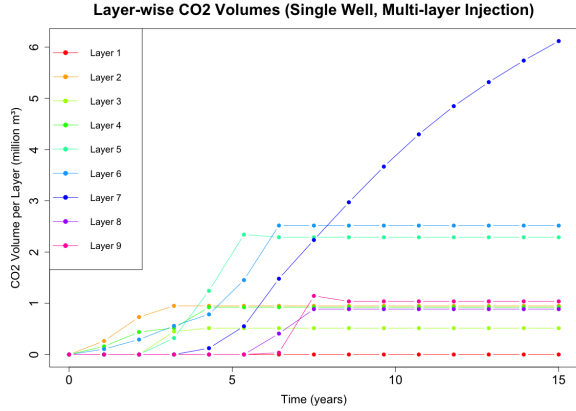
In the oil and gas industry, the reservoir simulations generally need the resolution of coupled differential equations that depict pressure and fluid saturation with time, which can be computationally intensive. See Lie (2019) for a comprehensive discussion of these approaches. Various modeling methodologies may be employed based on the degree of physical detail incorporated.

Callioli Santi et al. (2025) discuss potential scenarios of vertical CO_2 migration between the nine layers at Sleipner. We adopt this formulation, focusing on pressure- or height-induced vertical migration between layers. As layers are filled up from below, the pressure increase and vertical migration to the next layer occur when the pressure exceeds the capillary threshold. A simplified fluid-flow simulator with embedded vertical migration, volume filling, and layer connectivity through graph-based calculations is used (Svee, 2026). This simulator is a multi-layer extension of the gravity-based flow outlined in (Andersen, 2025). The approach significantly reduces computational cost compared to full reservoir simulations while still capturing the key mechanisms relevant for optimization. The computational efficiency of the simulator of Svee (2026) enables repeated evaluations within a MOBO loop.

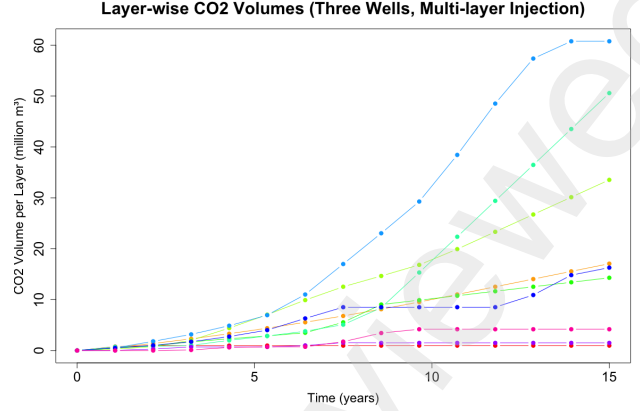
Two representative CO_2 injection scenarios are illustrated next. These scenarios run the simulator on different inputs, and they are designed to capture both single-well operation with vertical injection distribution and three-well operation with potential plume interaction. Together, they provide complementary perspectives on how operational configurations influence storage behavior and system response.

Figure 2a shows the setting with injection in multiple layers. The injection runs over 14 years at a time-varying rate that peaks at 1.0 Mt yr^{-1} before gradually declining. The injected mass is converted into volumetric rates using a constant CO_2 density of 425 kg m^{-3} . The well disperses CO_2 vertically across multiple strata at a specified lateral location, with 50% introduced into layer 2, 30% into layer 4, and 20% into layer 6. This multi-layer injection technique could preserve caprock integrity by preventing injection in the upper layers, hence facilitating the examination of vertical plume dispersion and storage distribution. Simulator results show that the CO_2 volumes increase over time, at different rates in the various layers, partly caused by vertical migration but also by injection rates.

Figure 2b shows the three-well CO_2 injection scenario was simulated under closed lateral boundary conditions using a custom reservoir configuration with layer-specific leakage properties. Here, the wells are arranged in a clustered triangular pattern near the center of the domain, with one well located at the grid center and the other two moved to the east and north, respectively. Each well operated with injection schedules over a 14-year period, with injection rates ranging from 0.8 to 1.0 Mt yr^{-1} . The injection CO_2 was distributed across multiple reservoir layers using predefined layer-specific fractions. This configuration allows evaluation of plume interaction and pressure interference arising from closely spaced injection wells and their impact on CO_2 storage behavior.



(a) Reference case at Sleipner



(b) Customized reference case

Figure 2: CO_2 flow simulations for the reference and customized cases using the Julia-based simulator (Svee, 2026).

3.3 Objective Functions

Designing geological CO_2 storage systems demands the identification of injection and operational parameters that guarantee long-term containment while preserving technological efficiency and economic viability across prolonged operational durations. These goals are inherently conflicting: increasing injected volumes may improve storage utilization but simultaneously elevate leakage risk and overall cost. The resulting design problem, therefore, requires explicit treatment of trade-offs between storage performance, containment integrity, and economic efficiency.

In this study, we formulate the injection design task as a three-objective optimization problem over a mixed (categorical-continuous) decision vector $X \in \mathcal{X}$. All objectives are evaluated over a fixed simulation horizon of length T , and project duration is not treated as an optimization variable. This allows the analysis to focus on physical containment behavior and cost performance under a prescribed operational period.

The multi-objective optimization problem is written as

$$\max_{X \in \mathcal{X}} F(X) = (f_1(X), -f_2(X), -f_3(X)), \quad (1)$$

where f_1 is maximized, while f_2 and f_3 are minimized. The resulting Pareto fronts are analyzed to identify trade-offs between retained storage, leakage risk, and economic cost.

3.3.1 Objective 1: maximize retained CO_2 storage

The first objective is to maximize the amount of retained CO_2 within the storage system over the simulation period. For a given decision vector X , which specifies the injection strategy through well configuration, injection rate, and layer allocation over time, the simulator returns the total retained CO_2 volume in the storage formation at each time step. We denote this quantity by $V_{\text{tot}}(t)$.

Instead of focusing solely on the maintained volume at the final simulation time, we establish the target as the cumulative retained volume during the whole time horizon T . This favors injection strategies that achieve storage early and maintain retention throughout the simulation period. The

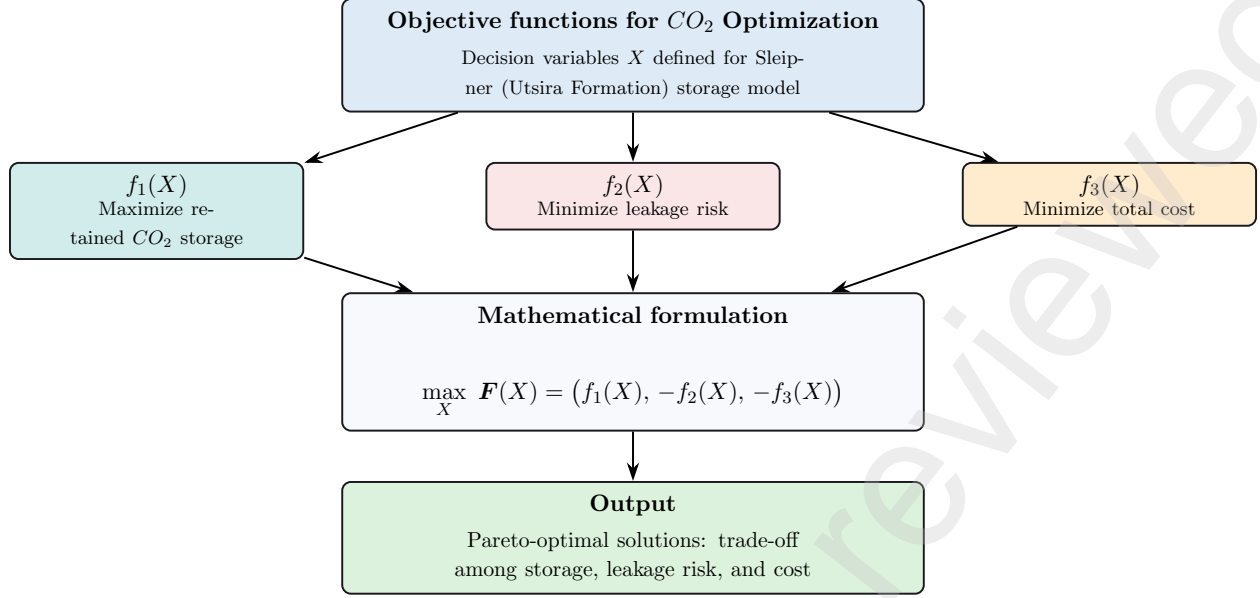


Figure 3: Three-objective optimization framework for CO_2 injection design.

first objective is therefore written as

$$f_1(X) = \int_0^T V_{\text{tot}}(t) dt \approx \sum_{k=1}^{T/\Delta t} V_{\text{tot}}(t_k) \Delta t, \quad (2)$$

where Δt is the simulation time step, taken as one year in this study.

In the implementation, $V_{\text{tot}}(t_k)$ is obtained directly from the simulator output at each time step, and $f_1(X)$ is evaluated as the discrete sum of these retained-volume values over the 14-year simulation horizon.

3.3.2 Objective 2: minimize leakage proxy

Direct quantification of physical leakage through faults, fractures, or legacy wells would require additional site-specific data and more detailed modeling assumptions. For the present optimization study, we therefore adopt a practical leakage proxy based on a global volume balance.

For a given design X , let $V_{\text{inj}}(X)$ denote the total injected CO_2 volume over the simulation horizon T . This is computed by summing the injected volumes across all wells, layers, and time steps:

$$V_{\text{inj}}(X) = \int_0^T Q(t) dt \approx \sum_{k=1}^{T/\Delta t} Q(t_k) \Delta t, \quad (3)$$

where $Q(t)$ is the total injection rate over all active wells and layers, and $\Delta t = 1$ year in the numerical implementation.

The leakage proxy is then defined as the positive difference between the cumulative injected volume and the retained CO_2 volume at the final simulation time:

$$f_2(X) = \max(0, V_{\text{inj}}(X) - V_{\text{tot}}(T)). \quad (4)$$

This quantity is zero if all injected CO_2 remains in the system at the end of the simulation, and increases when part of the injected volume is no longer retained. Although this proxy does not identify the exact migration pathway, it provides a simple and consistent measure of containment loss that is directly available from the simulator outputs and therefore suitable for comparative optimization.

3.3.3 Objective 3: minimize total project cost

The third objective represents the total project cost over the fixed simulation horizon T . The cost model is intentionally simple and follows the numerical implementation used in the optimization. It combines fixed project cost, well-dependent cost, variable cost proportional to total injected CO_2 volume, and monitoring cost over the simulation period.

The total cost is defined as

$$f_3(X) = C_{\text{fixed}} + N_w C_{\text{well}} + c_m V_{\text{inj}}(X) + c_{\text{mon}} T, \quad (5)$$

where C_{fixed} denotes fixed project overhead, N_w is the number of injection wells, and C_{well} is the cost per well. The term c_m is the marginal cost per unit of injected CO_2 , defined as

$$c_m = c_{\text{cap}} + c_{\text{trans}} + c_{\text{inj}},$$

where c_{cap} , c_{trans} , and c_{inj} represent capture, transport, and injection costs, respectively. The final term $c_{\text{mon}}T$ accounts for monitoring cost accumulated over the full simulation period.

This formulation implies that designs with more wells or higher total injection volume lead to higher project cost. The cost model is not intended as a detailed economic assessment, but rather as a consistent comparative objective that allows different injection strategies to be evaluated within the same optimization framework.

4 Multi-Objective Bayesian Optimization Methodology

Bayesian optimization (BO) is a data-efficient approach for optimizing objective functions that are expensive to evaluate, analytically intractable, or lack gradient information. Instead of performing an exhaustive search, BO constructs a probabilistic surrogate model of the objective responses and sequentially selects new designs using an acquisition function. The acquisition function balances two competing goals: *exploration*, which favors regions with high predictive uncertainty, and *exploitation*, which prioritizes designs that are predicted to perform well (Garnett, 2023).

4.1 Multi-objective problem formulation

These solutions are characterized through the concept of *Pareto dominance*. A design X_a is said to dominate another design X_b if it performs at least as well in all objectives and strictly better in at least one objective. A design is called *Pareto optimal* if no other feasible design dominates it. The set of all such non-dominated solutions forms the *Pareto set*, and its image in objective space is referred to as the *Pareto front* (Deb et al., 2002). Figure 4 illustrates this concept schematically.

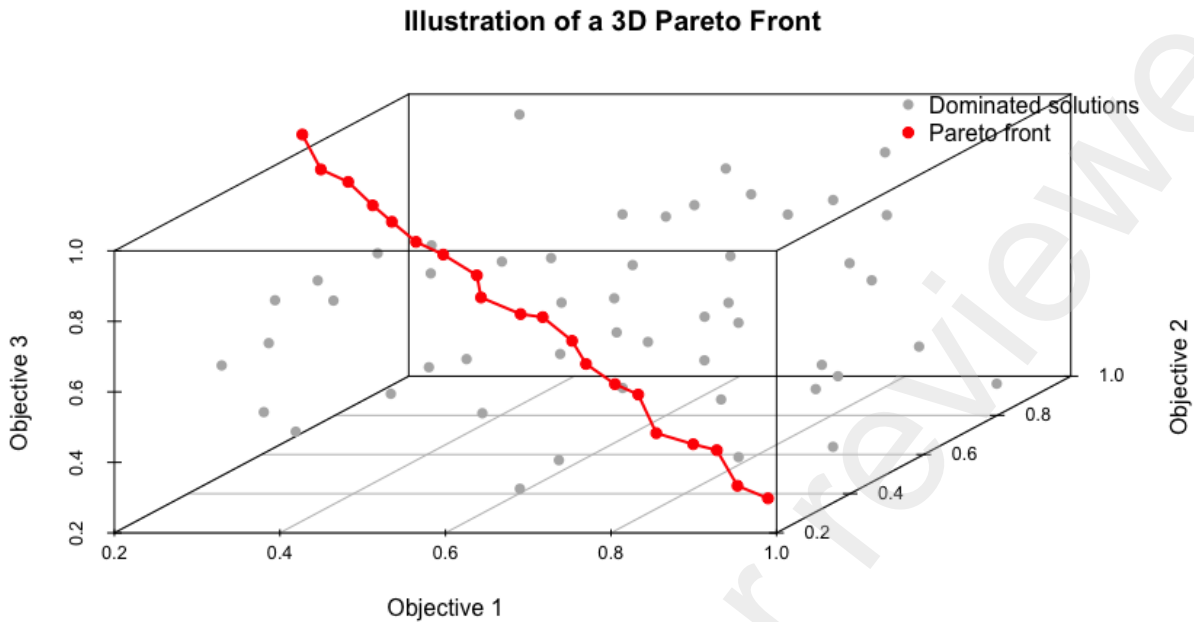


Figure 4: Illustrative example of Pareto dominance in a three-objective problem. The non-dominated solutions form the Pareto front, representing trade-offs among conflicting objectives.

In problems where evaluating the forward simulation model is computationally expensive, identifying the Pareto front through exhaustive search becomes impractical. MOBO addresses this challenge by combining surrogate modeling with acquisition strategies that guide the search toward promising regions of the design space while maintaining diversity among candidate solutions. This allows efficient approximation of the Pareto front within a limited evaluation budget.

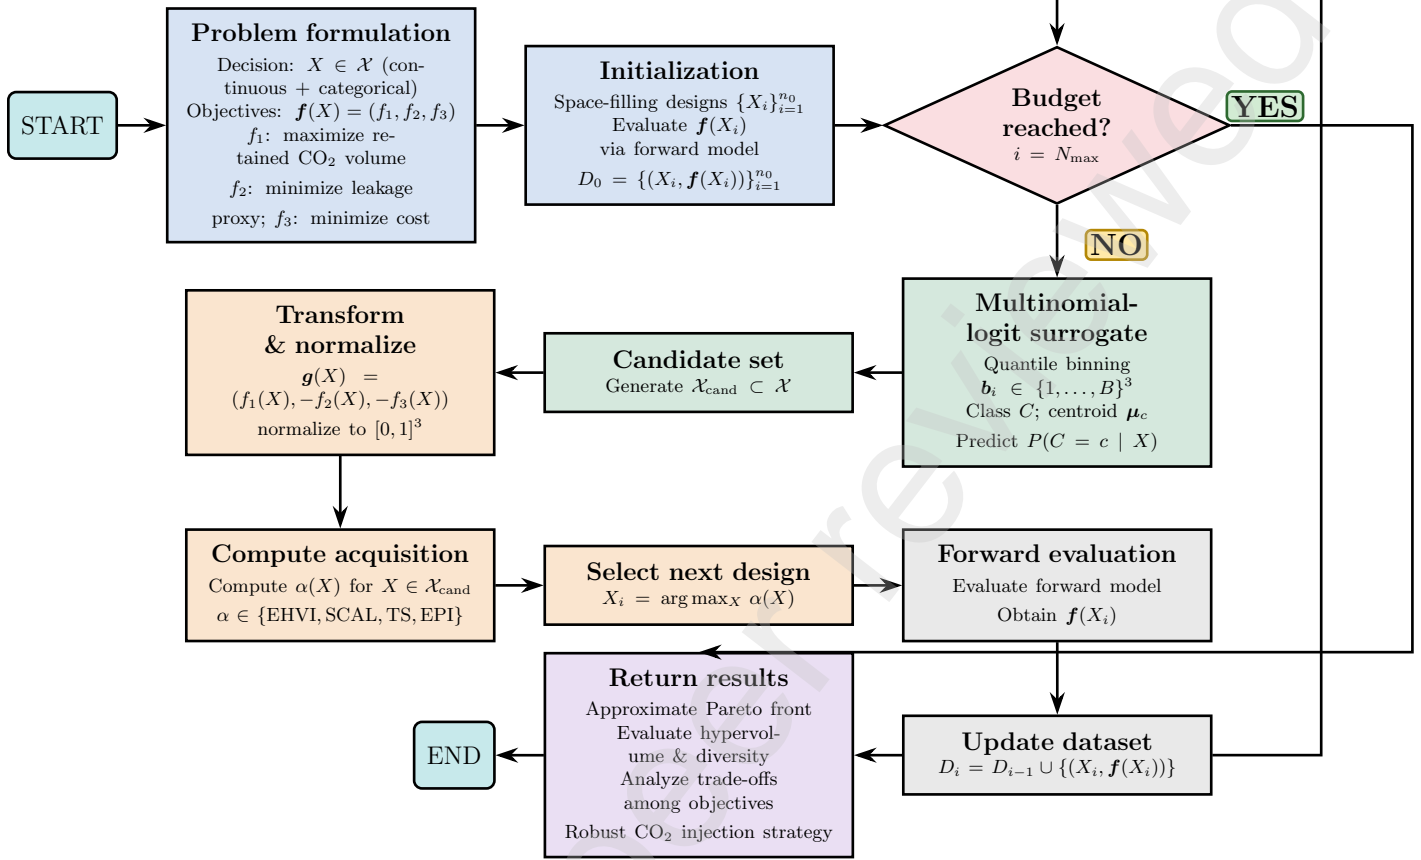


Figure 5: Workflow of the multi-objective Bayesian optimization (MOBO) framework.

Let $X \in \mathcal{X}$ denote a vector of decision variables, including both continuous and categorical components. The optimization problem considered in this study involves three objectives from Section 3.3:

For Pareto-based acquisition functions, objectives are transformed into a unified maximization framework:

$$\mathbf{g}(X) = (g_1(X), g_2(X), g_3(X)) = (f_1(X), -f_2(X), -f_3(X)). \quad (6)$$

Each component of $\mathbf{g}(X)$ is subsequently normalized to the unit interval using fixed bounds estimated from the initial design evaluations. This normalization ensures comparable scaling across objectives when evaluating acquisition functions.

4.2 Multinomial logit surrogate model

Many MOBO methods rely on GP surrogate models (Rasmussen and Williams, 2006). However, the present problem involves categorical decision variables and objective responses that may be non-smooth or non-Gaussian. To address these challenges, we adopt a classification-based surrogate model based on MNL regression, which is computationally efficient and naturally accommodates categorical predictors.

Quantile binning of objectives. Each objective function is discretized into B ordinal bins using empirical quantiles computed from the set of evaluated designs. For a design X_j with objective

vector $f(X_j) = (f_{1j}, f_{2j}, f_{3j})$ the discretization yields a bin-index vector

$$\mathbf{b}_j = (b_{1j}, b_{2j}, b_{3j}), \quad b_{kj} \in \{1, \dots, B\}.$$

The triplet \mathbf{b}_j is encoded as a single categorical class label representing a region of the multi-objective outcome space.

Multinomial logit formulation. The surrogate model learns the mapping from decision variables X to the probability distribution over observed outcome classes. For a candidate design X , the predicted probability of class c is

$$P(C = c | X) = \frac{\exp(\mathbf{w}_c^\top \tilde{X})}{\sum_{k=1}^G \exp(\mathbf{w}_k^\top \tilde{X})}, \quad (7)$$

where \tilde{X} denotes the encoded design vector, \mathbf{w}_c are class-specific parameters estimated by maximum likelihood, and G is the number of observed classes.

Centroid-based outcome representation. Each class c is associated with an empirical centroid $\boldsymbol{\mu}_c = (\mu_{1,c}, \mu_{2,c}, \mu_{3,c})$, computed as the mean objective vector of all evaluated designs assigned to that class. The surrogate therefore induces a discrete predictive distribution over objective outcomes, represented as a mixture over class centroids weighted by $P(C | X)$. This predictive distribution is sampled during acquisition evaluation, generating Monte Carlo realizations of objective vectors that are used to approximate acquisition values.

4.3 Acquisition functions

Acquisition functions convert surrogate predictions into decision rules that guide the selection of new designs for evaluation. All acquisition functions are evaluated in normalized maximization space and computed using Monte Carlo samples drawn from the MNL-induced predictive distribution.

4.3.1 Expected hypervolume improvement

Expected Hypervolume Improvement (EHVI) quantifies the expected increase in dominated hypervolume when a candidate design is evaluated (Emmerich et al., 2005). Let \mathcal{F} denote the current set of non-dominated solutions in normalized objective space and \mathbf{r}^{ref} a dominated reference point. The hypervolume is defined as

$$\text{HV}(\mathcal{F}) = \text{Vol} \left(\bigcup_{\mathbf{y} \in \mathcal{F}} [\mathbf{r}^{\text{ref}}, \mathbf{y}] \right). \quad (8)$$

Here, $\text{Vol}(\cdot)$ denotes the Lebesgue measure of a set in objective space. Each term $[\mathbf{r}^{\text{ref}}, \mathbf{y}]$ is the hyperrectangle spanned between the reference point \mathbf{r}^{ref} and a non-dominated solution $\mathbf{y} \in \mathcal{F}$, and the union $\bigcup_{\mathbf{y} \in \mathcal{F}}$ aggregates all such hyperrectangles to form the total dominated region in objective space. The reference point \mathbf{r}^{ref} is chosen to be dominated by all feasible solutions, ensuring the hypervolume is well-defined and finite. For a candidate design X , EHVI is then defined as

$$\text{EHVI}(X) = \mathbb{E}[\text{HV}(\mathcal{F} \cup \{\mathbf{y}(X)\}) - \text{HV}(\mathcal{F})], \quad (9)$$

where $\mathbf{y}(X)$ is the random objective vector obtained by evaluating X , and the expectation is approximated via Monte Carlo sampling from the predictive mixture distribution. A higher $\text{EHVI}(X)$ indicates that evaluating X is expected to expand the dominated hypervolume more, making it the preferred selection criterion when exploration of the Pareto front is the priority.

4.3.2 Thompson sampling

Thompson sampling (TS) selects candidates based on a single stochastic realization of the surrogate prediction. For each candidate X , a class $c^* \sim P(C | X)$ is sampled from the predictive class distribution, and the corresponding centroid $\boldsymbol{\mu}_{c^*}$ is used as a stochastic realization of $\boldsymbol{g}(X)$. The acquisition value is the resulting hypervolume gain:

$$\text{TS}(X) = \text{HV}(\mathcal{F} \cup \{\boldsymbol{\mu}_{c^*}\}) - \text{HV}(\mathcal{F}). \quad (10)$$

Repeated resampling across iterations yields a natural exploration–exploitation trade-off (Garnett, 2023).

4.3.3 Expected preference improvement

Expected Preference Improvement (EPI) incorporates stakeholder preferences through a scalar utility function (Paria et al., 2020). Using normalized maximization-space objectives $\tilde{\boldsymbol{g}}(X)$, the utility is defined as

$$u(X) = \sum_{k=1}^3 \lambda_k \tilde{g}_k(X), \quad \lambda_k \geq 0, \quad \sum_{k=1}^3 \lambda_k = 1. \quad (11)$$

Let u_{\max} denote the best observed utility value. The EPI acquisition is

$$\text{EPI}(X) = \mathbb{E}[\max(u(X) - u_{\max}, 0)], \quad (12)$$

with the expectation approximated via Monte Carlo sampling over surrogate-predicted classes.

4.3.4 Scalarized confidence bound

The Scalarized Confidence Bound (SCAL) acquisition balances expected performance and uncertainty in scalarized utility space. For each candidate X , Monte Carlo samples from the predictive mixture yield utility samples $\{u^{(s)}(X)\}_{s=1}^S$. The acquisition value is defined as

$$\text{SCAL-UCB}(X) = \bar{u}(X) + \kappa \text{sd}(u^{(s)}(X)). \quad (13)$$

where $\bar{u}(X)$ is the sample mean utility, $\text{sd}(\cdot)$ is the sample standard deviation, and $\kappa > 0$ controls the strength of exploration.

4.4 Bayesian optimization algorithm

The MOBO procedure alternates between surrogate model fitting, acquisition-based candidate selection, and evaluation using the forward simulation model. Figure 5 illustrates the overall workflow. The diagram follows a sequential loop: after an initialization phase using space-filling designs, the algorithm alternates between fitting the MNL surrogate, computing acquisition values over a candidate set, and evaluating the forward simulation model on the most promising design. Once the evaluation budget N_{\max} is exhausted, the loop terminates and the collected dataset is post-processed to extract the approximate Pareto front, representing the set of non-dominated trade-off solutions among the three objectives f_1 , f_2 , and f_3 .

Algorithm 1 MOBO with MNL surrogate

Require: Initial dataset $D_0 = \{(X_j, \mathbf{f}(X_j))\}_{j=1}^{n_0}$, evaluation budget N_{\max} , acquisition function $\alpha(\cdot)$

- 1: **for** $t = n_0 + 1$ **to** N_{\max} **do**
 - 2: Fit MNL surrogate using quantile-binned objective classes
 - 3: Generate candidate set $\mathcal{X}_{\text{cand}} \subset \mathcal{X}$
 - 4: **for all** $X \in \mathcal{X}_{\text{cand}}$ **do**
 - 5: Draw Monte Carlo samples from the surrogate predictive distribution
 - 6: Evaluate acquisition value $\alpha(X)$
 - 7: **end for**
 - 8: Select next design $X_t = \arg \max_{X \in \mathcal{X}_{\text{cand}}} \alpha(X)$
 - 9: Evaluate forward simulation model to obtain objective vector $\mathbf{f}(X_t)$
 - 10: Update dataset $D_t = D_{t-1} \cup \{(X_t, \mathbf{f}(X_t))\}$
 - 11: **end for**
 - 12: Extract non-dominated solutions from $\{\mathbf{f}(X_j)\}$ to obtain the approximate Pareto front
-

Algorithm 1 summarizes the overall MOBO workflow used in this study.

5 Empirical Evaluation

We evaluate the proposed MOBO framework in the Sleipner-inspired simulation environment. The primary objective is to evaluate how different acquisition strategies navigate the mixed-variable design space and identify high-quality trade-off solutions under a limited computational budget. All objective evaluations are obtained through the integrated R–Julia pipeline for the simulator. For each candidate design, a staged injection schedule is constructed and passed to the Julia-based simulator (Svee, 2026). We compute the three performance metrics: cumulative retained CO_2 storage (f_1), a leakage proxy (f_2), and total project cost (f_3). We compare the four acquisition strategies.

All of them are initialized using the same design and executed under an identical evaluation budget to ensure a consistent and fair comparison.

5.1 Experimental Setup

The optimization problem is defined over a mixed decision space:

$$X = (\text{scenario}, \text{rate}_{\text{mt}}, L_{\text{mid}}, L_{\text{top}}),$$

where the scenario determines the well configuration, the injection rate controls the annual CO_2 injection per well, and the layer indices define the staged allocation of injection across reservoir layers during different phases.

Each optimization run begins with an initial design, followed by sequential Bayesian optimization iterations. Each iteration generates a collection of possible solutions, which are then assessed using the acquisition function. The selected candidate is simulated, and the resulting observations are used to update the surrogate model. For consistency, objective vectors are transformed into a maximization setting $(f_1, -f_2, -f_3)$ and normalized objective vectors are transformed to the unit interval before computing hypervolume-based performance metrics.

5.2 Response-Surface Analysis

To gain insight into the structure of the design space, we first examine surrogate-predicted response surfaces over key decision variables. The observed design points in Figure 6 are overlaid on the surrogate-predicted response surfaces.

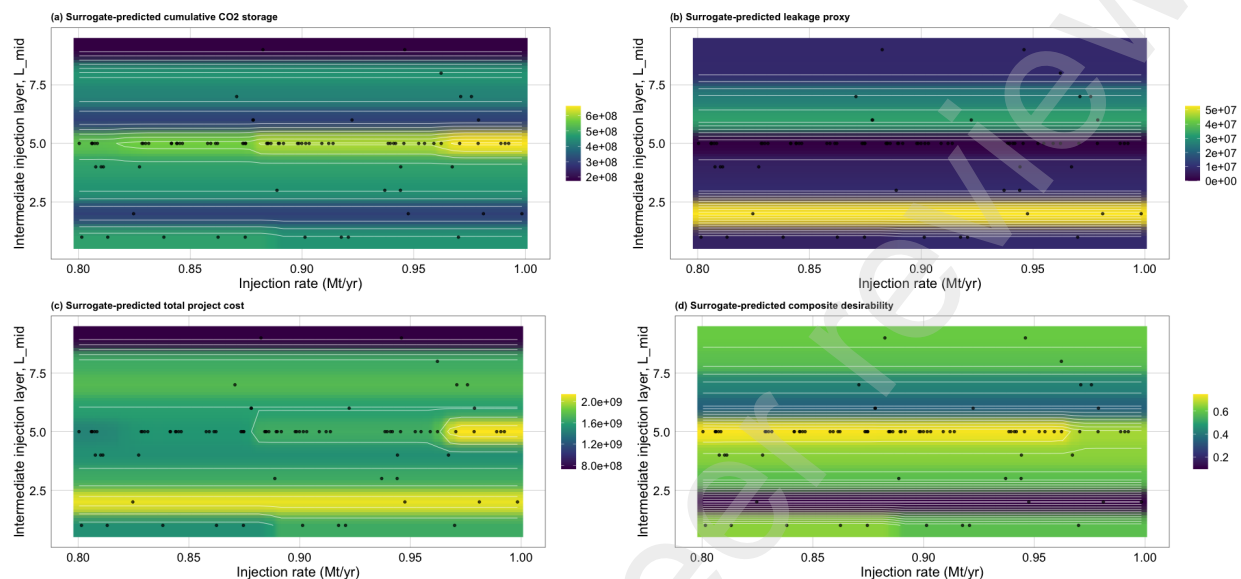


Figure 6: Surrogate-predicted response surfaces for (a) cumulative CO₂ storage, (b) leakage proxy, (c) total project cost, and (d) composite desirability as functions of injection rate and intermediate injection layer L_{mid} .

Cumulative CO₂ storage (panel a): Storage is dominated by the choice of intermediate injection layer rather than injection rate. The response surface exhibits a clear horizontal banding structure, with $L_{mid} \approx 5$ producing the highest storage values (approaching 6×10^8), while both low ($L_{mid} \lesssim 2$) and high ($L_{mid} \gtrsim 7.5$) layer indices yield substantially reduced storage. The injection rate has comparatively little influence on storage within the explored range of 0.80–1.00 Mt yr⁻¹, as evidenced by the near-vertical contour lines. The concentration of evaluated design points around $L_{mid} = 5$ confirms that the optimization rapidly identifies this layer as the preferred injection depth for maximizing storage.

Leakage proxy (panel b): The leakage surface displays the sharpest and most clearly stratified structure of all four objectives. Low leakage (dark purple, approaching zero) is confined to a narrow band around $L_{mid} \approx 5$, while both shallow ($L_{mid} \lesssim 2.5$) and deep ($L_{mid} \gtrsim 7$) layers produce very high leakage values (yellow, up to 5×10^7). This sharp transition suggests that the intermediate layer acts as a critical caprock boundary: injection near $L_{mid} = 5$ places CO₂ beneath an effective seal, whereas injection at other depths exposes it to vertical migration pathways. As with storage, injection rate has minimal influence on leakage, reinforcing that vertical placement is the dominant control variable for containment.

Total project cost (panel c): Unlike storage and leakage, cost exhibits a two-dimensional sensitivity to both layer index and injection rate. The lowest costs (dark purple, around 8×10^8) occur at low layer indices and low injection rates, while the highest costs (yellow, around 2×10^9) are concentrated at high layer indices and high injection rates. A notable local cost feature is visible near $L_{mid} = 5$ and injection rate ≈ 0.95 Mt yr⁻¹, where the contours show a localized cost ridge,

suggesting that the operationally preferred layer incurs a cost premium at higher injection rates. This creates a direct tension with the storage and leakage objectives, which both favor $L_{\text{mid}} \approx 5$, and illustrates why cost cannot be minimized simultaneously with the other two objectives.

Composite desirability (panel d): The composite desirability surface synthesizes the three individual objectives and reveals a narrow favorable region centered around $L_{\text{mid}} \approx 5$ across all injection rates, with desirability values reaching approximately 0.6–0.7 in this band. Outside this band, desirability drops sharply toward 0.2 or below, confirming the non-convex and highly layer-sensitive nature of the optimization problem. The relatively flat response along the injection-rate axis within the favorable band suggests that, once the optimal layer is identified, fine-tuning the injection rate yields only marginal additional gains in composite performance. The clustering of evaluated design points within this high-desirability band demonstrates that the MOBO algorithm successfully identifies and concentrates its sampling budget in the most promising region of the design space.

A three-dimensional view of the storage surface further highlights non-linear interactions between decision variables is shown in .

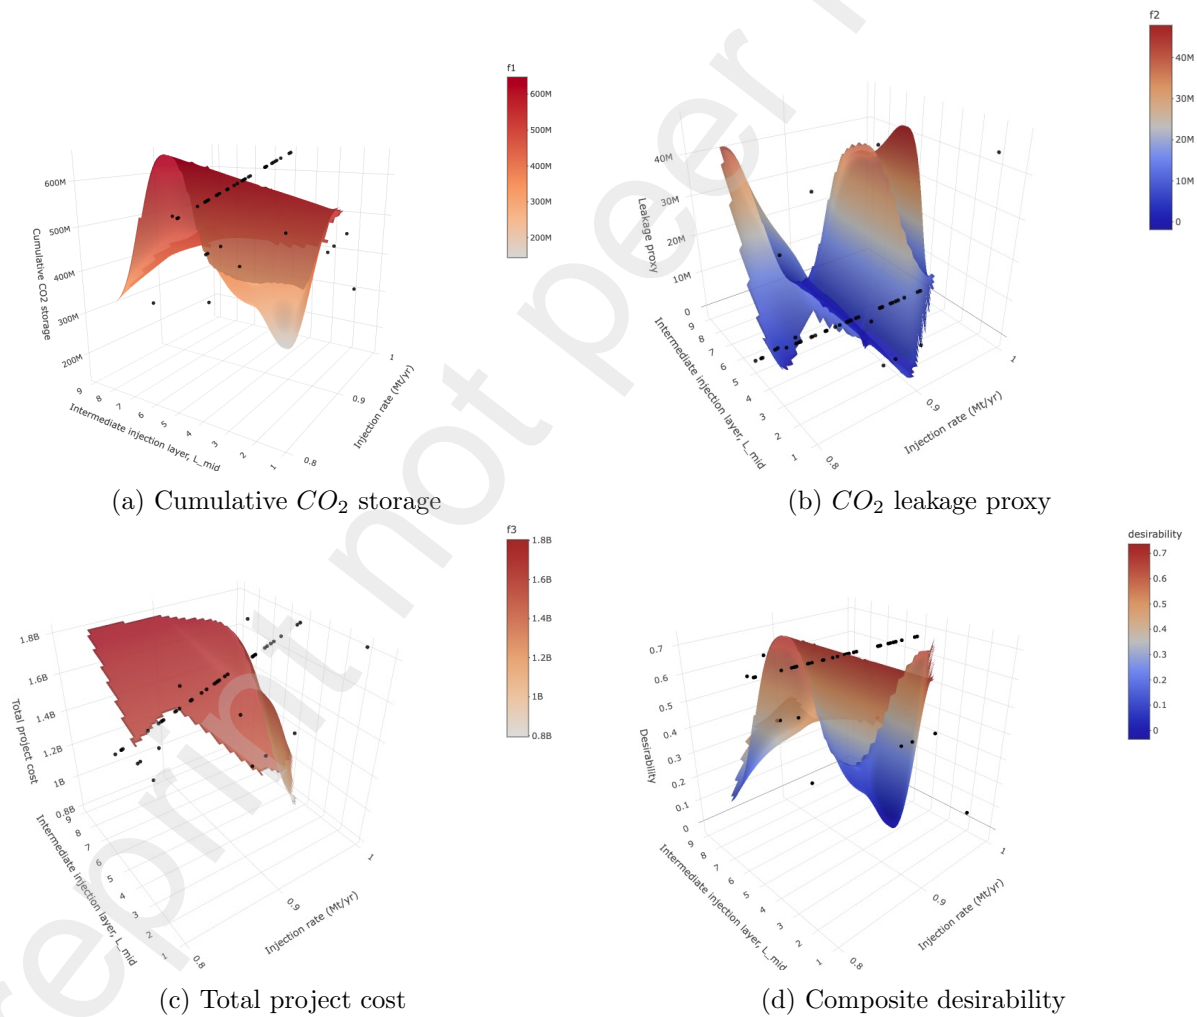


Figure 7: Three-dimensional surrogate response surfaces of storage, leakage, cost, and desirability as functions of injection rate and intermediate injection layer.

Figure 7 presents the three-dimensional surrogate-predicted response surfaces with observed design points overlaid. As shown in Fig. 7(a), the storage response exhibits strong non-linearity with localized maxima at specific intermediate layers, indicating a pronounced sensitivity to vertical placement. In Fig. 7(b), the leakage proxy varies sharply across layers, suggesting the influence of vertical migration pathways and stratigraphic heterogeneity. In contrast, Fig. 7(c) shows that the total project cost changes more smoothly with injection rate and is less sensitive to layer selection. Finally, Fig. 7(d) highlights that the composite desirability is confined to a narrow region of favorable trade-offs, reflecting a non-convex optimization landscape and underscoring the need for adaptive surrogate-based optimization strategies.

5.3 Optimization Performance Across Acquisition Strategies

The convergence behavior of the acquisition strategies is evaluated using hypervolume, which captures both the convergence quality and diversity of the Pareto front. Figure 8 compares the optimization performance of different acquisition strategies using hypervolume-based metrics.

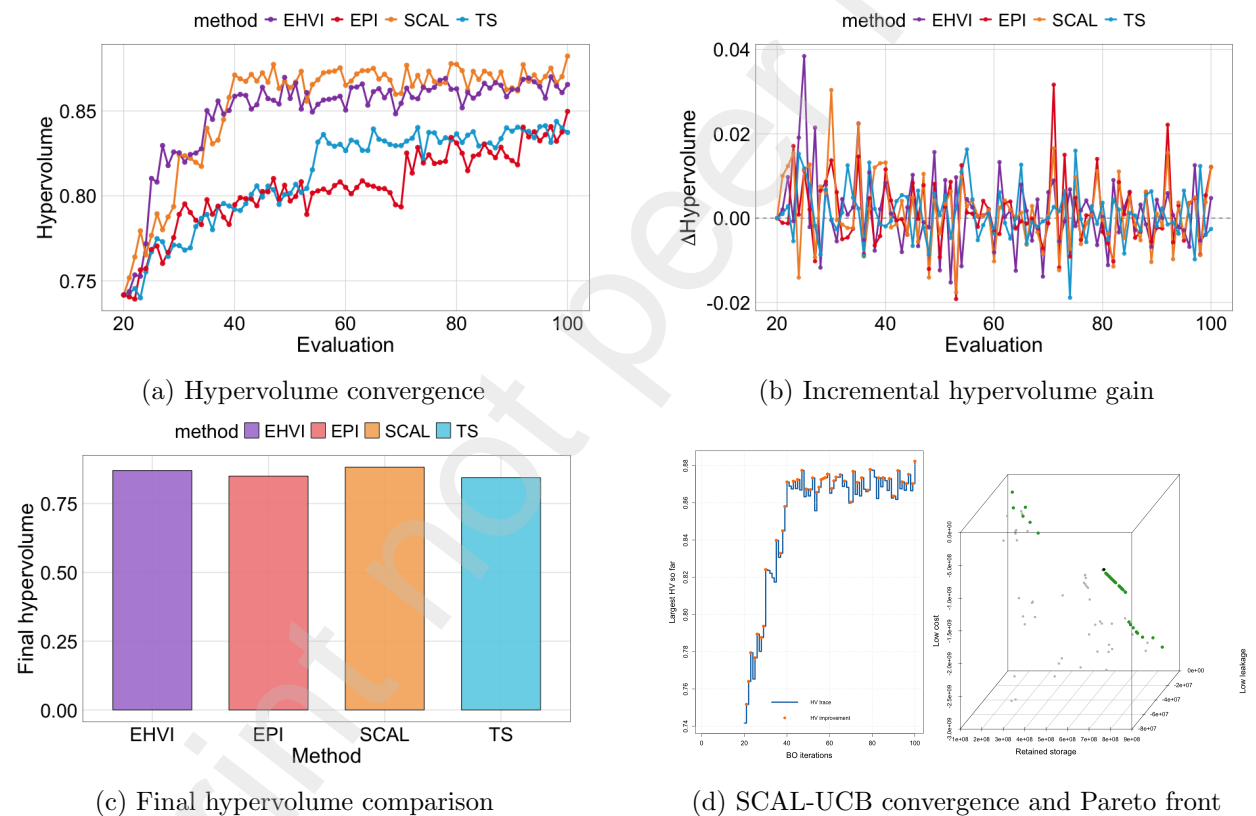


Figure 8: Hypervolume-based performance comparison of acquisition strategies (EHVI, EPI, SCAL-UCB, and TS).

Hypervolume convergence (panel a): All four strategies start from a similar hypervolume of approximately 0.74 at evaluation 20 and improve rapidly through the first 30–40 evaluations. EHVI converges earliest, reaching a stable plateau around 0.85–0.86 by evaluation 35 and maintaining it consistently thereafter. SCAL-UCB converges slightly later but ultimately achieves the highest hypervolume, fluctuating around 0.87–0.88 in the later stages, reflecting its continued discovery

of new non-dominated solutions throughout the budget. EPI and TS lag behind, stabilizing at approximately 0.83–0.84, with TS showing a notable late-stage improvement around evaluation 55–60 before plateauing. The persistent gap between SCAL-UCB and the remaining strategies suggests that the UCB exploration bonus continues to identify hypervolume-improving designs even at later iterations, whereas EHVI, EPI, and TS saturate earlier.

Incremental hypervolume gain (panel b): The incremental gains ΔHV fluctuate around zero for all strategies after the initial exploration phase, confirming that most evaluations beyond iteration ~ 40 yield only marginal improvements to the Pareto front. EHVI produces the largest single gain spike (approximately $+0.039$ near evaluation 25), consistent with its hypervolume-targeted acquisition function making an early decisive improvement. EPI shows a comparably large late spike near evaluation 80 ($+0.031$), suggesting it occasionally identifies high-impact designs even in later iterations. SCAL-UCB and TS produce more moderate but more frequent positive increments, reflecting their broader exploration behavior. Notably, all strategies also produce negative ΔHV values, which occur when a newly evaluated point is dominated and does not expand the front, highlighting the inherent difficulty of consistently improving near a well-converged Pareto front.

Final hypervolume (panel c): The bar chart confirms the ranking observed in panel (a): SCAL-UCB achieves the highest final hypervolume (≈ 0.885), followed closely by EHVI (≈ 0.865), with TS and EPI both reaching approximately 0.84. The differences are modest in absolute terms but consistent across the full optimization run, suggesting that SCAL-UCB’s exploration advantage compounds over many iterations rather than arising from a single lucky evaluation.

SCAL-UCB convergence and Pareto front (panel d): The individual SCAL-UCB convergence trace confirms a rapid improvement phase between evaluations 20–45, punctuated by frequent hypervolume-improving steps (red dots), followed by a more stable plateau above 0.86 with occasional further gains up to 0.88 at the end of the budget. The corresponding 3D Pareto front (right) reveals that SCAL-UCB identifies two geometrically distinct clusters of non-dominated solutions: a tight cluster at high retained storage ($\sim 7\text{--}9 \times 10^8$) with moderate cost and near-zero leakage and a more dispersed set at lower storage values covering a wider range of cost and leakage trade-offs. This two-cluster structure reflects SCAL-UCB’s phase transition from broad exploration to focused exploitation of the high-storage, low-leakage region identified around $L_{\text{mid}} \approx 5$, consistent with the decision-variable analysis.

5.4 Structure of the Objective Space

The trade-offs between objectives are illustrated through pairwise projections.

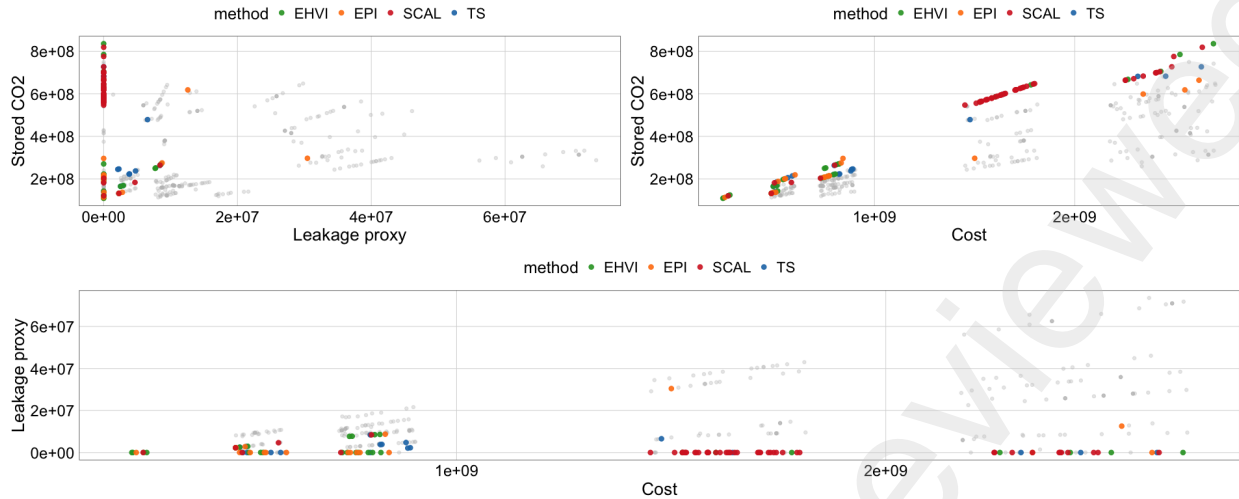


Figure 9: Pairwise objective trade-offs between storage, leakage, and cost with Pareto-optimal solutions highlighted.

Figure 9 illustrates the trade-offs between the optimization objectives using pairwise projections. Here, the colored markers represent Pareto-optimal solutions obtained by different acquisition strategies, while gray points denote dominated designs. As shown in Fig. 9(a), higher CO_2 storage is generally associated with increased leakage risk, although some strategies identify regions with improved trade-offs. Fig. 9(b) indicates a clear positive correlation between storage and cost, reflecting the economic burden of achieving higher storage capacities. In Fig. 9(c), low-leakage solutions are concentrated within specific cost ranges, suggesting constraints imposed by geological and operational factors. The distribution of Pareto points highlights the inherent conflict among objectives, confirming that simultaneous maximization of storage and minimization of both cost and leakage is not feasible. Moreover, different strategies explore distinct regions of the Pareto front, indicating variability in search behavior and solution diversity.

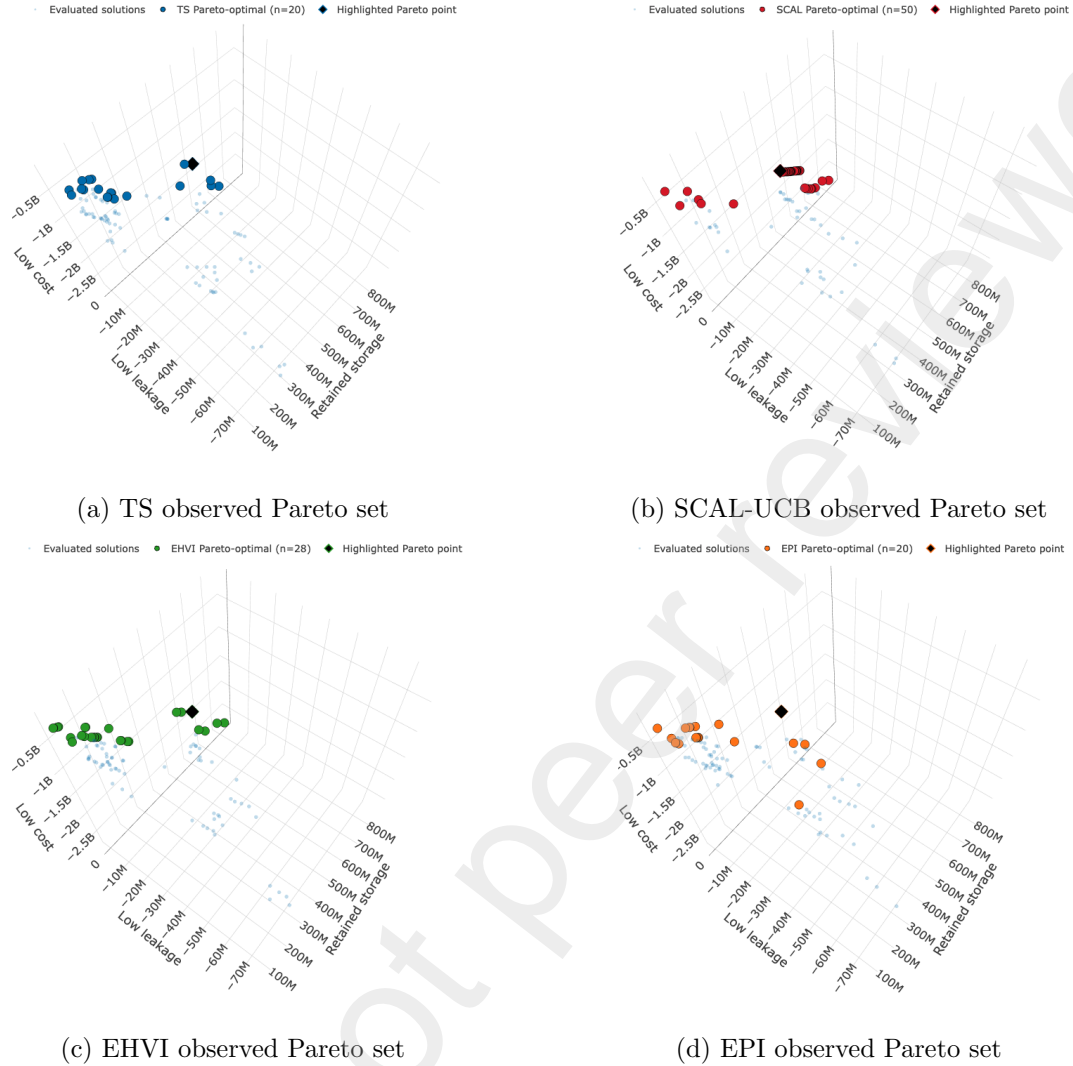


Figure 10: Observed Pareto sets in the three-dimensional objective space for different acquisition strategies.

The three-dimensional Pareto front provides a global view of non-dominated solutions. Figure 10 compares the observed Pareto sets obtained using different acquisition strategies. As shown in Fig. 10(a-d), each subplot presents evaluated solutions and the corresponding Pareto-optimal points in the objective space defined by retained CO_2 storage, leakage proxy, and project cost. The Pareto fronts reveal clear trade-offs among competing objectives, where higher storage is generally associated with increased cost and leakage risk. SCAL-UCB (Fig. 10(b)) produces a denser and more uniformly distributed Pareto front, indicating improved exploration and coverage of the objective space. In contrast, EHVI and TS (Fig. 10(a,c)) identify fewer but well-converged solutions, while EPI (Fig. 10(d)) shows comparatively limited diversity. These differences highlight the influence of acquisition strategy on exploration efficiency and the quality of the approximated Pareto front.

To further explore method-specific behavior and cross-method dominance, we analyze both the Pareto front derived from the SCAL-UCB strategy and a comprehensive Pareto dominance comparison among all acquisition methods. The SCAL-UCB Pareto front offers a comprehensive

perspective on how a singular acquisition method navigates and approximates the trade-off surface, emphasizing its capacity to discern a dense and well-organized array of non-dominated solutions. Conversely, the global Pareto dominance visualization consolidates solutions from all methods and highlights cross-method dominating points, providing an extensive view of relative performance and convergence attributes. Collectively, these representations offer complementary perspectives on individual method efficacy and overall optimization quality.

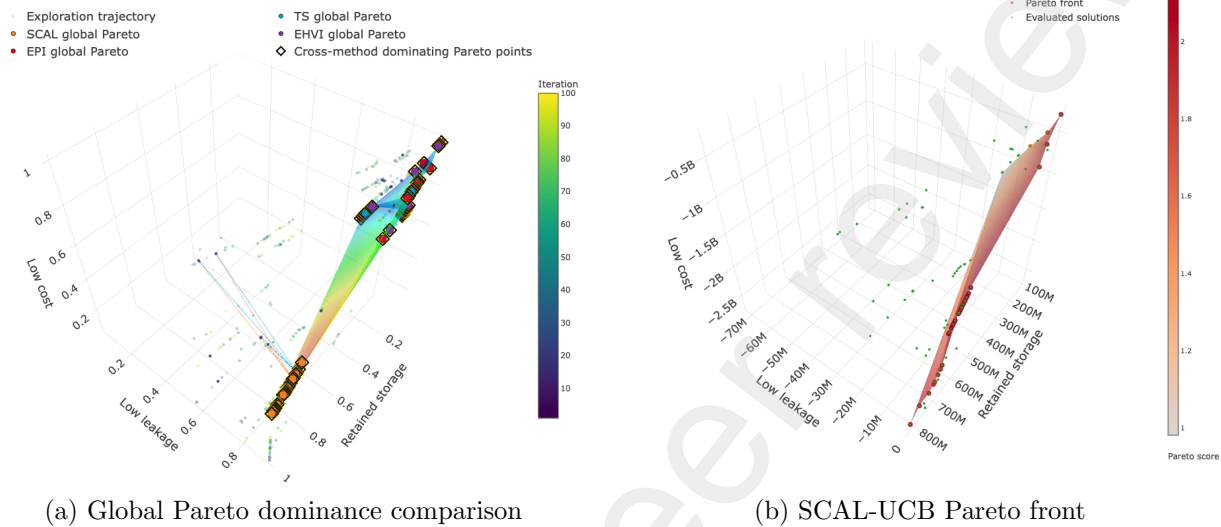


Figure 11: Comparison of SCAL-UCB Pareto front and global Pareto dominance across acquisition strategies.

The SCAL-UCB Pareto front has a well-defined and continuous structure, with solutions that cover a wide range of retained storage while requiring relatively little leakage and cost. The concentration of points along a smooth frontier implies that SCAL-UCB successfully balances exploration and exploitation, allowing it to densely populate high-quality parts of the objective space. The forefront advances towards high-storage solutions while preserving competitive costs and leakage levels, indicating a significant ability to identify beneficial trade-offs.

The global Pareto dominance analysis reveals that SCAL-UCB comprises a substantial portion of the dominant solutions, especially in areas characterized by extensive retained storage and minimal cost. Although EHVI and TS reveal competitive solutions, their coverage is less comprehensive, and EPI seems more conservative, exhibiting fewer dominant points. The existence of cross-method dominant spots signifies areas where specific procedures surpass others, highlighting the importance of acquisition function design in facilitating effective exploration. These findings demonstrate the efficacy of SCAL-UCB in attaining both diversity and dominance inside the Pareto set.

A complementary representation is provided using parallel coordinates, which reveals correlations and trade-off structures that are not easily observable in low-dimensional projections.

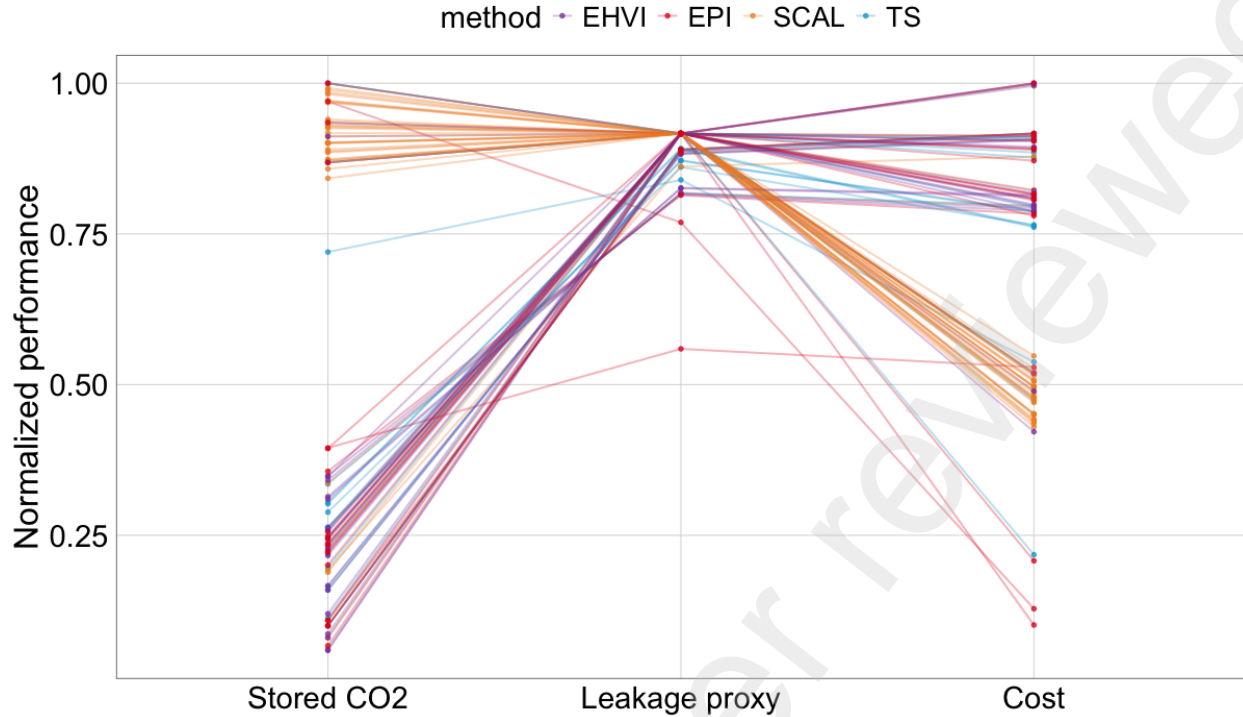


Figure 12: Parallel-coordinate representation of Pareto-optimal solutions in normalized objective space.

Figure 12 presents a parallel-coordinate visualization of all Pareto-optimal solutions across the four acquisition strategies, where each line represents a single non-dominated design and the three axes correspond to normalized stored CO₂, leakage proxy, and project cost (all normalized to [0, 1] where higher values indicate better performance for each objective).

Storage–leakage trade-off: The most prominent feature of the plot is the strong crossing pattern between the stored CO₂ and leakage proxy axes. Solutions with high storage (normalized values ≥ 0.85) converge toward a leakage proxy value of approximately 0.90, indicating that the best-storing designs achieve consistently low leakage. However, solutions with lower storage (normalized values 0.10–0.40) fan out across a wide range of leakage values, suggesting that reducing storage does not necessarily guarantee low leakage and that these configurations correspond to suboptimal layer placements away from $L_{\text{mid}} \approx 5$.

Leakage–cost trade-off: Between the leakage and cost axes, lines diverge sharply, revealing a strong conflict between these two objectives. Solutions that achieve near-zero leakage (leakage proxy ≈ 0.90) spread across the full range of cost values (0.10–1.00), confirming that low leakage does not constrain cost. A subset of lines, predominantly from EPI (red), drops to very low cost values (≤ 0.20), but these same solutions tend to have lower storage, indicating that EPI identifies some low-cost, low-storage configurations that the other strategies do not prioritize. The majority of lines from SCAL-UCB and EHVI cluster in the cost range 0.75–0.90, suggesting these strategies favor high-performance solutions that carry a moderate cost premium.

Strategy-level differences: SCAL-UCB (orange) contributes the densest bundle of lines, particularly in the high-storage region, consistent with its larger Pareto set size observed in Figure 14. EHVI (purple) produces a similarly dense bundle but with slightly less spread in the cost axis. EPI (red) is the most heterogeneous, spanning the widest range of storage and cost values, including

several low-storage, low-cost solutions not present in other strategies. TS (cyan/blue) contributes fewer lines overall and tends to occupy the moderate storage and cost range, with one notable outlier at low storage (≈ 0.30) and low cost (≈ 0.75). Overall, the parallel-coordinate plot confirms that no single strategy dominates across all three objectives simultaneously, and that the choice of acquisition function shapes not only the quantity but also the qualitative character of the trade-off solutions discovered.

5.5 Distribution of Objective Values

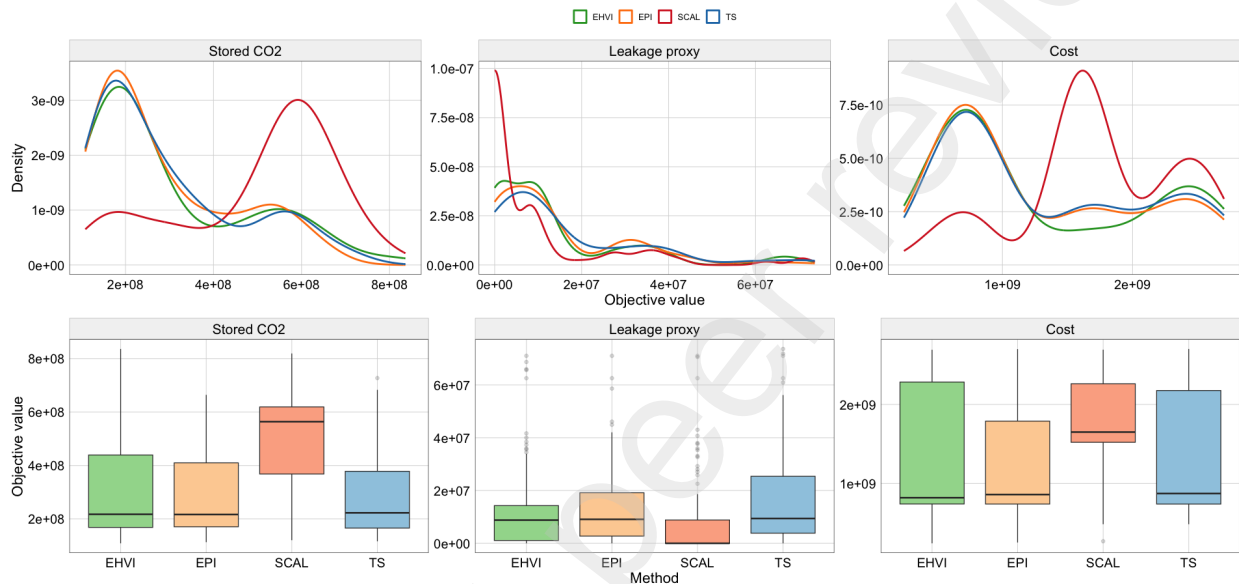


Figure 13: Distribution of objective values across acquisition strategies using density estimates and boxplots.

The distribution of evaluated objective values provides additional insight into the behavioral tendencies of each acquisition strategy. Figure 13 compares these distributions across all three objectives using both kernel density estimates (top row) and boxplots (bottom row).

SCAL-UCB exhibits a distinctly bimodal density for stored CO₂, with peaks at both low and high storage values, and the highest median storage among all strategies. This bimodality is characteristic of an Upper Confidence Bound-style strategy: the UCB bonus term actively encourages the surrogate to sample in uncertain regions, driving the algorithm to oscillate between well-understood low-storage configurations and promising but uncertain high-storage regions. The broad interquartile range in the boxplot confirms this wide exploration. For leakage, SCAL-UCB maintains relatively low values despite its exploratory nature, suggesting that the scalarization component effectively penalizes high-leakage designs even when exploring aggressively. The higher cost variability reflects the trade-off: aggressive exploration of storage inevitably visits expensive injection configurations.

Thompson Sampling (TS) shows the widest spread in cost among all strategies, as visible in both the density tail and boxplot whiskers. This is consistent with the stochastic nature of TS: by sampling a full realization of the surrogate at each iteration rather than optimizing a deterministic acquisition value, TS introduces randomness that prevents systematic exploitation of low-cost regions. The resulting cost distribution is therefore more diffuse. Storage and leakage distributions

for TS are similar to EHVI and EPI, indicating that its stochasticity primarily manifests in cost space rather than in the primary performance objectives.

EHVI and EPI both produce notably more concentrated distributions across all three objectives, with smaller interquartile ranges in the boxplots and unimodal, narrower density curves. This is consistent with their exploitation-heavy nature: EHVI explicitly targets regions of objective space that expand the dominated hypervolume, which naturally focuses sampling near already identified high-performing Pareto-optimal regions. Similarly, EPI (Expected Probability of Improvement) rewards designs that exceed the current Pareto front with high probability, again reinforcing convergence toward a narrow high-performance region. The cost of this focused behavior is reduced diversity in the evaluated dataset, as the strategies rarely venture into unexplored but potentially valuable regions of the design space.

Overall, the distributional differences directly reflect the exploration-exploitation balance encoded in each acquisition function: UCB-based scalarization and Thompson Sampling prioritize breadth at the cost of consistency, while hypervolume- and improvement-based strategies (EHVI, EPI) prioritize depth and convergence at the cost of diversity.

5.6 Pareto Front Evolution

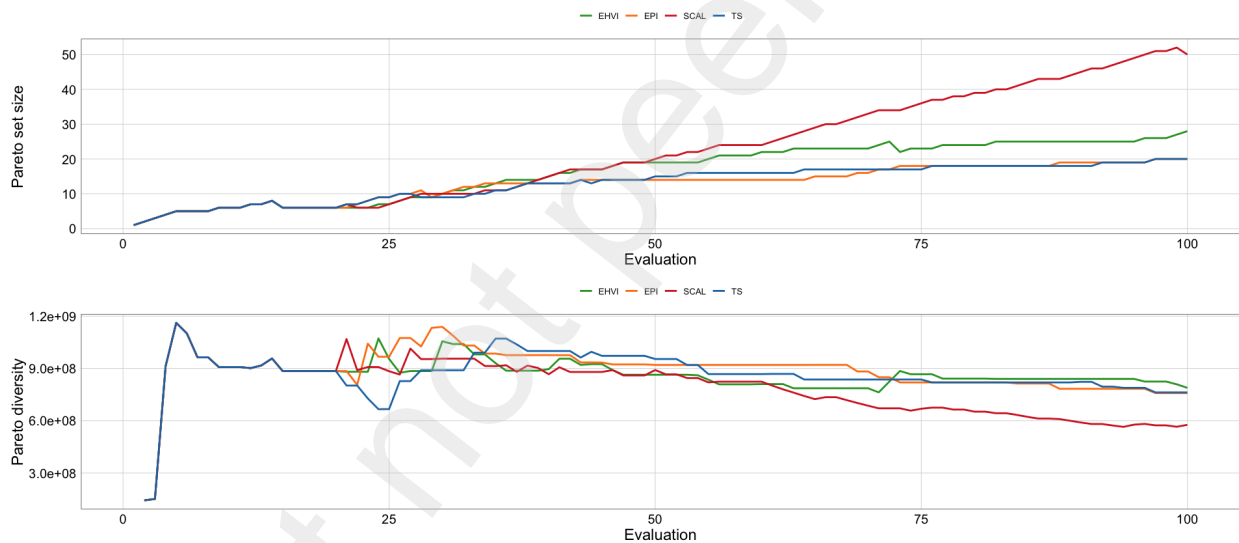


Figure 14: Evolution of Pareto set size and diversity over the optimization process.

We next analyze how the Pareto set evolves over the course of optimization. Figure 14 illustrates the evolution of Pareto set size (top) and diversity (bottom) over evaluations.

Pareto set size (top): The number of non-dominated solutions grows steadily for all strategies. SCAL-UCB accumulates the largest Pareto set by far, reaching approximately 50 non-dominated solutions by evaluation 100, compared to roughly 25–30 for EHVI and TS and around 20 for EPI. This is a direct consequence of the UCB exploration bonus, which drives SCAL-UCB to evaluate designs in diverse and previously unvisited regions of the objective space, many of which turn out to be non-dominated. EHVI and TS grow at a similar moderate rate, while EPI accumulates non-dominated solutions most slowly, consistent with its more conservative, improvement-focused sampling behavior.

Pareto diversity (bottom): All strategies exhibit a rapid initial rise in diversity during

the first ~ 25 evaluations, reflecting broad exploration of the objective space from the space-filling initialization. After this phase, diversity stabilizes and then gradually declines for all methods. Contrary to what the Pareto set size suggests, SCAL-UCB does *not* maintain higher diversity over time; instead, its diversity curve falls below those of EHVI, EPI, and TS in the later stages of optimization. This apparent paradox can be explained by the fact that while SCAL-UCB discovers many non-dominated solutions, these solutions tend to cluster in a few high-storage, low-leakage regions that the UCB bonus repeatedly revisits, reducing the average spacing between Pareto points. In contrast, EHVI and EPI, despite accumulating fewer non-dominated solutions overall, distribute them more evenly across the Pareto front, leading to comparatively higher diversity at later evaluations. TS lies between the two extremes, maintaining moderate diversity throughout due to its stochastic sampling mechanism. Overall, these results suggest that a large Pareto set size does not necessarily imply high diversity and that hypervolume- and improvement-based strategies may yield a better-spread approximation of the true Pareto front despite their smaller set sizes.

5.7 Decision-Space Characteristics

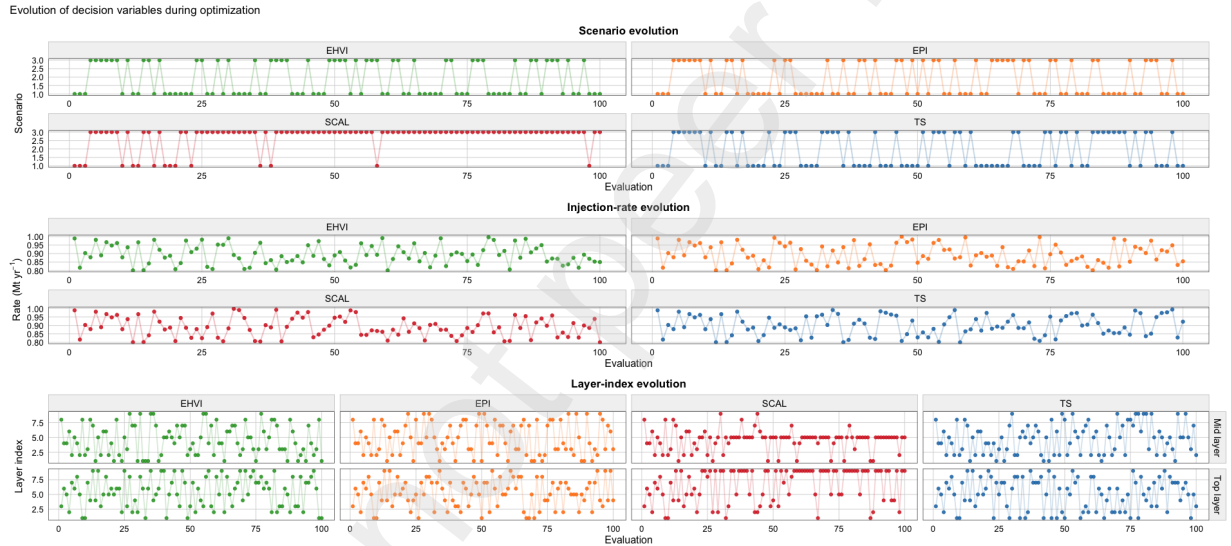


Figure 15: Evolution of decision variables across acquisition strategies during optimization.

The evolution of decision variables provides insight into how each acquisition strategy navigates the design space over time. Figure 15 shows the evolution of scenario selection (top), injection rate (middle), and mid- and top-layer indices (bottom) across all four strategies.

Scenario selection (top row): EHVI and EPI both continue switching between all three scenarios throughout the full 100 evaluations, with no clear convergence toward a single scenario. This reflects their hypervolume- and improvement-based objectives, which require sampling across diverse geological configurations to adequately cover the Pareto front. TS shows a similar pattern of persistent switching, consistent with its stochastic sampling mechanism. SCAL-UCB, by contrast, converges strongly toward Scenario 3 after approximately evaluation 25, with only occasional returns to Scenarios 1 and 2 in the early phase. This early lock-in indicates that SCAL-UCB’s UCB bonus identifies Scenario 3 as the most uncertain and high-reward geological configuration and subsequently concentrates its budget there, trading scenario diversity for deeper exploitation of a single geological setting.

Injection-rate evolution (middle row): All four strategies continue to explore the injection rate across the full range (0.80–1.00 Mt yr⁻¹) throughout optimization, with no clear convergence toward a preferred rate. This persistent exploration is consistent with the injection rate being a continuous variable that interacts with all three objectives simultaneously: higher rates increase stored CO₂ but also raise leakage risk and cost, so no single value dominates across the Pareto front. SCAL-UCB and TS show slightly more irregular switching between extreme values, reflecting their more aggressive exploration tendencies, while EHVI and EPI exhibit marginally smoother trajectories, consistent with their more exploitation-focused behavior.

Layer-index evolution (bottom row): The most striking differences across strategies appear in the layer-index panels. EHVI, EPI, and TS all explore mid- and top-layer indices broadly throughout optimization, sampling across the full range of available layers with no strong clustering. This suggests that these strategies do not identify a dominant vertical placement and instead continue probing diverse layer combinations to map the trade-off between storage and leakage across depth. SCAL-UCB, however, exhibits a qualitatively different pattern: after an initial exploratory phase, it converges strongly toward a narrow band of mid-layer indices (around index 5) and similarly clusters top-layer selections, maintaining these choices with high consistency from evaluation ~ 30 onward. This convergence implies that SCAL-UCB identifies specific layer combinations as robustly high-performing under Scenario 3, and its UCB bonus no longer assigns sufficient uncertainty reward to alternative layer configurations to override this preference. The non-uniformity of SCAL-UCB across all three variable types is therefore a coherent signature of its behavior: early broad exploration followed by confident, concentrated exploitation of a specific geological and operational configuration.

Overall, the decision-variable trajectories reveal that EHVI, EPI, and TS maintain broad and relatively uniform exploration of the design space throughout optimization, whereas SCAL-UCB undergoes a clear phase transition from exploration to exploitation, locking onto a preferred scenario, layer combination, and injection regime. This explains why SCAL-UCB accumulates the largest Pareto set but with lower late-stage diversity: it finds many non-dominated solutions within a narrow region of the design space rather than spreading evaluations evenly across it.

5.8 Summary of Optimization Results

Table 1 summarizes the findings. The results show that SCAL-UCB has the highest final hypervolume and the biggest Pareto set. This shows that it is very good at finding a wide range of high-quality trade-off solutions. EHVI and TS also do well, but EPI gives a more stable and conservative set of solutions.

Table 1: Summary statistics of optimization results across acquisition strategies.

Method	Evaluations	Best f_1	Min f_2	Min f_3	Pareto Points	Final HV
EHVI	100	8.36×10^8	0	2.43×10^8	28	0.870
EPI	100	6.64×10^8	0	2.53×10^8	20	0.850
SCAL	100	8.19×10^8	0	2.70×10^8	50	0.882
TS	100	7.27×10^8	0	4.85×10^8	20	0.844

Table 2: Summary statistics of the Pareto-optimal solutions obtained during optimization.

Metric	Minimum	Maximum
Rate ($rate_{mt}$)	0.800	0.998
L_{mid}	1	9
f_1	2.64×10^8	6.48×10^8
f_2	0	4.30×10^7
f_3	8.01×10^8	1.81×10^9

Total Pareto solutions: 47

Table 2 provides metrics of the optimal solutions. The empirical results show that the proposed MOBO framework can find injection strategies that strike a balance between storage performance, containment security, and economic cost with only a few simulator evaluations.

6 Concluding Remarks

This study presented a MOBO framework for the design of geological CO_2 injection strategies under three competing goals: maximizing retained storage, minimizing a leakage-related proxy, and minimizing total project cost. The problem was formulated over a mixed decision space comprising well configuration, injection rate, and staged allocation of injection across reservoir layers. By coupling the optimization procedure to a simulation-based forward model through an integrated R-Julia workflow (Svee, 2026), we were able to evaluate candidate strategies efficiently while preserving the main physical mechanisms governing storage, migration, and cost behavior.

This work’s primary contribution is the adaptation of MOBO to a CCS setting in which decision variables are not solely continuous and objective responses may exhibit discontinuity. The research employed a MNL surrogate model to adeptly handle the categorical characteristics of the design variables while maintaining computational efficiency. Within this framework, four acquisition strategies — TS, SCAL-UCB, EHVI, and EPI were analyzed, employing consistent initialization, normalization, and evaluation budgets. This made it possible to reliably compare how different search algorithms move through the design space and get close to the Pareto front.

The empirical results demonstrate that the proposed framework can identify substantial trade-offs with a relatively limited number of simulator evaluations. SCAL-UCB had the best overall performance of all the acquisition strategies that were tested. It attained the final hypervolume and produced the most extensive Pareto set, signifying an efficient equilibrium between the search of doubtful areas and the utilization of promising solutions. EHVI demonstrated strong performance, yielding competitive high-quality solutions, while TS was successful yet exhibited greater variability across the optimization trajectory. In contrast, EPI exhibited a more cautious behavior, resulting in a smaller and more concentrated array of trade-off solutions.

The resulting Pareto fronts provide useful insight into the structure of the CCS design problem. High-storage solutions were generally associated with increased project cost, while low-leakage designs were concentrated in narrower regions of the decision space. This suggests that containment performance is strongly controlled by specific operational choices, especially those related to vertical injection allocation. The response-surface and decision-space analyses support this interpretation: injection rate remained relatively broadly explored, whereas the layer variables exhibited much stronger structure, with high-performing solutions clustering around particular layer combinations. In other words, vertical placement appears to be one of the dominant controls on storage performance and containment behavior in the present simulator setting.

The optimization outcomes indicate that no singular injection approach is uniformly superior. Conversely, several operational regimes may excel based on the significance attributed to storage, leakage, and cost. While single-well configurations are prevalent among optimal solutions, multi-well designs also feature on the Pareto front, suggesting that more wells may enhance performance in certain instances, albeit their increased infrastructure costs. Similarly, the optimal injection rates were frequently elevated but not always at their peak, highlighting the necessity of balancing storage efficiency with migration-related drawbacks and cost increases.

Multiple extensions would enhance the framework significantly. A reasonable subsequent action is to integrate geological uncertainty more explicitly, for instance, via ensembles of reservoir realizations or stochastic representations of critical subsurface characteristics. The established objectives could be broadened to encompass pressure-management criteria, regulatory or monitoring limits, and long-term risk indicators. Furthermore, it may be beneficial to investigate alternative surrogate models, particularly those capable of embodying more intricate prediction structures while maintaining the capacity to manage mixed-variable decision spaces. Ultimately, multi-fidelity or adaptive sampling solutions may enhance computing efficiency as the forward model becomes more costly or as higher-fidelity physics are incorporated.

CRedit authorship contribution statement

Muhammad Amir Saeed: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization.

Jo Eidsvik: Supervision, Methodology, Writing—Review & Editing, Validation.

Antonio Candelieri: Supervision, Methodology, Writing—Review & Editing, Validation.

Code and data availability

The code and data supporting the findings of this study are available from the corresponding author upon reasonable request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author statement

All authors have contributed to the development of this work and the preparation of the manuscript. All authors have read and approved the final version of the manuscript.

Acknowledgements

The authors acknowledge support from the Centre for Geophysical Forecasting (CGF, grant no. 309960) at the Norwegian University of Science and Technology (NTNU). We are grateful to Elling Svee for providing help with the simulator used in this study. Muhammad Amir Saeed acknowledges the University of Milano-Bicocca for providing doctoral scholarship support, as well as the Erasmus+ Traineeship Program, which made this research possible. The authors also acknowledge

the Norwegian University of Science and Technology (NTNU) for providing support during the Erasmus mobility period.

References

- Alqahtani, A., He, X., Yan, B., and Hoteit, H. (2023). Uncertainty analysis of CO₂ storage in deep saline aquifers using machine learning and Bayesian optimization. *Energies*, 16:1684.
- Ament, S., Daulton, S., Eriksson, D., Balandat, M., and Bakshy, E. (2023). Unexpected improvements to expected improvement for Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 36. Curran Associates.
- Andersen, O. A. (2025). Topography-based surface water modeling in Julia, with support for infiltration and temporal developments. *Journal of Open Source Software*, 10(109):7785.
- Bandilla, K. W., Celia, M. A., and Leister, E. (2014). Impact of model complexity on CO₂ plume modeling at Sleipner. *Energy Procedia*, 63:3405–3415. 12th International Conference on Greenhouse Gas Control Technologies, GHGT-12.
- Bartoli, N., Lefebvre, T., Lafage, R., Saves, P., Diouane, Y., Morlier, J., Bussemaker, J. H., and Donelli, G. (2025). Multi-objective Bayesian optimization with mixed-categorical design variables for expensive-to-evaluate aeronautical applications. *arXiv preprint arXiv:2504.09930*. <https://arxiv.org/abs/2504.09930>.
- Callioli Santi, A., Ringrose, P., Eidsvik, J., and Haugdahl, T. A. (2025). Invasion percolation Markov chains – a probabilistic framework for assessing vertical CO₂ migration. *International Journal of Greenhouse Gas Control*, 142:104338.
- Candelieri, A., Ponti, A., and Archetti, F. (2024). Fair and green hyperparameter optimization via multi-objective and multiple information source Bayesian optimization. *Machine Learning*, 113:2701–2731.
- Carruthers, D. J. (2003). AAPG/Datapages discovery series no. 7, multidimensional basin modeling, chapter 3: Modeling of secondary petroleum migration using invasion percolation techniques. *AAPG/Datapages Discovery Series*.
- Daulton, S., Balandat, M., and Bakshy, E. (2020). Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 9851–9864. Curran Associates.
- Daulton, S., Balandat, M., and Bakshy, E. (2021). Parallel Bayesian optimization of multiple noisy objectives with expected hypervolume improvement. In *Advances in Neural Information Processing Systems*, volume 34, pages 2187–2200. Curran Associates.
- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197.
- Emmerich, M., Beume, N., and Naujoks, B. (2005). An EMO algorithm using the hypervolume measure as selection criterion. *Evolutionary Computation*, 15(3):411–436.
- Equinor / CO₂DataShare (2019). Sleipner 2019 benchmark model. <https://co2datashare.org/dataset/sleipner-2019-benchmark-model>. Sleipner CO₂ Reference Dataset. Open dataset.

- Fotias, S. P., Ismail, I., and Gaganis, V. (2024). Optimizing CO₂ injection rates in CCS: A Bayesian approach for realistic business efficiency. In *Proceedings of the GET 2024 Conference*. EAGE. Paper 204.
- Garnett, R. (2023). *Bayesian optimization*. Cambridge University Press.
- Lie, K.-A. (2019). *An introduction to reservoir simulation using MATLAB/GNU Octave: User guide for the MATLAB Reservoir Simulation Toolbox (MRST)*. Cambridge University Press.
- Liu, J., Liu, J., Zhu, Y., Sun, W., Zhang, D., and Pan, H. (2025a). Multi-objective optimization for efficient CO₂ storage under pressure buildup constraint in saline aquifer. *Applied Energy*, 382:125175.
- Liu, M., Wang, S., Zhang, K., Li, J., and Sun, H. (2023). Multiobjective optimization of CO₂ injection under geomechanical risk in high water cut oil reservoirs using artificial intelligence approaches. *Scientific Reports*, 13:11605.
- Liu, Y., Zhang, K., Chen, B., Zheng, F., and Qin, Z. (2025b). Deep learning–assisted multiobjective optimization of geological CO₂ storage performance under geomechanical risks. *SPE Journal*, 30(4):2073–2091.
- Loke, M. H. et al. (2025). Optimizing carbon capture and storage (CCS) design using multi-objective optimization and dynamic IPR-OPR nodal analysis. In *4th Carbon Capture and Storage Conference Asia Pacific*.
- Lu, X., Nakamura, T., Chowdhary, P., Bhatt, A., Zhou, Z., and Bhatt, S. (2022). Bayesian optimization for field-scale geological carbon storage. *Engineering*, 18:202–216.
- Nibbering, D. (2024). A high-dimensional multinomial logit model. *Journal of Applied Econometrics*, 39(3):481–497.
- Paria, B., Kandasamy, K., and Póczos, B. (2020). A flexible framework for multi-objective bayesian optimization with preference information. *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, USA.
- Ringrose, P. (2020). *How to store CO₂ underground: Insights from early-mover CCS Projects*. Springer.
- Saeed, M. A. and Candelieri, A. (2026). Bayesian optimization for categorical and mixed variables using a multinomial logit surrogate. *Algorithms*, 19(5):361.
- Saves, P., Diouane, Y., Bartoli, N., Lefebvre, T., and Morlier, J. (2023). A mixed-categorical correlation kernel for Gaussian process. *Neurocomputing*, 550:126349.
- Svee, E. (2026). CO₂BatchFill.jl. <https://github.com/ellingsvee/CO2BatchFill.jl>. Accessed: 2026-03-19.
- Zhu, M. and Bemporad, A. (2025). Global and preference-based optimization with mixed variables using piecewise affine surrogates. *Journal of Optimization Theory and Applications*, 204:26.

Table 3: Summary of notation for decision variables, model parameters, and objective functions used in the simulation-based CCS optimization framework.

Symbol	Description
X	Decision vector defining the injection strategy
scenario	Injection-well configuration (number of wells)
rate	CO_2 injection rate per well (Mt/yr)
L_{mid}	Layer index used during the mid-stage injection period
L_{top}	Layer index used during the late-stage injection period
\mathcal{X}	Feasible design space of admissible decisions
T	Total simulation time horizon (years)
Δt	Simulation time step (1 year)
t	Continuous simulation time
N_{stor}	Number of storage layers
i	Index of storage layer, $i = 1, \dots, N_{\text{stor}}$
A_i	Areal extent of storage layer i
ϕ_i	Porosity of layer i
$h_{i,\text{ref}}$	Reference thickness of layer i
p_{ref}	Reference pressure
$p_i(t)$	Pressure in layer i at time t
r	Vertical redistribution ratio controlling upward migration
d	Effective thickness scaling factor
$V_i(t)$	CO_2 volume stored in layer i at time t
$V_{\text{stored}}(t; X)$	Stored CO_2 volume at time t for design X
$V_{\text{tot}}(t)$	Total retained CO_2 volume across all layers at time t
$Q(t)$	Total CO_2 injection rate at time t
$V_{\text{inj}}(X)$	Total injected CO_2 volume accumulated over $[0, T]$
$f_1(X)$	Total stored CO_2 over the simulation horizon (maximize)
$f_2(X)$	Leakage proxy (minimize)
$f_3(X)$	Total project cost (minimize)
$\tilde{y}(X)$	Transformed objective vector $(f_1, -f_2, -f_3)$ used for maximization
C_{fixed}	Fixed project overhead (permitting, site preparation)
N_w	Number of injection wells
C_{well}	Cost per injection well
c_{cap}	Capture cost per unit CO_2 injected
c_{trans}	Transport cost per unit CO_2 injected
c_{inj}	Injection cost per unit CO_2 injected
c_m	Lumped marginal cost, $c_m = c_{\text{cap}} + c_{\text{trans}} + c_{\text{inj}}$
c_{mon}	Monitoring cost per unit time