energies

MDPI

Article

Ensemble Surrogates and NSGA-II with Active Learning for Multi-Objective Optimization of WAG Injection in CO_2-EOR

Yutong Zhu [1,2], Hao Li [3,*], Yan Zheng [4], Cai Li [1,2], Chaobin Guo [1,2] and Xinwen Wang [1,2,5]

[1] Chinese Academy of Geological Sciences, Beijing 100037, China; zhuyt2024@126.com (Y.Z.)
[2] Technology lnnovation Center for Carbon Sequestration and Geological Energy Storage, Ministry of Natural Resources, Beijing 100037, China
[3] Shaanxi Yanchang Petroleum (Group) Co., Ltd., Xi'an 710075, China
[4] Shaanxi Yanchang Petroleum (Group) Co., Ltd. Gas Field Company, Xi'an 716099, China
[5] School of Water Resources and Environment, China University of Geosciences (Beijing), Beijing 100083, China
* Correspondence: lihao007@sxycpc.com

Abstract

CO_2-enhanced oil recovery (CO_2-EOR) with water-alternating-gas (WAG) injection offers the dual benefit of boosted oil production and CO_2 storage, addressing both energy needs and climate goals. However, designing CO_2-WAG schemes is challenging; maximizing oil recovery, CO_2 storage, and economic returns (net present value, NPV) simultaneously under a limited simulation budget leads to conflicting trade-offs. We propose a novel closed-loop multi-objective framework that integrates high-fidelity reservoir simulation with stacking surrogate modeling and active learning for multi-objective CO_2-WAG optimization. A high-diversity stacking ensemble surrogate is constructed to approximate the reservoir simulator. It fuses six heterogeneous models (gradient boosting, Gaussian process regression, polynomial ridge regression, k-nearest neighbors, generalized additive model, and radial basis SVR) via a ridge-regression meta-learner, with original control variables included to improve robustness. This ensemble surrogate significantly reduces per-evaluation cost while maintaining accuracy across the parameter space. During optimization, an NSGA-II genetic algorithm searches for Pareto-optimal CO_2-WAG designs by varying key control parameters (water and CO_2 injection rates, slug length, and project duration). Crucially, a decision-space diversity-controlled active learning scheme (DCAF) iteratively refines the surrogate: it filters candidate designs by distance to existing samples and selects the most informative points for high-fidelity simulation. This closed-loop cycle of "surrogate prediction → high-fidelity correction → model update" improves surrogate fidelity and drives convergence toward the true Pareto front. We validate the framework of the SPE5 benchmark reservoir under CO_2-WAG conditions. Results show that the integrated "stacking + NSGA-II + DCAF" approach closely recovers the true tri-objective Pareto front (oil recovery, CO_2 storage, NPV) while greatly reducing the number of expensive simulator runs. The method's novelty lies in combining diverse stacking ensembles, NSGA-II, and active learning into a unified CO_2-EOR optimization workflow. It provides practical guidance for economically aware, low-carbon reservoir management, demonstrating a data-efficient paradigm for coordinated production, storage, and value optimization in CO_2-WAG EOR.

Keywords: CO_2-enhanced oil recovery; water-alternating-gas injection; multi-objective optimization; surrogate modeling; active learning

check for updates

Academic Editor: Albert Ratner

Received: 17 November 2025
Revised: 5 December 2025
Accepted: 15 December 2025
Published: 16 December 2025

Citation: Zhu, Y.; Li, H.; Zheng, Y.; Li, C.; Guo, C.; Wang, X. Ensemble Surrogates and NSGA-II with Active Learning for Multi-Objective Optimization of WAG Injection in CO_2-EOR. *Energies* **2025**, *18*, 6575. https://doi.org/10.3390/en18246575

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

In recent years, as the dual challenges of global climate change and energy security have intensified, carbon capture, utilization, and storage (CCUS) technologies have drawn increasing attention [1,2]. Among them, CO_2-enhanced oil recovery (CO_2-EOR), a key component of CCUS, not only boosts the recovery of mature fields but also enables subsurface storage of CO_2 [3]. Statistics indicate that CO_2-EOR accounts for roughly 20% of enhanced oil recovery projects worldwide and can deliver an additional 15–25% oil recovery. Meanwhile, studies have found that as much as 60% of the injected CO_2 can be trapped in the reservoir, demonstrating significant mitigation potential [4]. As a principal implementation of CO_2-EOR, water-alternating-gas (WAG) injection suppresses gas channeling and improves sweep efficiency [5–7]. Experimental and numerical results show that, in low-permeability reservoirs, carbonated water (CO_2-enriched water)–CO_2 alternating injection, compared with conventional CO_2 flooding or standard WAG, not only yields a markedly higher recovery but also achieves a higher CO_2 storage fraction, underscoring the synergy of CO_2-WAG in enhancing recovery while sequestering CO_2 [8,9].

In the CO_2-EOR/WAG domain, injection-scheme design is commonly posed as a simulation-constrained (often high-dimensional) optimization problem, in which operational controls (e.g., water/CO_2 injection rates, WAG timing and slug length, and project duration) are tuned to balance multiple conflicting objectives, such as oil recovery, CO_2 storage ratio, and economic performance. Recent surveys and reviews on intelligent reservoir optimization highlight that this class of problems shares core characteristics—strong nonlinearity, nonconvex constraints, and expensive forward simulation—and is closely related to broader well-placement/production-control optimization and history-matching workflows [10–13].

Methodologically, two optimization paradigms are the most frequently used. Model-based approaches (e.g., sequential quadratic programming and gradient/ensemble-gradient methods) can be efficient for continuous controls when reasonably smooth response surfaces or gradients are available [14–16]. In contrast, multi-objective evolutionary algorithms (MOEAs), such as NSGA-II and MOPSO, are widely adopted because they naturally handle nonconvexity and mixed discrete–continuous couplings while generating a diverse Pareto set in a single run [15–18]. These algorithms have been extensively applied to CO_2-EOR/CO_2-WAG co-optimization studies that balance oil recovery and CO_2 storage across different reservoir settings and operating conditions [14,19–21]. However, a global multi-objective search typically requires a large number of objective evaluations; when each evaluation calls a high-fidelity 3D simulator, the computational burden becomes the primary bottleneck.

To reduce simulator calls, surrogate-assisted optimization has become a central research direction. A wide range of surrogate models—spanning response surfaces, Gaussian process/kriging models, support vector regression, tree ensembles, and neural networks—have been embedded into optimization loops to approximate simulator outputs at a negligible cost [22–25]. Recent developments further explore (i) multi-fidelity and transfer-learning surrogates to leverage inexpensive low-fidelity simulations alongside high-fidelity data [26,27] and (ii) ensemble surrogates (e.g., stacking or selective ensembles) to hedge against model mis-specification and improve robustness across heterogeneous response behaviors [26,28–30]. Meanwhile, deep-learning surrogates that represent spatiotemporal flow patterns (e.g., U-Net/Transformer-based models and graph-network surrogates) are emerging for CO_2-EOR and CO_2 storage applications, enabling the rapid screening of well configurations and control strategies [31–33].

Because a surrogate trained on a fixed initial design may be inaccurate near optimal regions (and may suffer from a distribution shift as the optimizer explores new regions),

adaptive sampling and active learning strategies are increasingly used to iteratively refine surrogates during optimization [34–37]. Typical infill criteria include uncertainty, discrepancy, or hybrid infill rules designed to maximize the information gained from each new simulator run. This "surrogate prediction → high-fidelity correction → model update" paradigm is particularly important for multi-objective problems, where surrogate bias can distort dominance relations and yield false Pareto fronts.

Finally, there is a growing emphasis on techno-economic and low-carbon decision making. Economic metrics such as net present value (NPV) and CO_2-related costs/credits are essential for deployable CCUS-EOR projects but are still frequently treated as ex-post indicators or constraints, rather than coequal objectives, in multi-objective workflows [38]. Recent CO_2-WAG/CCUS studies show that explicitly including economic objectives (and, in some cases, emissions metrics) can materially change the recommended operating strategies compared to purely technical objectives [21,28,39,40]. To make the above landscape explicit and to clarify the research gaps, Table 1 compares mainstream approaches for simulation-constrained optimization of CO_2-EOR/CO_2-WAG schemes.

Table 1. Comparison of mainstream approaches for optimization and the associated research gaps.

Approach Family	Typical Workflow	Pros	Limits/Gaps	Best for
Full-physics, model-based opt.	Full-physics, model-based opt.	Full-physics, model-based opt.	Full-physics, model-based opt.	Full-physics, model-based opt.
Direct sim-based MOEA	NSGA-II/MOPSO evaluates simulator per candidate	Handles nonconvexity; diverse Pareto set	Very high eval budget (esp. 3D)	Benchmarks/ big compute
Static SA-MOEA (single surrogate)	Offline DOE → train 1 proxy → MOEA on proxy	Large speed-up	Proxy drift near opt.; OOD errors; false Pareto risk	Early screening/ simple cases
Ensemble/MF SA-MOEA	Multi-proxy or MF (stack/select/ transfer) + MOEA	More robust; better bias–variance	Still coverage-limited; ensemble diversity often not explicit	Strong nonlinearity; multi-output objs
Closed-loop surrogate + AL	Proxy opt. → infill → simulate → update	Data-efficient; accuracy near Pareto	Infill clustering; weak decision-space coverage	Tight sim budgets/ expensive sims
Techno-economic/ low-carbon tri-obj.	Joint: recovery–storage–NPV/(emissions)	Decision-relevant; avoids uneconomic "optima"	Econ uncertainty; often simplified/post hoc	Field planning/ CCUS investments

Table note: opt. = optimization; hi-fi = high-fidelity; sim = simulator/simulation; grad = gradient-based; EG = ensemble gradient; MO = multi-objective; SQP = sequential quadratic programming; MOEA = multi-objective evolutionary algorithm; SA-MOEA = surrogate-assisted MOEA; MF = multi-fidelity; DOE = design of experiments; proxy = surrogate model; opt. = optimum/optimal; OOD = out-of-distribution; AL = active learning; infill = adaptive sampling points; obj(s) = objective(s); NPV = net present value; CCUS = carbon capture, utilization, and storage; EOR = enhanced oil recovery; WAG = water-alternating-gas.

Based on Table 1, three gaps are particularly relevant for practical CO_2-WAG design under tight simulation budgets: (G1) many surrogate-assisted workflows still rely on a single proxy or do not explicitly manage model diversity, making multi-objective predictions vulnerable to local bias and out-of-distribution errors; (G2) common active-learning criteria emphasize uncertainty or objective-space improvement and can produce clustered samples, leading to inefficient decision-space coverage; and (G3) economic feasibility is not consistently integrated, since the NPV is still often omitted as a peer objective in tri-objective optimization. These gaps motivate a closed-loop framework that combines a high-diversity stacking surrogate, global Pareto search, and explicitly diversity-controlled active learning, while treating the NPV as a coequal objective.

To address these challenges, this study proposes a closed-loop multi-objective optimization framework for the design of typical CO_2-WAG schemes. Under a limited high-fidelity simulation budget, the framework simultaneously optimizes three objectives—oil recovery factor, CO_2 storage ratio, and NPV—using four control variables as decision parameters: water injection rate, CO_2 injection rate, single-cycle WAG duration (slug

length), and total project duration. A global search is performed with NSGA-II. The main contributions are as follows: (1) High-diversity stacking ensemble surrogate. Six heterogeneous base learners—gradient boosting decision trees (GBDT), Gaussian process regression with a Matérn kernel (GPR), second-order polynomial ridge regression, k-nearest neighbors (KNN), generalized additive models (GAM), and radial basis function kernel SVR (RBF-SVR)—are linearly fused by a ridge-regression meta-learner. The original decision variables are explicitly included alongside base-learner outputs to enhance robustness and interpretability. Scheme-grouped fivefold cross-validation (Group K-Fold) and out-of-fold predictions are used to assess generalization and suppress CV information leakage caused by temporal dependence. (2) Efficient NSGA-II coupling with convergence-based termination. The stacking surrogate is embedded into NSGA-II fitness evaluations, using a population size of p = 100 (=25 $\times$ the number of decision variables) to ensure adequate diversity and resolution in the 4D decision space. Instead of terminating NSGA-II after a fixed number of generations, we adopt a convergence-based stopping rule (stagnation of the Pareto set), while keeping $G_{\max}$ = 100 only as a conservative upper bound. SBX crossover and polynomial mutation are used to maintain exploration and population diversity. (3) Diversity-controlled active filtering (DCAF). A two-stage sampling strategy is proposed: first, an "acceptable distance loss" criterion filters out candidate points that are too close to existing samples; second, from the remaining candidates, truly isolated and highly representative points are selected for high-fidelity evaluation. This strategy increases coverage and information gain, enabling a targeted exploration of model-sensitive regions with limited new data and forming a closed loop of "surrogate prediction–high-fidelity correction–model update." (4) Numerical validation. The framework is validated on the public SPE5 light-oil reservoir model under CO_2-WAG conditions. Experimental design and performance evaluation are conducted over the four control variables and three objectives. Comparisons against the true simulated Pareto front and conventional baselines demonstrate that the proposed "stacking + NSGA-II + DCAF" framework more closely approximates the true Pareto front while substantially reducing the number of high-fidelity simulator evaluations. The results provide guidance on parameter ranges for CO_2-WAG design and quantitative decision support for balancing production and carbon-mitigation benefits.

Overall, the proposed method offers an intelligent design paradigm for optimizing CO_2-EOR/WAG injection, achieving coordinated improvements in oil recovery, CO_2 storage, and economic performance, and furnishing theoretical and methodological support for related engineering practice.

2. Methodology

This section follows the overall workflow of "stacking ensemble surrogate $\rightarrow$ multi-objective evolutionary optimization $\rightarrow$ active-learning sampling $\rightarrow$ numerical validation". We first build a high-accuracy, diversity-aware stacking surrogate (Section 2.1), then perform three-objective global optimization using NSGA-II (Section 2.2). In the subsequent sections, we introduce a decision–space–diversity-oriented active-learning strategy, DCAF (Section 2.3), and validate the framework on a reservoir case (Section 2.4).

2.1. Stacking Ensemble Surrogate Modeling

High-fidelity reservoir simulation can accurately capture multiphase flow and phase behavior during WAG development; however, frequent simulator calls in multi-objective optimization are prohibitively expensive and do not scale to large, long-horizon parameter searches. To this end, we approximate the simulator responses with machine learning surrogates and adopt a stacking ensemble strategy that reduces the cost of each fitness

evaluation while maintaining high predictive accuracy and robustness across the parameter space. Stacking uses a meta-learner to learn from the predictions of multiple base models, exploiting complementary error patterns to reduce variance and thereby improve overall accuracy and robustness.

(1) Model portfolio and diversity design. Guided by a "low-bias–high-capacity" full-spectrum philosophy, we construct a pool of 14 base learners spanning linear/polynomial models, kernel and basis-function methods, instance-based methods, tree ensembles, generalized additive models (GAM), and feedforward neural networks (ANN), thus covering major model families from parsimonious linear structures to high-capacity nonlinear approximators. Tree-based models capture strong nonlinearities and high-order interactions; the kernel and nearest-neighbor methods emphasize local complexity; linear/polynomial models and GAM strike a balance between interpretability and fit; and ANN provides a universal approximation capability [31–33]. This heterogeneous pool supplies rich and complementary "modular building blocks" for subsequent stacking [26,29,33]. For intuition, schematic diagrams of representative model families are provided in Figures 1–6.

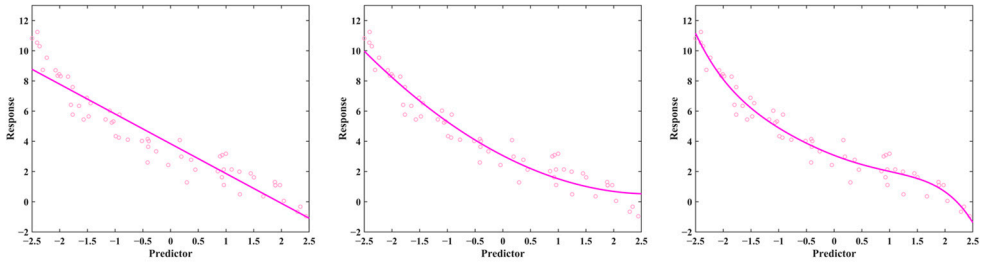


Figure 1. Schematic illustration of linear/polynomial regression.

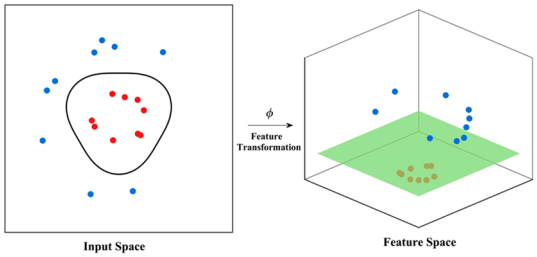


Figure 2. Kernel method schematic. Dots in different colors represent samples with different response levels/classes. In the original input space (2D), these samples are not directly separable. After a kernel-induced feature transformation to a higher-dimensional space, the classes become separable; the green diamond-shaped surface represents the separating decision surface (hyperplane) used to discriminate different response levels/classes in the transformed feature space.

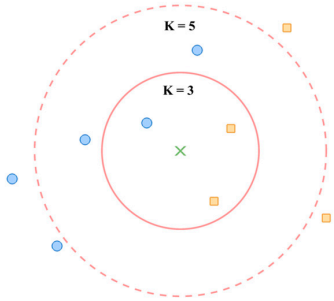


Figure 3. KNN neighborhood schematic. Training samples are shown with different colors (and symbols) to indicate different response levels/classes. The green marker denotes the query sample to be predicted. The circle centered at the query sample indicates the local neighborhood used to identify the k nearest neighbors. Predictions are obtained by (distance-weighted) averaging of the neighbors for regression, or by majority voting for classification.

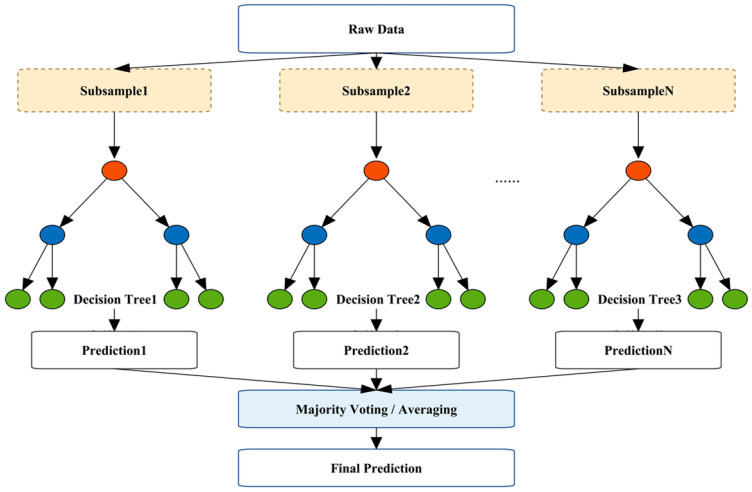


Figure 4. Tree ensemble schematic. Node colors indicate different node types/levels in the trees: orange circles denote root nodes, blue circles denote internal split (decision) nodes, and green circles denote terminal leaf nodes that output predictions.

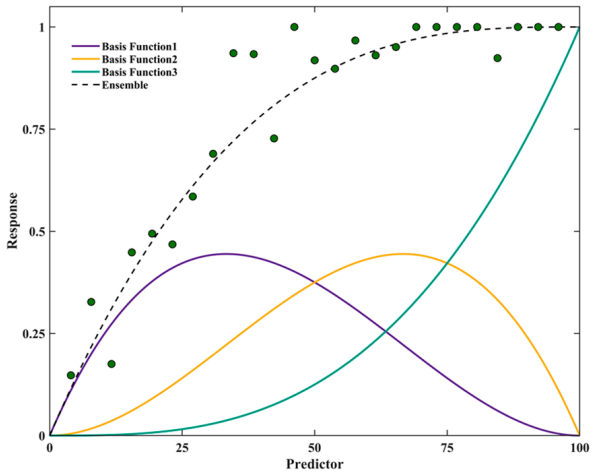


Figure 5. GAM schematic. Dots denote the training observations (or sample points/partial residuals) used to fit each smooth term. Colored smooth curves represent the fitted univariate functions $f_j\left(x_j\right)$ for different predictors. The GAM prediction is obtained by summing all smooth components together with an intercept term.

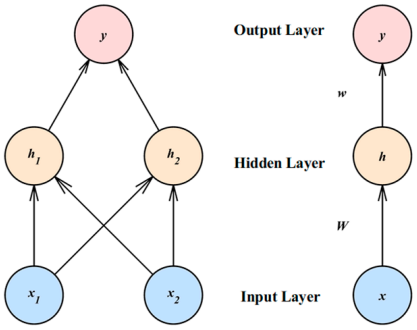


Figure 6. Feedforward neural network schematic. The letters denote the standard network components: x_{i} are the input features (input-layer neurons), h_{i} are hidden-layer neurons, and y is the network output. ω and W denote connection weights. Arrows indicate the forward propagation of information through weighted connections from one layer to the next.

Linear/polynomial models (e.g., second-order polynomial ridge regression, Poly2Ridge). Linear regression assumes an (approximate) linear relation between the response, y, and the feature vector, x:

$$y = \beta_0 + \beta^T x + \varepsilon \tag{1}$$

Polynomial regression first applies a feature expansion $\phi(x)$ (e.g., quadratic terms and interactions) and then performs linear regression in the expanded feature space. Polynomial ridge adds L2 regularization:

$$\min_{\beta} \parallel y - \Phi\beta \parallel_2^2 + \lambda \parallel \beta \parallel_2^2 \tag{2}$$

These models are stable and interpretable (often higher bias but lower variance), which helps to capture smooth global trends and provides a low-variance "skeleton" for ensembling (Figure 1).

The kernel and basis-function methods (e.g., Matérn kernel GPR, RBF-SVR). A basis-function model represents the response as a linear combination of nonlinear bases $\{\phi_j\}$:

$$\hat{y}(x) = \sum_{j=1}^{J} w_j \phi_j(x) \tag{3}$$

Kernel methods avoid explicitly constructing ϕ by using a kernel

$$k(x, x') = \langle \phi(x), \phi(x') \rangle \tag{4}$$

enabling flexible nonlinear regression under small-sample regimes. In this work, Matérn kernel GPR provides a high-accuracy backbone, while RBF-SVR adds complementary locality control (Figure 2).

Instance-based learning (K-nearest neighbors, KNN). KNN is a "lazy learner" that predicts by local averaging in feature space. For a query x, let $\mathcal{N}_K(x)$ be its K nearest neighbors; a typical regression form is as follows:

$$\hat{y}(x) = \frac{1}{K} \sum_{i \in \mathcal{N}_K(x)} y_i \tag{5}$$

(or a distance-weighted average). KNN captures irregular local structures in dense regions and often complements smooth/global models (Figure 3).

Tree models and ensembles (Random Forest, GBDT). A decision tree recursively partitions the feature space into regions and fits piecewise models. Random Forest reduces variance by bagging, whereas gradient boosting constructs an additive model:

$$F_t(x) = F_{t-1}(x) + \eta\, h_t(x) \tag{6}$$

where h_t is a weak learner, fitted to the residual/negative gradient. Tree ensembles are strong at capturing nonlinear interactions in tabular data; here, GBDT serves as a key accuracy backbone (Figure 4).

Generalized additive models (GAM). GAM extend generalized linear models by replacing linear terms with smooth univariate functions:

$$g(\mathbb{E}[y \mid x]) = \beta_0 + \sum_j f_j(x_j) \tag{7}$$

where f_j are spline-based smoothers. GAM balance flexibility and interpretability, contributing structured, regularized nonlinearity to the ensemble (Figure 5).

Feedforward neural networks (ANN). A feedforward network composes linear transforms with nonlinear activations. For a one-hidden-layer network:

$$\hat{y}(x) = \omega\sigma(W_x + b_1) + b_2 \tag{8}$$

and deeper networks iterate this composition. ANN offer a high capacity but are more data- and tuning-sensitive, so they are included in the pool as high-capacity candidates (Figure 6).

Overall, these learners span the bias–variance and interpretability–capacity spectra and thus provide the diversity that stacking can exploit: when base learners exhibit complementary (weakly correlated) errors, the meta-learner can combine them to reduce variance and correct systematic biases.

(2) Dataset construction and grouped cross-validation. The surrogate takes the four decision variables (see Section 2.4) as inputs together with time—namely, water-injection rate, CO_2-injection rate, WAG half-cycle duration, total WAG duration, and time—and outputs the oil recovery factor, CO_2 storage ratio, and net present value (NPV). We generate 100 development schemes via Latin hypercube sampling and record responses over a 10-year forecast horizon with a 0.5-year time step, yielding 2000 time-stamped samples [34]. To avoid information leakage between the training and test sets, we adopt fivefold Group K-Fold cross-validation, grouping all time-series samples from the same scheme into the same fold. Features in the training set are standardized using z-score normalization, and the same transform is applied to the test folds. For each of the three targets, we train single-output surrogates for all 14 base learners, yielding 14×3 base models; their out-of-fold (OOF) predictions are concatenated into a 2000-row sequence to provide unbiased estimates of generalization performance and to serve as the foundation for stacking and error analysis.

(3) Single-model performance and base-learner selection. On the full OOF set, we compute RMSE, MAE, R^2, NRMSE, MAPE, and SMAPE for each target–learner pair and summarize the mean and standard deviation of RMSE, MAE, and R^2 across folds to characterize accuracy and stability; we also measure training and inference time to quantify efficiency. Results show that GBDT attains the highest accuracy for oil recovery and CO_2 storage ($R^2 \approx 0.99$, NRMSE $\approx$ 2–3%), while Matérn kernel GPR performs best for NPV ($R^2 \approx 0.987$, NRMSE $\approx$ 2.3%). A Friedman test indicates that the average rankings of GBDT and Matérn-GPR are significantly better than those of linear models, ANN, and simple kernel regressors, with p-values well below 0.01, confirming statistically significant differences in RMSE among learners. To further quantify coverage of "hard samples", we compute a sample-level win rate; GBDT and Matérn-GPR rank among the top across all three targets (Figures 7–9), whereas KNN and Poly2Ridge—though slightly lower in overall ranking—exhibit high win rates in localized regions, revealing important complementarity. A 14×14 residual–correlation matrix constructed from OOF residuals (Figures 10–12) shows that residuals within the tree-ensemble family are highly correlated—indicating near-redundant error structures—whereas Poly2Ridge, KNN, and GAM display low residual correlation with the tree models and with GPR, providing the structural diversity that is crucial for ensembling.

Balancing accuracy, stability, diversity, and efficiency, we adopt a three-stage procedure to select base learners: first, we remove learners with inadequate overall accuracy or negative R^2 on any target, based on Friedman rankings and win rates; second, prune tree and kernel models are highly redundant with GBDT and Matérn-GPR, according to residual correlations, so we retain only representative backbone models; and third, when performance is similar, we prefer learners with a lower computational cost and smaller CV variability, keeping the ensemble to 5–7 base learners to avoid over-complexity.

The final stacking set comprises GBDT, Matérn kernel GPR, Poly2Ridge, KNN, GAM, and RBF-kernel SVR. In this composition, GBDT and Matérn-GPR provide the accuracy backbone; Poly2Ridge and GAM contribute a low-bias, interpretable structure; and KNN and RBF-SVR strengthen the representation of locally complex patterns.

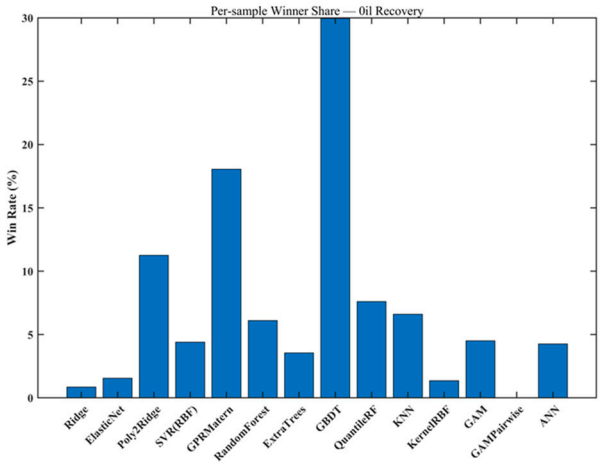


Figure 7. Win rate of each model in predicting oil recovery.

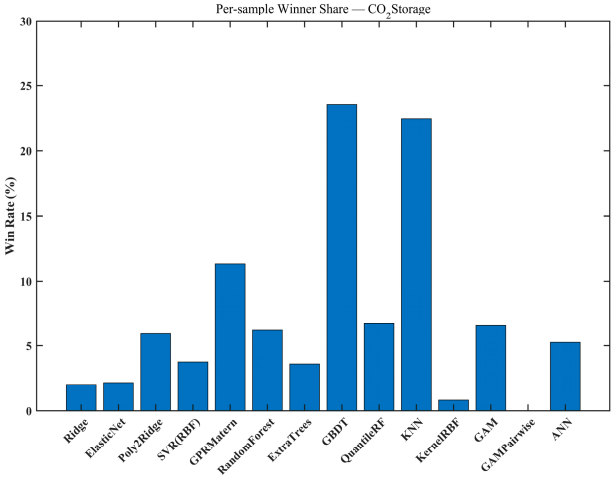


Figure 8. Win rate of each model in predicting CO_2 storage.

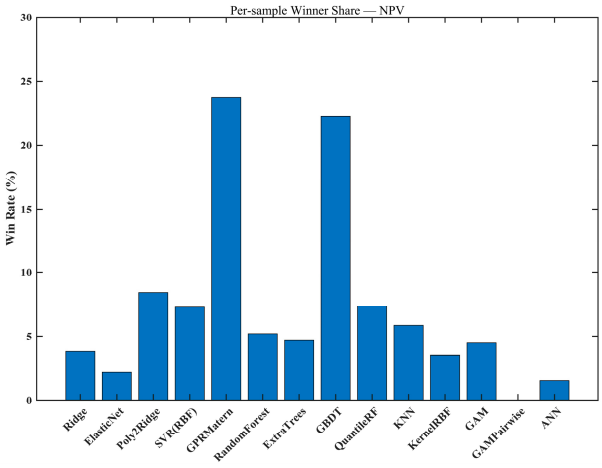


Figure 9. Win rate of each model in predicting NPV.

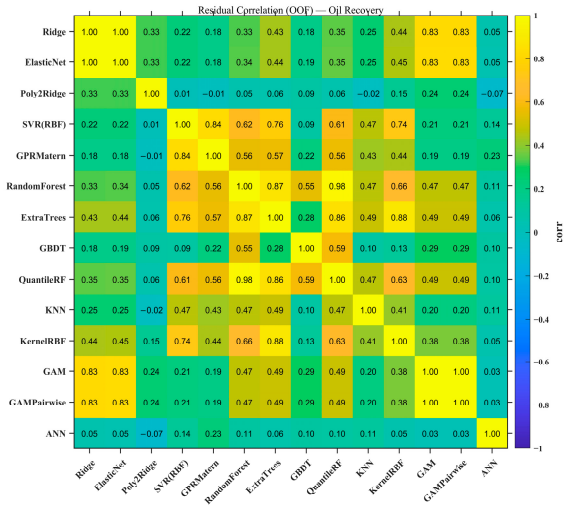


Figure 10. Residual correlation matrix for oil recovery predictions from each model.

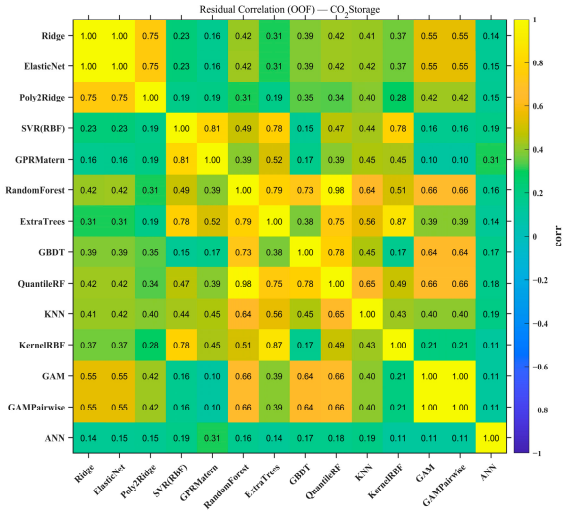


Figure 11. Residual correlation matrix for CO_2 storage predictions from each model.

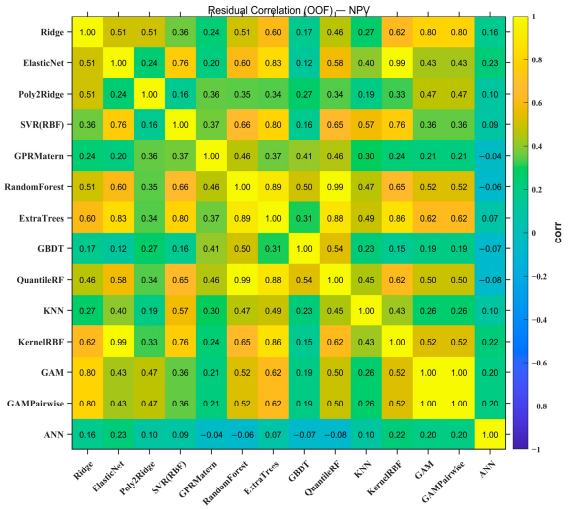


Figure 12. Residual correlation matrix for NPV predictions from each model.

(4) Stacking architecture and a "two-phase + cyclic" workflow. Stacking (stacked generalization) is a layered ensemble in which first-layer heterogeneous base learners $\{f_m\}$ predict responses from the original features, and a second-layer meta-learner, g, learns how to combine these predictions, based on their out-of-sample behavior [41,42] (Figure 13).

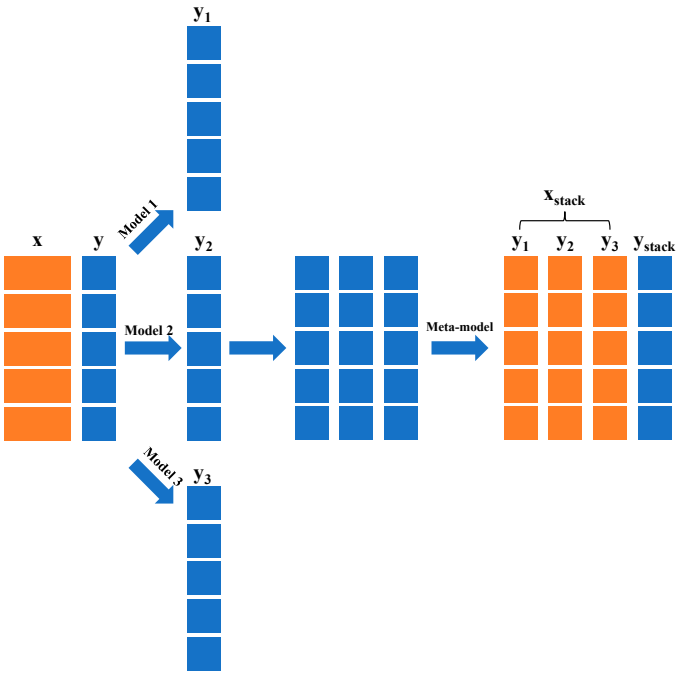


Figure 13. Two-level stacking architecture schematic. Colors distinguish inputs and predictions: orange indicates the input data x fed into the first-level base learners (original features), whereas blue indicates the corresponding predicted values produced by the base learners. The second-level meta-learner is trained using the first-level predictions as meta-features, fuses these meta-features, and outputs the final prediction.

Let the true response at sample i be $y_i = y(x_i)$. The m-th base learner yields

$$\hat{y}_{m,i} = f_m(x_i) = y_i + \varepsilon_{m,i} \tag{9}$$

Stacking forms a meta-feature vector

$$z_i = [\hat{y}_{1,i}, \ldots, \hat{y}_{M,i}, x_i] \tag{10}$$

and trains a meta-learner

$$\hat{y}_i = g(z_i) \tag{11}$$

If g is linear ridge regression, then

$$\hat{y}_i = w_0 + \sum_{m=1}^{M} w_m \hat{y}_{m,i} + w_x^\top x_i \tag{12}$$

with the regularized objective

$$\min_w \sum_i (y_i - \hat{y}_i)^2 + \lambda \parallel w \parallel_2^2 \tag{13}$$

The ridge penalty stabilizes the combination and prevents extreme weights; the key gain comes from exploiting complementary error patterns: base learners with a low residual correlation provide more independent information, enabling the meta-learner to reduce variance and correct systematic bias.

To avoid information leakage, the meta-learner is trained on out-of-fold (OOF) predictions. In K-fold CV ($K = 5$), each base model is trained on $K - 1$ folds and predicts the held-out fold; concatenating held-out predictions across folds yields an OOF prediction for every sample, which then becomes the second-layer input (Figure 14). In our time-stamped

data, this is implemented as a fivefold Group K-Fold so that all time steps from the same development scheme stay within the same fold.

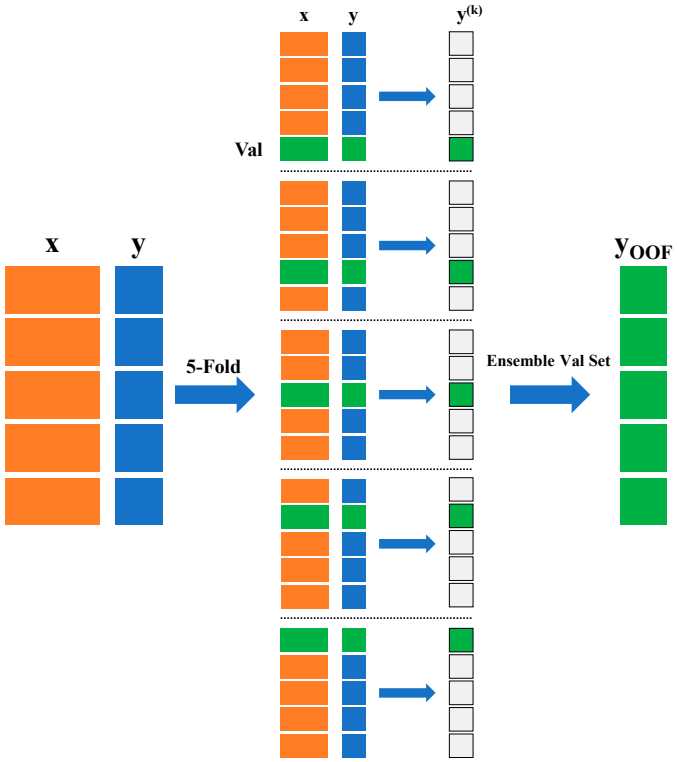


Figure 14. Fivefold OOF training workflow for stacking. Orange denotes the model input features x, and blue denotes the target values y. In each fold, the base learners are trained on the training folds and then predict the held-out validation fold (green) to generate out-of-fold (OOF) predictions. Concatenating OOF predictions across all five folds yields the complete meta-feature matrix/vector (green), which is subsequently used to train the second-level meta-learner.

Within the stacking framework, the six base learners make predictions for the three targets over the same feature space; their fivefold OOF predictions, together with the raw decision variables, are fed into the meta-learner. The meta-learner is ridge regression, which linearly fuses base-model outputs while explicitly retaining the main effects of the decision variables. L_2 regularization controls complexity, mitigates overfitting, and adaptively allocates weights to base learners. Passing the raw decision variables directly to the meta-learner enhances the physical interpretability and out-of-distribution robustness [43]: systematic local biases can be corrected by linear/polynomial models or KNN, whereas regions with strong nonlinearity and interactions are dominated by GBDT and Matérn-GPR.

In application, we use a "initial training → evolutionary optimization → active-learning correction" two-phase cyclic strategy. We first perform high-fidelity simulations for 100 Latin-hypercube samples to build the initial training set, train the 14 base learners and the stacking meta-learner, and compute OOF evaluations. NSGA-II then generates candidate solutions for which the stacking surrogate rapidly predicts the multi-objective responses; combined with DCAF, representative points in decision space are selected for new high-fidelity simulations. The newly acquired samples are incrementally merged into the training set and the surrogate is updated. The iterative stopping criterion is that, for five consecutive generations, the mean prediction error for each objective falls below 1% and the maximum error below 3%. This closed loop—surrogate prediction → high-fidelity correction → model update—substantially reduces simulator calls while progressively improving approximation to the true Pareto front, and it can be readily shifted to accommodate changes in reservoir conditions or development objectives.

2.2. Multi-Objective Optimization Algorithm (NSGA-II)

Building on the stacking surrogate of Section 2.1, we employ the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) [44] to conduct three-objective global optimization of reservoir development schemes and to elucidate the trade-offs among oil recovery, CO_2 storage, and NPV. Through fast non-dominated sorting, crowding distance, tournament selection, and elitism, NSGA-II preserves the population diversity while driving the Pareto front toward high-quality regions [15–17,45]. In this work, the decision vector contains D = 4 optimized control variables (Section 2.4); accordingly, the population size is selected to scale with D, and we set p = 100 (=25D) to provide sufficient sampling resolution in the 4-D decision space and to maintain a well-spread tri-objective nondominated set under the crowding-distance mechanism. Because fitness evaluations are performed using the stacking surrogate, this population size remains computationally affordable, while the expensive high-fidelity runs are controlled by the proposed DCAF sampling strategy. We use simulated binary crossover (SBX; crossover probability 0.8) and polynomial mutation (mutation probability 0.2, corresponding to roughly one mutation per individual, on average). Importantly, the NSGA-II run is not terminated by a predetermined generation count alone: the evolution is stopped once the Pareto set shows no significant improvement for several consecutive generations (e.g., negligible change in a Pareto quality indicator, such as hypervolume), with $G_{\max}$ = 100 imposed only as a safeguarding upper bound. In the surrogate-assisted closed-loop optimization of this study, the overall workflow terminated earlier (generation 27), when the predictive errors stabilized according to the accuracy-based convergence criterion described in Section 2.1. Embedding the stacking surrogate into the NSGA-II fitness evaluations enables the efficient discovery of solution sets that perform well across all three objectives and are evenly distributed—providing a robust basis for subsequent trade-off analysis and scheme selection [46,47].

2.3. Diversity-Controlled Active Sampling in Decision Space

Within the "stacking surrogate + NSGA-II" optimization framework, we adopt active learning to incrementally update the surrogate under a limited budget of high-fidelity simulations/experiments [22,35–37]. If one relies solely on random sampling or a single uncertainty-based rule, candidate points tend to cluster locally in decision space, producing many near-duplicate neighbors that waste the evaluation budget and weaken coverage of unexplored regions. To address this, we propose DCAF (diversity-controlled active filtering), an active-learning strategy tailored to decision-space diversity: we first remove redundant candidates via distance-based filtering and then perform outlier retention among the deleted points to "rescue" a very small number of truly isolated candidates that may be globally informative. The result is a high-fidelity evaluation set that balances diversity with the protection of rare but important samples [37].

Let the decision space be $X \subset \mathrm{R}^n$, with $n = 4$ in this study. For any two design vectors, $x_i, x_j \in X$, define the Euclidean distance:

$$d(x_i, x_j) = \sqrt{\sum_{k=1}^{n} \left(x_i^{(k)} - x_j^{(k)} \right)^2} \tag{14}$$

(1) Redundancy removal by distance filtering with acceptable loss.

In active-learning round t, let P_t be the candidate set produced by NSGA-II and L the set of already evaluated (labeled) high-fidelity samples. The goal is to remove those candidates that are too close to L (and thus provide limited information gain) from P_t,

while retaining candidates that are sufficiently separated from the existing samples. For any $x \in P_t$, define its minimum distance to the labeled set as follows:

$$d_{\min}(x) = \min_{y \in L} d(x, y) \tag{15}$$

Smaller $d_{\min}(x)$ implies less incremental information. Introduce a distance threshold, D_t, and apply the following rule: $d_{\min}(x) < D_t \Rightarrow$ delete x (redundant), $d_{\min}(x) \geq D_t \Rightarrow$ keep x.

To avoid ad hoc choices of D_t, we determine it adaptively via an "acceptable distance-loss" parameter $\beta \in (0,1)$. Let $\mid P_t \mid = N_t$. For each candidate, $x^{(i)} \in P_t$, compute:

$$\delta_i = d_{\min}\left(x^{(i)}\right), i = 1, \ldots, N_t \tag{16}$$

and form the minimum distance sequence $\left\{\delta_i\right\}_{i=1}^{N_t}$. Sort it in ascending order:

$$\delta_{(1)} \leq \delta_{(2)} \leq \cdots \leq \delta_{(N_t)} \tag{17}$$

Let the total "distance mass" be $S_t^d = \sum_{i=1}^{N_t} \delta_{(i)}$. Define the cumulative share contributed by deleting the k closest samples as follows:

$$R_k = \frac{\sum_{i=1}^{k} \delta_{(i)}}{S_t^d}, k = 1, \ldots, N_t \tag{18}$$

Here, R_k quantifies the fraction of total distance mass removed if the threshold, D_t, is placed between $\delta_{(k)}$ and $\delta_{(k+1)}$, i.e., if $\left\{\delta_{(1)}, \ldots, \delta_{(k)}\right\}$ are deleted.

Given an acceptable loss, β(e.g., $\beta = 5\%$), find the following: $k_t^* = \max\{k \mid R_k \leq \beta\}$. If $1 \leq k_t^* < N_t$, deleting the first k_t^* nearest neighbors keeps the distance loss within the budget, β, whereas deleting one more would exceed it. Set the threshold within the following: $D_t \in \left(\delta_{(k_t^*)}, \delta_{(k_t^*+1)}\right)$, and, in this work, take the midpoint

$$D_t = \frac{\delta_{(k_t^*)} + \delta_{(k_t^*+1)}}{2} \tag{19}$$

rounded to the nearest 0.001 to align the numerical implementation with manuscript reporting. In the edge case, $R_1 > \beta$, let $k_t^* = 0$ (i.e., the candidates are already relatively dispersed with respect to L), and set $D_t = \delta_{(1)}$.

Partition the candidates into the retained and redundant subsets:

$$P_t^{\text{keep}} = \{x \in P_t \mid d_{\min}(x) \geq D_t\}, P_t^{\text{del}} = \{x \in P_t \mid d_{\min}(x) < D_t\} = P_t \setminus P_t^{\text{keep}} \tag{20}$$

Without outlier retention, all points in P_t^{keep} would be sent for high-fidelity evaluation in round t. In our case study, with $\beta = 5\%, 10\%, 20\%$, distance filtering removes approximately 26.1%, 36.6%, and 51.2% of candidates, respectively, effectively suppressing near-neighbor redundancy.

(2) Outlier retention based on the global candidate distribution and thresholding.

Minimum-distance filtering systematically removes candidates that are too close to L, but the deleted set P_t^{del} may contain a very small number of globally isolated candidates that are informative for the overall structure. Such points may be close to some labeled samples (small $d_{\min}$) yet lie in a low-density region of the current candidate distribution—near a cluster boundary or the entrance to a new region. To avoid discarding these, DCAF

performs outlier retention after distance filtering, but only within P_t^{del}. If no deleted sample meets the outlier threshold, the rescue set is empty and no extra samples are added.

We proceed as follows. First, we define an isolation (local density) indicator using both historical and current information. Let $\mathcal{B}_t = L \cup P_t$ be the reference set for the neighborhood search. For any $x \in P_t$, let $N_k(x)$ be its k nearest neighbors in B_t (with a preset k). Define the local mean neighbor distance:

$$\overline{d}_k(x) = \frac{1}{k} \sum_{y \in \mathcal{N}_k(x)} d(x, y) \tag{21}$$

Larger $\overline{d}_k(x)$ indicates greater isolation of x in the joint set B_t, suggesting lower sample density and a higher chance that x lies in an under-explored region or near an extreme boundary.

Next, compute $\overline{d}_k(x^{(i)})$ for all $x^{(i)} \in P_t$ (so $\mid P_t \mid = N_t$) and sort in ascending order:

$$\eta_{(1)} \leq \eta_{(2)} \leq \cdots \leq \eta_{(N_t)} \tag{22}$$

Introduce an outlier proportion parameter $\gamma \in (0, 1)$, representing the maximum fraction (about γN_t) of candidates we are willing to flag as "isolated outliers". Define the index:

$$q_t = \lceil (1 - \gamma)\, N_t \rceil \tag{23}$$

and the outlier threshold: $D_{out,t} = \eta_{(q_t)}$.

Statistically, $D_{\text{out},t}$ is the $(1 - \gamma)$ quantile of the isolation distribution over P_t; points with $\overline{d}_k(x) \geq D_{\text{out},t}$ lie in the top γ tail of isolation. Outlier retention is then applied only to the deleted set:

$$A_t = \left\{ x \in P_t^{\text{del}} \,\middle|\, \overline{d}_k(x) \geq D_{out,t} \right\} \tag{24}$$

The final high-fidelity sampling set for round t is $S_t = P_t^{\text{keep}} \cup A_t$.

If no deleted sample satisfies $\overline{d}_k(x) \geq D_{\text{out},t}$, then $A_t = \emptyset$ and $S_t = P_t^{\text{keep}}$. In practice, β and γ can be tuned so that $\mid S_t \mid$ matches the available evaluation budget: β controls the strength of the redundancy filtering, while γ controls the fraction of rescued isolated points among the deleted candidates.

(3) DCAF workflow and integration.

Combining distance filtering with outlier retention, the DCAF procedure in each iteration comprises the following:

(1) Inputs: Current candidate set P_t (generated by NSGA-II using the surrogate from the previous round), labeled set L, acceptable distance-loss β, outlier proportion γ, neighbor count k, and other hyperparameters.

(2) Distance filtering: compute $d_{\min}(x)$ for all $x \in P_t$; determine D_t from β via the cumulative-loss statistic; partition P_t into P_t^{keep} and P_t^{del}.

(3) Outlier retention: compute $\overline{d}_k(x)$ on P_t; determine the global outlier threshold $D_{\text{out},t}$; construct the rescue set $A_t = \left\{ x \in P_t^{\text{del}} \mid \overline{d}_k(x) \geq D_{\text{out},t} \right\}$; define the final sampling set $S_t = P_t^{\text{keep}} \cup A_t$ (if $A_t = \emptyset$, then $S_t = P_t^{\text{keep}}$).

(4) Evaluation and update: perform high-fidelity simulations/experiments for all $x \in S_t$; augment L with the new labeled data; update the surrogate parameters; use the updated stacking surrogate with NSGA-II to generate the next candidate set, P_{t+1}, and proceed to round $t + 1$.

By explicitly enforcing diversity and protecting rare informative samples, the proposed strategy reduces near-neighbor redundancy (increasing the information density of training data) while rescuing a few critical, globally isolated candidates that might otherwise be

discarded. The two components are complementary, improving both the spatial coverage and information value of the labeled set under a fixed high-fidelity budget. In turn, this enhances the global accuracy and robustness of the stacking surrogate and accelerates convergence toward the true Pareto front in concert with NSGA-II.

2.4. Experimental Design and Evaluation Metrics

Building on the framework in Sections 2.1–2.3, this section specifies the test reservoir, decision variables, and optimization objectives used to evaluate the effectiveness of the stacking surrogate + NSGA-II + DCAF co-optimization method.

Reservoir model and operating conditions: We adopt the public SPE5 light-oil reservoir model (Figure 15) as the benchmark case for evaluation [48]. The grid is 35 × 35 × 3 (3675 cells) over an area of approximately 1.1 km^2, with a thickness of 30 m and an average porosity of 30%. The model contains pronounced vertical heterogeneity: the three layers have markedly different horizontal permeabilities (about 500, 50, and 200 mD), while the vertical permeability is roughly 25–50 mD. Such layered contrast induces flow stratification and early gas channeling/override, which are typical challenges in gas-based EOR and therefore make WAG-type mobility control highly relevant. The reservoir fluid is a light volatile oil and the injected gas is pure CO_2 [49], allowing the case to retain the key compositional/phase-behavior effects that govern both displacement efficiency and CO_2 cycling. Key reservoir parameters are summarized in Table 2, and the reservoir-fluid and injected-gas compositions are given in Table 3. The well pattern consists of one injector and one producer, placed diagonally, providing a sufficiently long displacement path and a clear sweep-efficiency trade-off for comparing alternative WAG controls. The development scenario starts with primary depletion; after two years, CO_2-WAG (alternating water and CO_2 injection) is initiated for the forecast study. Upon completion of the WAG cycles, the injector is shut in and production continues under depletion. The simulation horizon is 10 years.

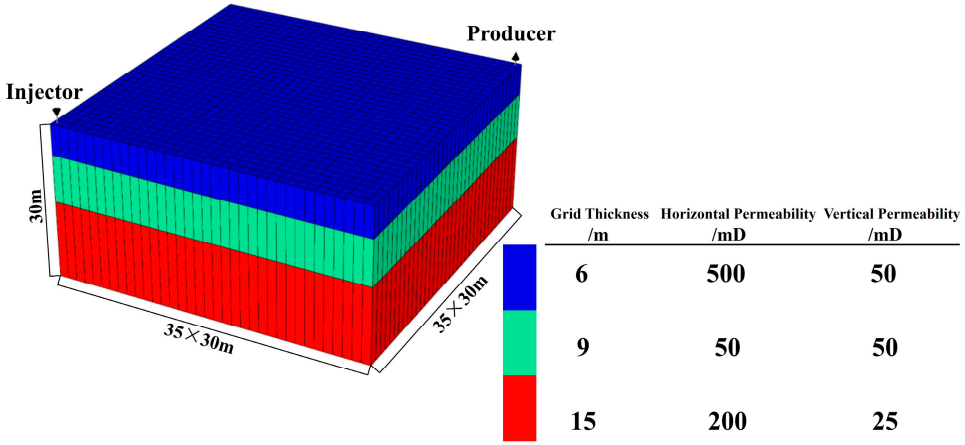


Figure 15. Schematic diagram of the SPE5 reservoir model.

Representativeness of the reservoir benchmark: Although SPE5 is a public benchmark model, it was originally designed for the Fifth Comparative Solution Project to evaluate miscible-flood simulators and has since been widely used to test algorithms under realistic miscible/near-miscible displacement physics. In the context of this paper, its representativeness stems from three aspects: (i) a strong vertical heterogeneity and permeability contrast that promote preferential flow and early gas breakthrough (a key motivation for WAG in field practice); (ii) light-oil compositional behavior with pure CO_2 injection, which captures the main CO_2-EOR mechanisms that affect both oil recovery and CO_2 storage ratio; and (iii) a simple but canonical injector–producer configuration that isolates the impact

of operational controls (rates, slug length, and duration) on recovery–storage–economics trade-offs, making it an efficient and transparent testbed for validating surrogate-assisted multi-objective optimization.

Table 2. Main parameters of the SPE5 reservoir model.

Parameter	Numerical Value	Unit
Reservoir area	1.1	km^2
Reservoir depth	2540	m
Reservoir thickness	(6, 9, 15)	m
Reservoir temperature	72	°C
Original formation pressure	27.5	MPa
Average porosity	30	%
Horizontal permeability	(500, 50, 200)	mD
Vertical permeability	(50, 50, 25)	mD
Rock compressibility	7.25×10^{-3}	MPa^{-1}
The compressibility of water	4.8×10^{-3}	MPa^{-1}
Initial water saturation	20	%
Number of grids	$35 \times 35 \times 3$	/
Grid size	30×30	m

Table 3. Composition of reservoir fluids and injected gas.

Composition	Reservoir Fluid Mole Fraction (%)	Injected Gas Mole Fraction (%)
CO_2	0	100
C1	50	0
C3	3	0
C6	7	0
C10	20	0
C15	15	0
C20	5	0

Mechanism of CO_2-WAG flooding: CO_2-WAG improves performance by combining mobility control and CO_2-driven microscopic displacement. Water slugs reduce effective gas mobility through relative permeability/saturation effects, mitigating viscous fingering and gravity override, and diverting the displacement front into less-swept zones; CO_2 slugs enhance microscopic recovery through oil swelling, viscosity reduction, and component extraction, while helping to maintain reservoir pressure. Therefore, injection rates and slug length jointly govern sweep efficiency, breakthrough behavior, and CO_2 cycling, which explains why they are selected as the key decision variables in this study [5–7].

The carbon storage pathway is mainly considered in this paper. Consistent with typical CO_2-EOR assessments, we focus on operational (engineering) storage over the project time scale, i.e., net CO_2 retained in the reservoir at the end of the simulation. In this framework, retained CO_2 primarily occurs via (i) hydrodynamic/structural retention of free-phase CO_2, (ii) residual trapping strengthened by water slugs, and (iii) solubility trapping as CO_2 dissolves into oil and water; long-term mineral trapping is not considered within the 10-year horizon [3,7,19,50].

Decision variables. Four control parameters are optimized in the WAG design: water-injection rate, CO_2-injection rate, WAG half-cycle length (slug length), and total WAG duration. The bounds (lower/upper limits) and initial values are listed in Table 4, spanning conservative to aggressive configurations. Low rates/short cycles represent conservative designs, whereas high rates/long cycles represent aggressive designs. This setup allows for

the optimizer to search across combinations of injection intensity and storage contribution to identify favorable trade-offs among recovery, storage, and economics.

Table 4. Parameter ranges for water-alternating-gas (WAG) injection.

Parameter	Unit	Minimum Value	Median Value	Maximum Value
Water injection rate	STB/D	4000	12,000	20,000
Gas injection rate	MCF/D	12,000	16,000	20,000
WAG half-cycle	day	60	210	360
WAG duration	year	4	6	8

Optimization objectives. Three objectives are maximized: the oil recovery factor (ORF), CO_2 storage ratio (CSR), and net present value (NPV). The ORF is defined as the ratio of the end-of-period cumulative oil production to the original oil in place (OOIP). The CSR is the fraction of injected CO_2 retained in the subsurface at the end of the period. The NPV is the discounted sum of net revenues (total revenue minus total cost) over the project's life. The economic assumptions are as follows: oil price USD 70/bbl, water-injection cost USD 2 per metric ton, CO_2-injection cost USD 0.02 per standard cubic foot (scf), and an annual discount rate of 10%. Although fixed values are used here to provide a consistent baseline, the proposed closed-loop "surrogate + NSGA-II + active learning" framework is not tied to a specific price/cost setting and can adapt to oil price fluctuations and development cost changes by updating these economic parameters (or conducting multi-scenario evaluations) and re-optimizing to obtain the corresponding Pareto front. Increasing the ORF and CSR typically requires higher injection volumes, which can raise production and storage but also substantially increase costs, thereby compressing the NPV; consequently, the three objectives involve clear trade-offs. Multi-objective optimization considers these metrics within a unified framework to produce a Pareto set that characterizes the balance among oil production, CO_2 storage, and economic performance—covering solutions from high-recovery/low-storage to high-storage/low-recovery, as well as economically balanced options—thus, providing decision makers with a diversified set of optimized schemes.

3. Results and Discussion

3.1. Comparison of Prediction Accuracy of the Surrogate Models

To verify the effectiveness of the constructed stacking ensemble surrogate model (hereafter referred to as the "stacking model"), this study systematically compares its prediction accuracy with that of each individual base model for three performance indicators: oil recovery factor, CO_2 storage ratio, and net present value (NPV). Figures 16–18 present the coefficient of determination (R^2), RMSE, and MAE of each model for the three responses. The results show that the R^2 values of the stacking model for oil recovery factor, CO_2 storage ratio, and NPV reach 0.9996, 0.9949, and 0.9826, respectively, with an average R^2 of 0.9923, indicating that it can reproduce the numerical simulation results very well. Meanwhile, the stacking model exhibits overall superiority in both accuracy and stability: over the entire dataset, the mean relative errors of the three outputs are all lower than 1%, with values of approximately 0.42% for the oil recovery factor, 0.85% for the CO_2 storage ratio, and 0.82% for the NPV. This level of accuracy is overall better than or comparable to that of the best-performing single model and is generally superior to most base models.

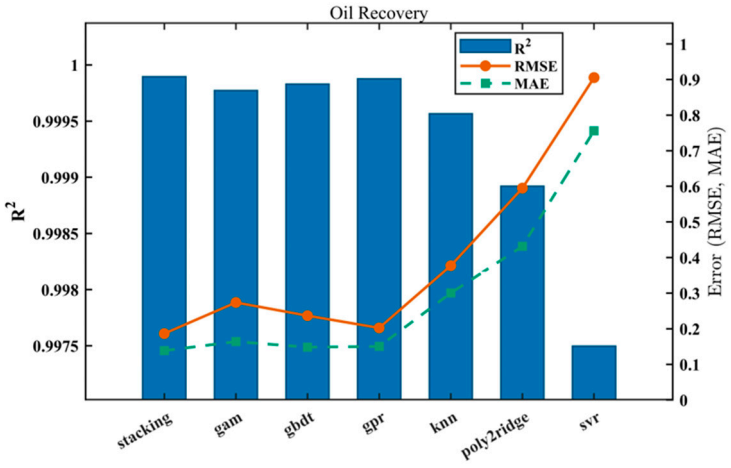


Figure 16. R^2, RMSE, and MAE of each model for predicting oil recovery.

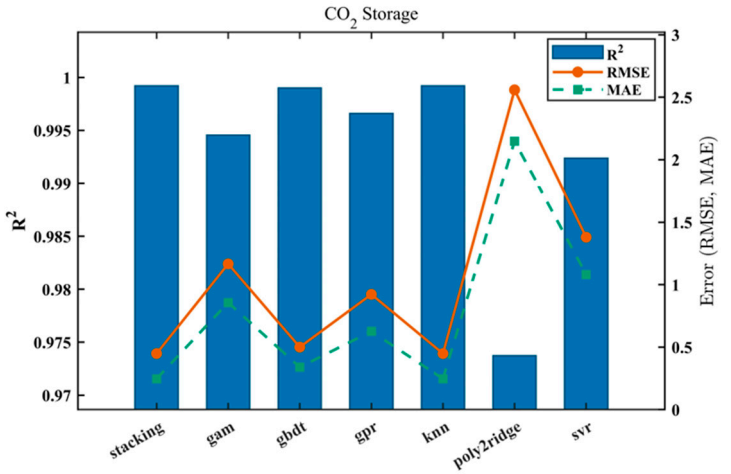


Figure 17. R^2, RMSE, and MAE of each model for predicting CO_2 storage.

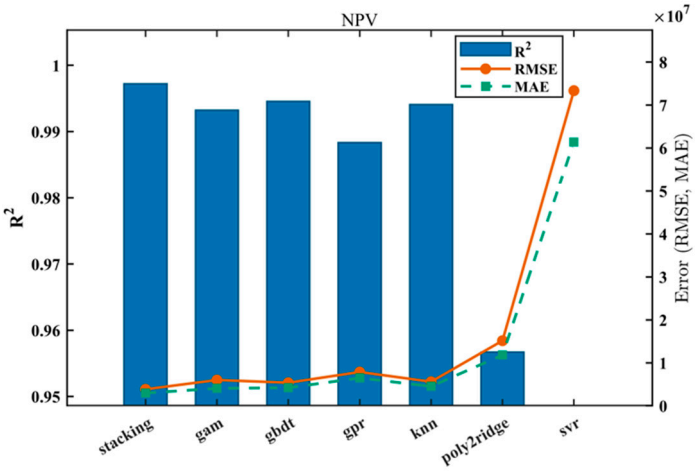


Figure 18. R^2, RMSE, and MAE of each model for predicting NPV.

Specifically, GBDT yields mean errors of about 0.41% and 0.75% for the oil recovery factor and NPV, respectively, representing a high level among similar methods, but its error for the CO_2 storage ratio is slightly higher, at about 1.41%. Although KNN achieves the lowest error for the CO_2 storage ratio (around 0.81%), its error for the oil recovery factor is about 0.59%, which is still higher than that of the stacking model. The mean errors of GAM and GPR are mostly around 1%, with acceptable accuracy but noticeable bias. Polynomial ridge regression (Poly2Ridge) has difficulty capturing nonlinear effects, leading

to mean errors of 1.24% and 2.28% for the oil recovery factor and NPV, respectively. SVR tends to suffer from overfitting or unstable hyperparameters, yielding a mean relative error of 11.17% for the NPV and a maximum error exceeding 50% for the oil recovery factor. These comparisons indicate that individual models struggle to perform well across all indicators simultaneously and thus cannot fully meet the prediction requirements for the three-objective problem.

By integrating heterogeneous base learners, the stacking model combines the strong generalization capability of GBDT, the local neighbor advantages of KNN, and the ability of GAM/GPR to fit smooth relationships, while effectively suppressing the extreme errors that are potentially caused by SVR and polynomial regression. In terms of error distribution, the stacking model is more robust: the maximum absolute relative error for any indicator is controlled within 16%—approximately 8.9% for the oil recovery factor, 15.8% for the CO_2 storage ratio, and 9.9% for the NPV—which is lower than the worst cases of several single models (for example, the maximum error of SVR reaches about 31.5% for the NPV and about 50.6% for the oil recovery factor). Therefore, the stacking ensemble surrogate model consistently outperforms the alternatives in accuracy and reliability and can serve as a high-precision and robust response surface evaluation tool for multi-objective optimization. The above evidence further validates the effectiveness of constructing surrogate models via ensemble learning: by integrating the complementary strengths of base models, the stacking model effectively alleviates the bias and high-variance issues in predicting complex nonlinear reservoir processes, thereby providing reliable support for subsequent optimization.

3.2. Optimization Convergence and Evolution of Model Error

Within the iterative framework of multi-objective optimization, the prediction error of the stacking ensemble surrogate model decreases progressively with iterations and tends to stabilize at later stages. Figure 19 compares the relationship between the predicted and simulated 10-year objective values at three stages—early (generation 1), intermediate (generation 14), and final (generation 27)—showing that incremental learning can continuously improve model accuracy and support algorithmic convergence. In the initial stage (generation 1), the mean relative errors are relatively high, due to the limited size of the training sample: approximately 2.35% for the oil recovery factor, 2.01% for the CO_2 storage ratio, and 4.88% for the NPV. As new solutions generated by the genetic algorithm in each generation are added to the training set, the coverage of the parameter space by the model improves and the errors shrink rapidly; by generation five, the mean relative errors of the oil recovery factor and the NPV are reduced to about 0.31% and 1.06%, respectively. By generation 10, the errors of all three indicators further decrease to about 0.19% for the oil recovery factor, 0.45% for the CO_2 storage ratio, and 0.52% for the NPV (the first two being below 0.5%). These results indicate that the continuous injection of new data effectively supplements the previously uncovered parameter regions, progressively eliminating systematic bias in the model and enabling prediction accuracy and optimization progress to improve in a coordinated manner. Figure 20 shows the evolution of the prediction error of the stacking ensemble surrogate model with the number of iterations.

It is worth noting that a brief increase in error occurs at generation 17. The mean relative error of the stacking model for the oil recovery factor is about 0.21% at generation 16 but rises to about 1.28% at generation 17; the mean relative errors of the CO_2 storage ratio and the NPV also increase from approximately 0.5% to above 1.5%. This phenomenon reflects the dynamic mechanism of interaction between the optimizer and the surrogate model: as the surrogate accuracy improves, the Pareto front obtained by the genetic algorithm approaches the true front. When generation 17 explores high-quality solutions that lie

outside the prior sample distribution (out of distribution, OOD), insufficient generalization of the surrogate in this local region leads to increased prediction bias, thereby raising the mean error of that generation. After incorporating the new samples from this generation into the training set, the model is quickly recalibrated; from generation 18 onward, the errors recover and continue to decline (for example, the mean relative error for the oil recovery factor is about 0.23% at generation 18). Although slight fluctuations remain for some indicators in generation 19 (e.g., about 1.4% for CO_2 storage ratio), the overall trend continues to improve and the prediction stability in the neighborhood of the Pareto front is gradually enhanced. By the time convergence is reached at generation 27, the mean relative errors of the three objectives stabilize at approximately 0.3–0.5%, indicating that the model has achieved a high prediction accuracy in the vicinity of the non-dominated solution set.

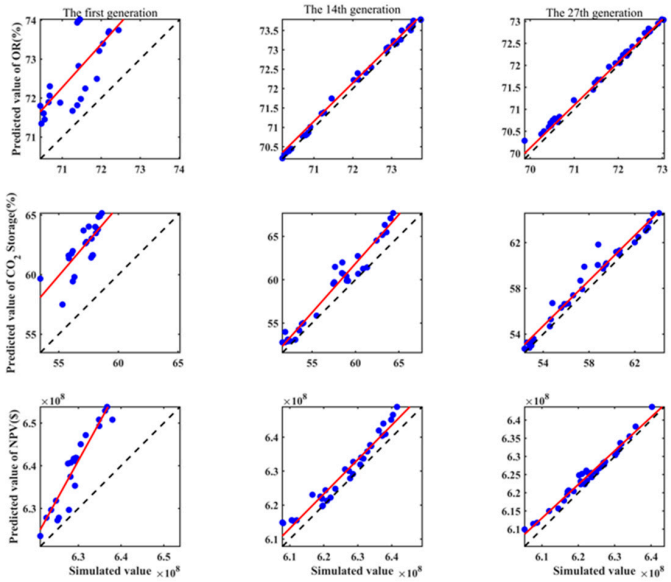


Figure 19. Comparison of predicted versus simulated values in the early, middle, and final stages, using the stacking surrogate model. The nine panels compare the predicted versus simulated 10-year objective values—oil recovery factor (ORF), CO_2 storage ratio (CSR), and net present value (NPV)—at three stages (early: generation 1; intermediate: generation 14; final: generation 27). Each dot represents one evaluated scheme (Stacking surrogate prediction plotted against the corresponding high-fidelity simulation result). The black dashed line denotes the 1:1 reference line (perfect agreement), and the red solid line denotes the least-squares fitted trend line in each panel.

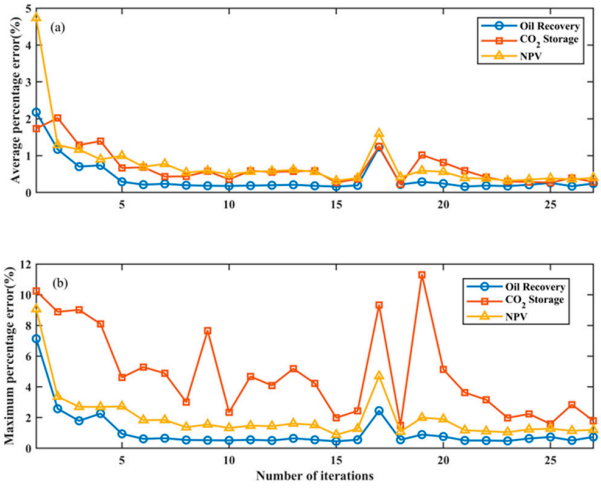


Figure 20. Evolution of the prediction error of the stacking surrogate model with the number of iterations: (**a**) mean error and (**b**) maximum error.

In summary, the coupled strategy of iterative optimization and incremental learning is effective in this study: the data from each generation are used to expand the training distribution and correct out-of-distribution bias. As a result, the optimization algorithm progressively approaches the global Pareto-optimal solution set while simultaneously improving the consistency of evaluations and the smoothness of convergence.

3.3. Pareto Solution Set and Multi-Objective Trade-Off Analysis

Through iterative optimization, a set of non-dominated CO_2-WAG designs is obtained. Based on the model described in Sections 2.1–2.4, each data point in Figure 21 is produced by inputting a four-dimensional WAG control vector x into the optimization–evaluation loop. NSGA-II is configured with a population size of p = 100 and a preset maximum number of generations, G = 100, with the stacking surrogate serving as the fitness evaluator for all individuals. When coupled with active learning, only K = 35 individuals per generation are selected and submitted to the high-fidelity reservoir simulator; the remaining P−K individuals only participate in genetic operations and do not trigger numerical simulations. The high-fidelity simulator outputs at the end of the 10-year forecast horizon (t = 10 years)—namely the ORF, CSR, and NPV (computed using the economic settings in Section 2.4)—are used as the coordinates in the three-objective space. In this study, the accuracy stopping criterion in Section 2.1 is satisfied at generation 27, so the algorithm terminates early: approximately P × 27 ≈ 2700 individuals are evaluated by the surrogate, and 27 × K = 945 development schemes are evaluated by high-fidelity simulation. Figure 21 therefore shows the final high-fidelity ("simulator-based") Pareto front obtained by merging the 945 simulated objective triplets from all 27 generations and applying non-dominated sorting; each point corresponds to one simulated CO_2-WAG design and its (ORF, CSR, NPV) values. Statistical analysis of this Pareto set reveals structural conflicts among the ORF, CSR, and NPV, such that the three objectives cannot be maximized simultaneously and trade-offs among them are required for decision making (Figure 21). To better visualize these trade-offs in a pairwise manner, Figure 22 presents the two-dimensional projections of the three-objective Pareto front (ORF–CSR, ORF–NPV, and CSR–NPV).

(1) Contradictory relationship between the ORF and the CSR.

The Pareto solution set shows an approximately linear negative correlation between the ORF and the CSR: an increase in the ORF is generally accompanied by a marked decrease in the CSR. A comparison of extreme solutions confirms this: when the ORF ≈ 73.0%, CSR ≈ 54.8%; conversely, when the CSR ≈ 100% (i.e., nearly all injected CO_2 is retained in the formation at the end of the period), the ORF is only ≈28.6%. This trade-off is mainly rooted in (i) the CO_2 mass-balance embedded in the definition of the CSR and (ii) the mobility and breakthrough behavior of CO_2 during WAG. In this study, the CSR is defined as the fraction of injected CO_2 that remains in the subsurface at the end of the project (Section 2.4). Neglecting leakage across the model boundary, it can be expressed as follows:

$$\mathrm{CSR} = 1 - \frac{M_{\mathrm{CO_2,prod}}}{M_{\mathrm{CO_2,inj}}} \tag{25}$$

where $M_{\mathrm{CO_2,inj}}$ denotes cumulative injected CO_2 and $M_{\mathrm{CO_2,prod}}$ (back-produced CO_2) denotes produced CO_2. This expression shows that, under the same accounting boundary, any operational change that increases CO_2 back-production directly lowers the CSR.

To push the ORF toward its upper extreme, the optimizer tends to select more "production-oriented" WAG settings (e.g., higher CO_2 injection rate, longer gas slugs/half-cycles, and/or longer active WAG duration). These settings enhance pressure support and CO_2–oil contact (oil swelling and viscosity reduction), improving displacement efficiency and thus, ORF [5–7]. However, CO_2 has a much lower viscosity than oil/water; increasing

the gas injection intensity or gas-slug length elevates the mobility ratio and gas relative permeability, which promotes gas channeling/viscous fingering and accelerates CO_2 breakthrough to the producer, even under WAG. After the breakthrough, maintaining high ORF typically involves substantial CO_2 cycling through repeated injection–production, meaning $M_{CO_2,prod}$ grows rapidly. Consequently, although high-ORF solutions may inject more total CO_2, a larger share is produced back, and the retained fraction (CSR) decreases. In contrast, "storage-oriented" designs that maximize the CSR must suppress CO_2 back-production by keeping the mobile gas phase away from the producer. This is generally achieved by lower gas injection intensity and shorter gas slugs (often with relatively stronger water support), which favors residual/solubility trapping at the end of cycles but reduces the extent and duration of effective CO_2–oil contact and pressure maintenance. As a result, the displacement benefit of CO_2 is weakened and the ORF decreases. Overall, because the ORF is improved by effective CO_2 cycling and production while the CSR penalizes CO_2 back-production by definition, the ORF and CSR are structurally conflicting on the Pareto set and cannot be maximized simultaneously.

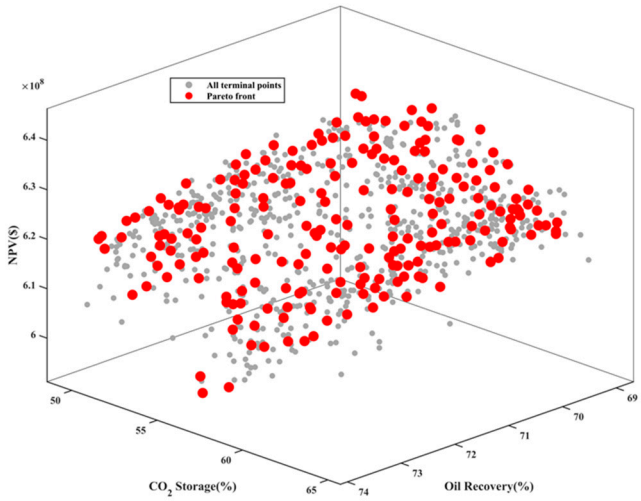


Figure 21. Distribution of the Pareto front in the three-objective space.

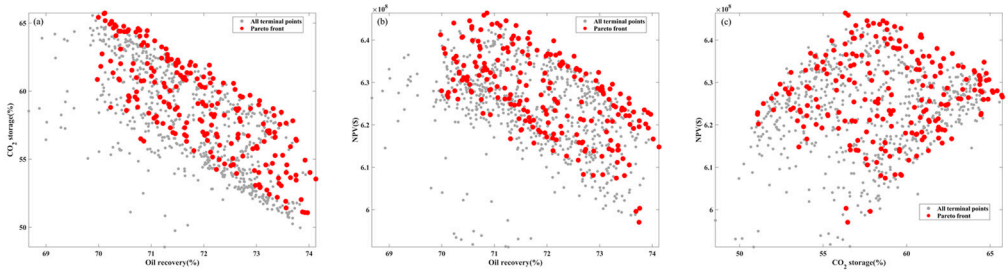


Figure 22. Two-dimensional projections of the three-objective Pareto front: (**a**) oil recovery vs. CO_2 storage, (**b**) oil recovery vs. NPV, (**c**) CO_2 storage vs. NPV.

(2) Trade-offs between the NPV and the techno-environmental objectives.

The NPV exhibits a moderately positive overall correlation with the ORF (correlation coefficient $\rho \approx 0.53$), but the relationship is not strictly monotonic. In the low-to-moderate ORF range, the increase in revenue from the additional oil production dominates, and the NPV increases with the ORF. Once the ORF exceeds a certain threshold, however, sustaining further production requires higher CO_2 injection and operational intensity, while the marginal gain in oil recovery diminishes and costs rise; as a result, the growth of the NPV slows and may even decline. Accordingly, in this study, the optimal NPV solution does not correspond to the maximum ORF: the maximum NPV is NPV $\approx 7.73 \times 10^8$ (about

773 million in the given monetary unit), with ORF ≈ 63.9% and CSR ≈ 73.0%; in contrast, the maximum ORF solution has NPV ≈ 6.22×10^8 and CSR ≈ 54.8%. When the ORF increases from ≈63.9% to ≈73.0%, the NPV decreases by about 1.51×10^8 (≈−19.5%), while the CSR drops by about 18.2 percentage points (from ≈73.0% to ≈54.8%). This indicates that pursuing extreme recovery beyond a moderate ORF weakens the economic returns.

In contrast, the NPV shows a weak negative correlation with the CSR ($\rho \approx -0.39$). Taking the NPV-optimal solution as a reference, if the CSR is increased from ≈73% to 100% (switching to the maximum CSR solution), the NPV decreases from ≈7.73×10^8 to ≈6.05×10^8—a reduction of about 1.68×10^8 (≈−21.7%)—while the ORF simultaneously drops by about 35.3 percentage points to ≈28.6%. These results quantitatively characterize the tension between the environmental objective and economic returns, pushing the CSR to its extreme substantially compresses development profits. In this sense, the NPV serves as a "regulator", linking the technological benefit (ORF) and environmental benefit (CSR): on the one hand, it benefits from the additional production associated with the moderately increased ORF; on the other hand, under excessive gas injection or stringent storage constraints, it is squeezed by both higher costs and reduced production. The Pareto solution set thus provides a set of comparable alternatives, enabling decision makers to implement transparent and traceable trade-offs among the three objectives, according to their preferences.

3.4. Validation of the Effectiveness of the DCAF Sampling Strategy

The proposed DCAF active learning strategy demonstrates a clear improvement in surrogate accuracy and Pareto front approximation under limited evaluation budgets. Table 5 summarizes the performance of the stacking surrogate model and final solution set for different sampling strategies: the baseline (no diversity control) versus DCAF with distance-loss thresholds β = 5%, 10%, and 20%. Lower β values correspond to preserving more "distance information" (i.e., stricter diversity filtering), resulting in more candidates being evaluated, whereas higher β aggressively filters out more redundant points. As expected, the baseline (no filtering) required the most high-fidelity simulations (945 evaluations over 27 iterations) to converge, while DCAF markedly reduced this number. For instance, with β = 20% (allowing 20% distance loss), DCAF evaluated only 461 samples—51.2% fewer evaluations than the baseline—and even the β = 5% case (most conservative filtering) used 698 samples (a 26.1% reduction) to achieve convergence, saving the substantial simulation budget.

Table 5. Comparison of the performance of the stacking model at the final prediction stage, under different sampling strategies.

Sampling Strategy	Sampling Quantity	Average Error of ORF (%)	Average Error of CSR (%)	Average Error of NPV (%)
baseline	945	0.395	2.90	0.691
distance-loss thresholds β = 5%	698	0.317	2.69	0.675
distance-loss thresholds β = 10%	599	0.347	2.84	0.692
distance-loss thresholds β = 20%	461	0.323	2.98	0.689

Despite using far fewer training points, DCAF achieved equal or better surrogate accuracy compared to the baseline. The stacking model's predictive error for oil recovery (one of the key objectives) was actually lower with DCAF. Under the baseline sampling, the surrogate's average relative error in oil recovery was about 0.39%, with a worst-case error of ~3.6% on a newly sampled design. In contrast, all DCAF variants maintained average errors of around 0.32–0.35%, and the worst-case error dropped to ~3%. This indicates that

by avoiding clustered, low-information samples, DCAF provided more informative and diverse training data, enabling the surrogate to learn the global response surface more accurately. Notably, even the most aggressive filtering (β = 20%, least samples) did not degrade the model—its oil recovery error (~0.323%) was on par with β = 5% (~0.317%) and was actually slightly better than the no-filter baseline. Similar trends were observed for the other two objectives (CO_2 storage and NPV), where the stacking model's prediction errors under DCAF were comparable to the baseline, despite using significantly fewer data points. Overall, the surrogate trained with DCAF was both highly accurate and robust, validating that diversity-controlled sampling effectively improves learning efficiency.

The Pareto front quality obtained by the surrogate + NSGA-II optimization was also greatly improved by DCAF's sample efficiency. Using the full dataset (baseline) as a reference for the "true" Pareto front, we compare the final solution sets from each strategy. With all 945 evaluations, the baseline naturally identified the complete Pareto front. Impressively, the DCAF strategies achieved an almost identical Pareto-optimal set while using far fewer samples (Figure 23). The β = 5% DCAF run, for example, consumed only 74% of the evaluation budget, yet its final Pareto front almost coincided with the true front, with differences in objective values on the order of 0.1% or less. Even in the β = 10% and β = 20% cases, the Pareto solutions obtained still spanned the major trade-off extremes across all objectives. More importantly, the Pareto-optimal points missed by DCAF were largely those very close to ones it found; quantitatively, the inverse generational distance (IGD) from the DCAF fronts to the global front is extremely small (e.g., averaging < 0.002 in normalized objective space for β = 10% and <0.001 for β = 5%). In practical terms, the DCAF-optimized fronts are virtually indistinguishable from the true Pareto front, despite using significantly fewer simulations. This confirms that DCAF's diversity-driven sampling, by maintaining broad coverage of the decision space, allows the surrogate-assisted NSGA-II to converge to the global Pareto-optimal solution set more efficiently. The results overwhelmingly validate the effectiveness of the DCAF strategy: it accelerates convergence to a high-fidelity Pareto front and enhances surrogate model accuracy, all while respecting a limited evaluation budget.

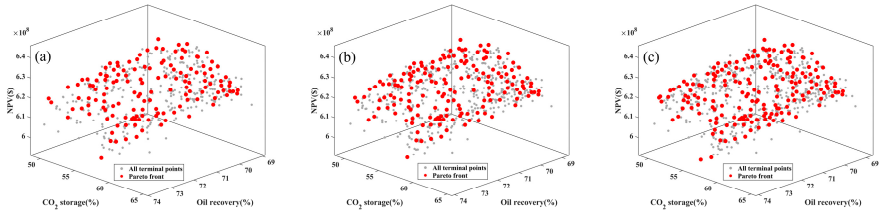


Figure 23. Comparison of the distribution of the Pareto front in the three-objective space under different sampling strategies: (**a**) distance-loss thresholds β = 20%, (**b**) distance-loss thresholds β = 10%, and (**c**) distance-loss thresholds β = 5%.

4. Conclusions

This study developed a closed-loop, multi-objective optimization framework for CO_2-WAG design that couples a diversity-aware stacking surrogate with NSGA-II and a decision-space diversity-controlled active learning strategy (DCAF). On the SPE5 benchmark, the framework jointly optimizes the oil recovery factor (ORF), CO_2 storage ratio (CSR), and net present value (NPV) using four operational controls, and it converges to a Pareto set that closely approximates the simulator truth while markedly reducing high-fidelity evaluations—thereby providing transparent, quantitative decision support for balancing production and carbon-mitigation benefits.

The following conclusions were drawn:

(1) Closed-loop optimization delivers accurate tri-objective fronts under tight simulation budgets. Embedding the stacking surrogate into NSGA-II, together with an iterative "surrogate prediction → high-fidelity correction → model update" loop, yields evenly distributed non-dominated solutions and a Pareto front that tracks the simulator truth; the loop stops when, for five successive generations, each objective's mean error falls below 1% and the maximum error below 3%, ensuring stable convergence near the front.

(2) The stacking surrogate attains high fidelity and robustness for all three objectives. Across the ORF, CSR, and NPV, the surrogate achieves $R^2 = 0.9996$, 0.9949, and 0.9826, with mean relative errors of ~0.42%, ~0.85%, and ~0.82%, respectively; error distributions are well controlled (max absolute relative error $\leq$ ~16%). Iterative learning further drives the mean errors in the Pareto neighborhood to ~0.3–0.5%, evidencing reliability for in-loop fitness evaluation.

(3) DCAF markedly improves the sample efficiency without degrading the front quality. Relative to a no-filter baseline requiring 945 simulator calls, DCAF reduces high-fidelity evaluations to 698 (-26.1%, $\beta = 5\%$) and 461 (-51.2%, $\beta = 20\%$), while maintaining slightly lower average prediction errors and producing Pareto fronts that are practically indistinguishable from the truth (normalized IGD on the order of 10^{-3}).

(4) Quantified trade-offs reveal structural tensions among production, storage, and economics. The ORF and CSR are negatively correlated on the Pareto set (e.g., ORF $\approx$ 73.0% with CSR $\approx$ 54.8% versus CSR $\approx$ 100% with ORF $\approx$ 28.6%). The NPV-optimal solution ($\approx 7.73 \times 10^8$) sits at a moderate ORF $\approx$ 63.9% and CSR $\approx$ 73.0%; pushing the ORF to $\approx$73.0% reduces the NPV by $\approx$19.5% and the CSR by $\approx$18.2 pp, while forcing the CSR to 100% lowers the NPV by $\approx$21.7% and ORF by $\approx$35.3 pp. These results position the NPV as the economic "regulator" that mediates between techno-productive and environmental gains.

(5) Decision value and transferability. By treating the ORF, CSR, and NPV as coequal objectives and enforcing decision-space diversity during sampling, the framework yields traceable, engineering-ready option sets and parameter ranges for CO_2-WAG design, and it can be readily adapted to shifts in reservoir conditions or economic objectives—supporting CCUS-aligned, low-carbon development pathways.

Limitations and outlook. This study's findings are based on the public SPE5 benchmark model with a one-injector–one-producer configuration. Extending the approach to field-scale, multi-well settings and incorporating additional uncertainties (e.g., price dynamics) are natural next steps. Even so, the proposed "stacking + NSGA-II + DCAF" loop already constitutes a practical, data-efficient optimizer for CO_2-EOR/WAG that unifies production, storage, and value—providing a replicable template for digital, low-carbon reservoir management.

Author Contributions: Conceptualization, Y.Z. (Yutong Zhu), H.L. and X.W.; methodology, Y.Z. (Yutong Zhu); software, Y.Z. (Yutong Zhu), C.L. and C.G.; validation, Y.Z. (Yutong Zhu), C.L. and C.G.; formal analysis, Y.Z. (Yutong Zhu) and Y.Z. (Yan Zheng); investigation, Y.Z. (Yutong Zhu), C.L. and C.G.; data curation, C.L. and C.G.; writing—original draft preparation, Y.Z. (Yutong Zhu); writing—review and editing, H.L., Y.Z. (Yan Zheng), X.W., C.L. and C.G.; visualization, Y.Z. (Yutong Zhu) and C.L.; supervision, H.L. and C.L.; project administration, H.L.; funding acquisition, H.L. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the China National Science Foundation (grant No. U2244215, U2344226, 42372286, 42002255, 42302297), the China Geological Survey Project (DD20221819), and the Fundamental Research Funds of the Chinese Academy of Geological Sciences (JKY202413, JKYQN202306).

Data Availability Statement: The data are available from the corresponding author upon reasonable request.

Conflicts of Interest: Author Hao Li was employed by the company Shaanxi Yanchang Petroleum (Group) Co., Ltd. Author Yan Zheng was employed by the company Shaanxi Yanchang Petroleum (Group) Co., Ltd. Gas Field Company. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Zhao, Z.-Y.; Yao, S.; Yang, S.-P.; Wang, X.-L. Under Goals of Carbon Peaking and Carbon Neutrality: Status, Problems, and Suggestions of CCUS in China. *Environ. Sci.* **2023**, *44*, 1128–1138.
2. Yang, Y. Technology progress and development direction of carbon capture, oil-flooding and storage in China. *Acta Perolei Sin.* **2024**, *45*, 325–338.
3. Wang, P.-T.; Wu, X.; Ge, G.; Wang, X.; Xu, M.; Wang, F.; Zhang, Y.; Wang, H.; Zheng, Y.J.S.; Transition, T.F.E. Evaluation of CO_2 enhanced oil recovery and CO_2 storage potential in oil reservoirs of petroliferous sedimentary basin, China. *Sci. Technol. Energy Transit. (STET)* **2023**, *78*, 3. [CrossRef]
4. Middleton, R.S.; Levine, J.S.; Bielicki, J.M.; Viswanathan, H.S.; Carey, J.W.; Stauffer, P.H. Jumpstarting commercial-scale CO_2 capture and storage with ethylene production and enhanced oil recovery in the US Gulf. *Greenh. Gases Sci. Technol.* **2015**, *5*, 241–253. [CrossRef]
5. Kulkarni, M.M.; Rao, D.N. Experimental investigation of miscible and immiscible Water-Alternating-Gas (WAG) process performance. *J. Pet. Sci. Eng.* **2005**, *48*, 1–20. [CrossRef]
6. Bocoum, A.O.; Rasaei, M.R. Multi-objective optimization of WAG injection using machine learning and data-driven Proxy models. *Appl. Energy* **2023**, *349*, 121593. [CrossRef]
7. Ding, S.; Wen, F.; Wang, N.; Zhang, Y.; Lu, R.; Gao, Y.; Yu, H. Multi-objective optimization of CO_2 enhanced oil recovery and storage processes in low permeability reservoirs. *Int. J. Greenh. Gas Control* **2022**, *121*, 103802. [CrossRef]
8. Li, Y. Technical advancement and prospect for CO_2 flooding enhanced oil recovery in low permeability reservoirs. *Pet. Geol. Recovery Effic.* **2020**, *27*, 1–10.
9. Wu, G.; Zhao, Z.; Wu, B. CO_2 flooding development models and economic benefit evaluation of dfferent types of reservoirs in subei basin. *Pet. Reserv. Eval. Dev.* **2021**, *11*, 864–870.
10. Du, X.; Salasakar, S.; Thakur, G. A comprehensive summary of the application of machine learning techniques for CO_2-enhanced oil recovery projects. *Mach. Learn. Knowl. Extr.* **2024**, *6*, 917–943. [CrossRef]
11. Wang, L.; Yao, Y.; Luo, X.; Adenutsi, C.D.; Zhao, G.; Lai, F. A critical review on intelligent optimization algorithms and surrogate models for conventional and unconventional reservoir production optimization. *Fuel* **2023**, *350*, 128826. [CrossRef]
12. Zhao, Y.; Luo, R.; Li, L.; Zhang, R.; Zhang, D.; Zhang, T.; Xie, Z.; Luo, S.; Zhang, L. A review on optimization algorithms and surrogate models for reservoir automatic history matching. *Geoenergy Sci. Eng.* **2024**, *233*, 212554. [CrossRef]
13. Rostamian, A.; de Moraes, M.B.; Schiozer, D.J.; Coelho, G.P. A survey on multi-objective, model-based, oil and gas field development optimization: Current status and future directions. *Pet. Sci.* **2025**, *22*, 508–526. [CrossRef]
14. Nguyen, Q.M.; Onur, M.; Alpak, F.O. Multi-objective optimization of subsurface CO_2 capture, utilization, and storage using sequential quadratic programming with stochastic gradients. *Comput. Geosci.* **2024**, *28*, 195–210. [CrossRef]
15. Fonseca, R.M.; Reynolds, A.C.; Jansen, J.D. Generation of a Pareto front for a bi-objective water flooding optimization problem using approximate ensemble gradients. *J. Pet. Sci. Eng.* **2016**, *147*, 249–260. [CrossRef]
16. Fu, J.; Wen, X.-H. Model-based multiobjective optimization methods for efficient management of subsurface flow. *SPE J.* **2017**, *22*, 1984–1998. [CrossRef]
17. Safarzadeh, M.A.; Motahhari, S.M. Co-optimization of carbon dioxide storage and enhanced oil recovery in oil reservoirs using a multi-objective genetic algorithm (NSGA-II). *Pet. Sci.* **2014**, *11*, 460–468. [CrossRef]
18. Isebor, O.J.; Durlofsky, L.J. Biobjective optimization for general oil field development. *J. Pet. Sci. Eng.* **2014**, *119*, 123–138. [CrossRef]
19. Kashkooli, S.B.; Gandomkar, A.; Riazi, M.; Tavallali, M.S. Coupled optimization of carbon dioxide sequestration and CO_2 enhanced oil recovery. *J. Pet. Sci. Eng.* **2022**, *208*, 109257. [CrossRef]
20. Liu, J.; Meng, F.; Zhao, H.; Xu, Y.; Wang, K.; Shi, C.; Chen, Z. Optimization of CO_2 EOR and geological sequestration in high-water cut oil reservoirs. *J. Pet. Explor. Prod. Technol.* **2024**, *14*, 1491–1504. [CrossRef]
21. Rodrigues, H.W.; Mackay, E.J.; Arnold, D.P. Multi-objective optimization of CO_2 recycling operations for CCUS in pre-salt carbonate reservoirs. *Int. J. Greenh. Gas Control* **2022**, *119*, 103719. [CrossRef]

22. Wang, L.; Zhang, L.; Deng, R.; Qu, J.; Wang, H.; Zhang, L.; Zhao, X.; Xu, B.; Lv, X.; Adenutsi, C.D. Active learning based surrogate ensemble assisted multi-objective optimization framework for reservoir water-flooding optimization. *J. Pet. Explor. Prod. Technol.* **2025**, *15*, 40. [CrossRef]
23. You, J.; Ampomah, W.; Sun, Q. Development and application of a machine learning based multi-objective optimization workflow for CO_2-EOR projects. *Fuel* **2020**, *264*, 116758. [CrossRef]
24. Amar, M.N.; Zeraibi, N.; Jahanbani Ghahfarokhi, A.J. Applying hybrid support vector regression and genetic algorithm to water alternating CO_2 gas EOR. *Greenh. Gases: Sci. Technol.* **2020**, *10*, 613–630. [CrossRef]
25. Shahkarami, A.; Mohaghegh, S. Applications of smart proxies for subsurface modeling. *Pet. Explor. Dev.* **2020**, *47*, 372–382. [CrossRef]
26. Wang, L.; Yao, Y.; Zhang, L.; Adenutsi, C.D.; Zhao, G.; Lai, F. An intelligent multi-fidelity surrogate-assisted multi-objective reservoir production optimization method based on transfer stacking. *Comput. Geosci.* **2022**, *26*, 1279–1295. [CrossRef]
27. Jiang, S.; Durlofsky, L.J. Use of multifidelity training data and transfer learning for efficient construction of subsurface flow surrogate models. *J. Comput. Phys.* **2023**, *474*, 111800. [CrossRef]
28. Kanaani, M.; Sedaghat Kameholiya, A.; Amarzadeh, A.; Sedaee, B. Stacking Learning for Smart Proxy Modeling in CO_2–WAG Optimization: A Techno-Economic Approach to Sustainable Enhanced Oil Recovery. *ACS Omega* **2025**, *10*, 9563–9582. [CrossRef]
29. Wang, L.; Deng, R.; Zhang, L.; Qu, J.; Wang, H.; Zhang, L.; Zhao, X.; Xu, B.; Lv, X.; Adenutsi, C.D. A Novel Surrogate-Assisted Multi-Objective Well Control Parameter Optimization Method Based on Selective Ensembles. *Processes* **2024**, *12*, 2140. [CrossRef]
30. Salehian, M.; Sefat, M.H.; Muradov, K. Multi-solution well placement optimization using ensemble learning of surrogate models. *J. Pet. Sci. Eng.* **2022**, *210*, 110076. [CrossRef]
31. Zhuang, X.-Y.; Wang, W.-D.; Su, Y.-L.; Dai, Z.-X.; Yan, B.-C. Deep learning-assisted optimization for enhanced oil recovery and CO_2 sequestration considering gas channeling constraints. *Pet. Sci.* **2025**, *22*, 3397–3417. [CrossRef]
32. Tang, H.; Durlofsky, L.J. Graph network surrogate model for optimizing the placement of horizontal injection wells for CO_2 storage. *Int. J. Greenh. Gas Control* **2025**, *145*, 104404. [CrossRef]
33. Kazemi, A.; Esmaeili, M. Reservoir Surrogate Modeling Using U-Net with Vision Transformer and Time Embedding. *Processes* **2025**, *13*, 958. [CrossRef]
34. Borisut, P.; Nuchitprasittichai, A. Adaptive Latin hypercube sampling for a surrogate-based optimization with artificial neural network. *Processes* **2023**, *11*, 3232. [CrossRef]
35. Yu, J.; Jafarpour, B. Active learning for well control optimization with surrogate models. *SPE J.* **2022**, *27*, 2668–2688. [CrossRef]
36. Lye, K.O.; Mishra, S.; Ray, D.; Chandrashekar, P. Iterative surrogate model optimization (ISMO): An active learning algorithm for PDE constrained optimization with deep neural networks. *Comput. Methods Appl. Mech. Eng.* **2021**, *374*, 113575. [CrossRef]
37. Li, Z.-L.; Peng, S.-S.; Wang, T. A surrogate-based optimization design method based on hybrid infill sampling criterion. *Eng. Mech.* **2022**, *39*, 27–33.
38. Wang, X.; van't Veld, K.; Marcy, P.; Huzurbazar, S.; Alvarado, V. Economic co-optimization of oil recovery and CO_2 sequestration. *Appl. Energy* **2018**, *222*, 132–147. [CrossRef]
39. Gao, M.; Liu, Z.; Qian, S.; Liu, W.; Li, W.; Yin, H.; Cao, J. Machine-learning-based approach to optimize CO_2-WAG flooding in low permeability oil reservoirs. *Energies* **2023**, *16*, 6149. [CrossRef]
40. Yang, R.-F.; Zhang, W.; Liu, S.-C.; Yuan, B.; Wang, W.-D. Multi-objective optimization workflow for CO_2 water-alternating-gas injection assisted by single-objective pre-search. *Pet. Sci.* **2025**, *22*, 2967–2976. [CrossRef]
41. Wolpert, D.H. Stacked generalization. *Neural Netw.* **1992**, *5*, 241–259. [CrossRef]
42. Barton, M.; Lennox, B. Model stacking to improve prediction and variable importance robustness for soft sensor development. *Digit. Chem. Eng.* **2022**, *3*, 100034. [CrossRef]
43. Karpatne, A.; Atluri, G.; Faghmous, J.H.; Steinbach, M.; Banerjee, A.; Ganguly, A.; Shekhar, S.; Samatova, N.; Kumar, V. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Trans. Knowl. Data Eng.* **2017**, *29*, 2318–2331. [CrossRef]
44. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182–197. [CrossRef]
45. Zhou, J.; Wang, H.; Xiao, C.; Zhang, S. Hierarchical surrogate-assisted evolutionary algorithm for integrated multi-objective optimization of well placement and hydraulic fracture parameters in unconventional shale gas reservoir. *Energies* **2022**, *16*, 303. [CrossRef]
46. Li, H.; Xu, Z.; Wei, W. Bi-objective scheduling optimization for discrete time/cost trade-off in projects. *Sustainability* **2018**, *10*, 2802. [CrossRef]
47. Wang, Q.; Wang, L.; Huang, W.; Wang, Z.; Liu, S.; Savić, D.A. Parameterization of NSGA-II for the optimal design of water distribution systems. *Water* **2019**, *11*, 971. [CrossRef]
48. Killough, J.; Kossack, C. Fifth comparative solution project: Evaluation of miscible flood simulators. In Proceedings of the SPE Reservoir Simulation Conference, San Antonio, TX, USA, 1–4 February 1987; p. SPE–16000–MS.

49. Sandve, T.H.; Sævareid, O.; Aavatsmark, I. Dynamic PVT model for CO_2-EOR black-oil simulations. *Comput. Geosci.* **2022**, *26*, 1029–1043. [CrossRef]
50. Ampomah, W.; Balch, R.; Cather, M.; Will, R.; Gunda, D.; Dai, Z.; Soltanian, M. Optimum design of CO_2 storage and oil recovery under geological uncertainty. *Appl. Energy* **2017**, *195*, 80–92. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.