

Surrogate-Guided Generative Modeling with Physics Verification for Expensive Multiphysics Inverse Problems

Hyungjun Kim^{a,b}, Alessandro Persico^c, Logan Burnett^c, Benoit Forget^c, Majdi I. Radaideh^{*a,d}

^a*Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, MI 48109, United States*

^b*Korea Atomic Energy Research Institute, Daejeon, 34057, South Korea*

^c*Department of Nuclear Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, United States*

^d*Department of Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109, United States*

Abstract

High-fidelity multiphysics simulations provide physically rich information for nuclear reactor analysis, but their computational cost often limits the duration and diversity of conditions that can be practically simulated. This work introduces a physics-informed guided generative modeling framework that learns from limited high-fidelity computational fluid dynamics (CFD) data and expands them into a searchable distribution of candidate states for coupled-physics inverse problems. Tabular generative models learn assembly-level spatial patterns from high-fidelity CFD simulations, while a differentiable surrogate guides candidate generation toward improved responses from a downstream neutronics model. Importantly, the surrogate is used only for sampling-time guidance; every generated candidate is independently evaluated by SIMULATE3 (a neutronics solver), ensuring that optimization decisions remain grounded in verified physics calculations. Two strategies are developed to balance prior fidelity, exploration, and computational cost: a hybrid explorer-generator framework and a more economical rank-weighted generative framework. The methodology is demonstrated for assembly-level DLT field inference for fuel assembly bowing using the BEAVRS benchmark. Starting from a baseline detector root-mean-square (RMS) error mismatch of approximately 5.1%, brute-force evaluation of the CFD-informed dataset alone achieved 4.2%, whereas generative optimization reduced the best verified RMS to 1.86%. The rank-weighted TabDDPM framework reached 2.08% using substantially fewer SIMULATE3 evaluations, while TabDiff reached 3.12% with greater spatial smoothness. The results further show that improved detector agreement requires controlled exploration beyond direct reproduction of the CFD prior. Overall, the proposed framework establishes a general strategy for combining limited high-fidelity data, generative exploration, surrogate guidance, and continuous physics verification to enable computationally tractable inverse analysis of expensive multiphysics systems.

Keywords: Computational Fluid Dynamics, Reactor Physics, Machine Learning, Generative Models, Optimization

1. Introduction

High-fidelity multiphysics simulation is essential for resolving the strongly coupled phenomena governing nuclear reactor behavior, but coupling detailed physics models across multiple spatial and temporal scales remains computationally demanding [1, 2]. Computational fluid dynamics (CFD), for example, provides detailed descriptions of coolant transport, mixing, and flow heterogeneity that are difficult to obtain from lower-order thermal-hydraulic models, but at substantially greater computational cost [3, 4]. Consequently, generating sufficiently long and diverse high-fidelity CFD datasets for direct exploration of large inverse, optimization, or uncertainty spaces can become computationally prohibitive.

Generative modeling provides a potential means of leveraging such limited high-fidelity datasets more effectively. Rather than treating available CFD calculations only as a fixed collection of states, generative models can learn complex high-dimensional distributions and sample additional candidate fields that preserve statistical structure

*Corresponding Authors: Hyungjun Kim (hyungjki@umich.edu), Majdi I. Radaideh (radaideh@umich.edu)

learned from the underlying data [5]. This concept is particularly attractive for inverse and black-box optimization problems, where generative priors can be progressively conditioned or guided toward desired responses while restricting the search to regions supported by the learned distribution [6, 7, 8]. Diffusion-based posterior sampling further enables gradient-based guidance of a learned generative prior for nonlinear inverse problems [9]. These developments motivate the use of generative models to expand the searchable core inlet conditions state space represented by limited CFD calculations while retaining high-fidelity physics simulation for downstream verification.

This work addresses these challenges through a CFD-informed guided generative optimization framework for reconstructing assembly-level DLT fields as proxies for the neutronic effects induced by fuel assembly bowing. A limited high-fidelity CFD dataset is used as a physical prior from which tabular generative models learn assembly-level spatial patterns. The generative distribution is then progressively expanded and refined according to feedback from the neutronics calculation. A differentiable surrogate provides gradients during generation to steer candidate fields toward improved solutions, while the underlying reactor-physics solver remains responsible for verification. Importantly, surrogate predictions are never used as substitutes for the physics calculation: every generated candidate contributing to the optimization is independently evaluated using neutronics solvers, and only simulator-verified results are used for ranking and updating the search.

The approach is demonstrated using the BEAVRS benchmark [10], where assembly-level inlet conditions are reconstructed to reproduce measured in-core reaction-rate distributions. The BEAVRS benchmark represents a four-loop pressurized water reactor with detector measurements available at 58 of the 193 fuel assemblies. At hot zero-power (HZP) conditions, these measurements exhibit a quadrant power tilt of approximately 5%, which is not reproduced by standard neutronic simulations. Previous stochastic reconstruction using hill climbing and simulated annealing demonstrated that perturbing the assembly-level inlet conditions through the DLT card can reduce the detector root mean square error mismatch from approximately 5.1% to the 1–2% range [11]. The present work advances this inverse problem by replacing direct stochastic perturbation with generative search informed by high-fidelity CFD data and verified through neutronics calculations. The main contributions of this work are:

- A CFD-informed generative approach that learns from a limited high-fidelity dataset and generates new structured fields for coupled multiphysics evaluation.
- A simulator-in-the-loop optimization approach that combines generative exploration with surrogate-gradient guidance while retaining reactor physics solvers as the sole source of verified objective values.
- Two novel optimization frameworks are proposed and tested—a hybrid explorer-generator framework and a rank-weighted generative framework—for balancing exploration beyond the CFD prior, preservation of physically informed spatial structure, and computational cost.
- Demonstration of the framework on a reactor neutronics inverse problem with real experimental and simulation data, including quantitative assessment of detector reconstruction accuracy, spatial smoothness, similarity to the CFD prior, and computational efficiency.

The remainder of this paper is organized as follows. Section 2 reviews related work on generative modeling, machine learning, and black-box optimization in the areas of fluid mechanics and nuclear engineering, followed by the relevant CFD and guided-diffusion background in Section 3. Section 4 presents the proposed CFD-informed generative optimization frameworks and their coupling with neutronics solvers. Section 5 presents the numerical results, compares the optimization strategies, and discusses their computational and physical implications. Finally, the main conclusions and directions for future work are summarized in Section 6.

2. Related Works

Artificial intelligence and machine learning have been increasingly applied across nuclear engineering for reactor-state prediction, transient diagnosis, high-dimensional physics modeling, condition monitoring, control, and digital twins. Early neural-network applications demonstrated the feasibility of identifying nuclear power plant transients directly from plant variables [12], while subsequent work developed dynamic neural-network ensembles capable of jointly detecting, classifying, and predicting reactor transients. More recent studies have expanded these capabilities using deep learning; for example, transfer learning has been investigated to improve nuclear power plant fault diagnosis under changing operating conditions [13], and comparative evaluations of deep neural architectures have demonstrated the potential of convolutional–recurrent models for equipment fault diagnosis under limited and noisy data [14]. These developments complement nuclear-specific studies that combine simulation data with deep learning and uncertainty quantification [15], analyze reactor simulation outputs and uncertainty using data-driven methods [16], forecast loss-of-coolant-accident transients [17], and approximate high-dimensional full-core BWR neutronics [18]. Beyond prediction and diagnosis, AI is increasingly being integrated with physics-based models. Gong et al. combined reduced-order modeling and machine learning to construct a physics-informed digital twin capable of rapidly reconstructing high-dimensional reactor neutron-flux and power fields [19], while Song et al. developed an online machine-learning calibration strategy to improve the accuracy of nuclear power plant digital twins while maintaining real-time capability [20]. AI applications have also progressed toward autonomous decision making, including deep reinforcement learning for microreactor load-following control [21] and multistep criticality search and power shaping [22]. Variational digital twins provide another means of combining physical models and observations for inference [23], while symbolic learning has been explored to improve interpretability of AI models for advanced energy systems [24]. The breadth of this research is reflected in recent surveys by Jendoubi and Asad, who reviewed AI applications spanning nuclear plant monitoring, diagnostics, prediction, and operation [25], and by Ejigu et al., who examined the integration of AI and big-data methods for nuclear power plant control [26]. Collectively, this literature demonstrates substantial progress in using AI to predict, reconstruct, diagnose, and control nuclear systems. *However, most existing approaches learn a direct input–output mapping, reduced-order representation, diagnostic model, or control policy from the available data.*

Surrogate modeling provides an established approach for reducing the computational burden of engineering analyses in which repeated high-fidelity simulation is prohibitively expensive. Classical surrogate-based optimization demonstrated how response surfaces and other inexpensive approximations can substantially reduce the number of direct black-box evaluations required during engineering design [27, 28]. Within nuclear engineering, deep Gaussian processes have been used to construct probabilistic surrogates of advanced simulations while representing predictive uncertainty [29], while Gaussian-process surrogates have also been applied to reactor depletion calculations to rapidly predict discharged-fuel nuclide compositions together with gradients and prediction uncertainties [30]. Neural-network surrogates have been extended to high-dimensional reactor calculations; for example, Bei et al. reconstructed three-dimensional neutron-flux distributions and multiplication factors for a molten-salt reactor from high-fidelity transport calculations [31]. Multifidelity approaches provide another route for reducing the dependence on expensive simulations. Chakroborty et al. combined Gaussian-process corrections, active learning, and models of different fidelities for rare-event estimation, including a TRISO-fuel application, while substantially reducing high-fidelity evaluations [32]. Similarly, multifidelity simulations and statistical surrogates have been coupled for dynamic probabilistic risk assessment of nuclear power plants to reduce the cost of repeated accident-sequence calculations [33]. Recent nuclear applications have further combined multiple simulation fidelities for thermal-hydraulic prediction [34] and incorporated uncertainty-aware surrogates directly into Monte Carlo design optimization [35]. These studies demonstrate that learned approximations can substantially reduce dependence on computationally

intensive physics calculations. However, directly optimizing or selecting designs using a learned surrogate introduces an important vulnerability: optimization can move toward regions in which approximation errors are large, particularly outside the distribution represented by the training data. *This motivates a different role for the surrogate in which it supplies differentiable guidance for navigating an expensive black-box problem, while the underlying physics simulator remains the sole source of verified objective values used for ranking and final candidate selection.*

Optimization of expensive nuclear simulations has traditionally relied on evolutionary, stochastic, and surrogate-assisted approaches. Kumar and Tsvetkov, for example, combined genetic algorithms with regression analysis for reactor design optimization involving neutronic and thermal-hydraulic considerations [36], while large-scale neuroevolution has subsequently been demonstrated for BWR fuel-bundle optimization with coupled physics constraints [37]. OpenNeoMC further coupled OpenMC with the NEORL framework for general black-box optimization of particle-transport simulations [38], and surrogate-driven optimization has been used to reduce expensive Monte Carlo evaluations while accounting for predictive uncertainty [35]. More broadly, Bayesian optimization provides a principled framework for expensive black-box objectives by iteratively using a probabilistic surrogate and acquisition function to determine where costly evaluations should be performed [39]. Data-driven offline optimization attempts to reduce such evaluations further by learning exclusively from previously evaluated designs; Design-Bench formalized this setting across high-dimensional scientific and engineering problems [40], while Conservative Objective Models demonstrated the importance of controlling surrogate overestimation under distribution shift when optimization moves beyond the available data [41]. Generative optimization provides a complementary perspective by learning a probability distribution over candidate solutions rather than optimizing individual design variables directly. Conditioning by Adaptive Sampling progressively adapts a generative distribution toward desirable black-box outcomes while retaining information from an initial data distribution [6]. Diffusion models provide flexible high-dimensional generative priors [5], and diffusion posterior sampling has shown that such priors can be guided by measurement gradients for nonlinear inverse problems [9]. Denoising Diffusion Optimization Models subsequently extended diffusion generation to offline black-box optimization by learning conditional distributions over designs and their objective values [7]. These methods are particularly relevant to expensive engineering design, but offline model-based optimization remains fundamentally constrained by the coverage of the available dataset, as emphasized by Trabucco et al. [40]. *A key gap therefore remains in combining a generative prior learned from limited high-fidelity multiphysics data with surrogate-guided exploration and continuous black-box physics verification, allowing the search to move beyond the finite simulation dataset without treating either generated samples or surrogate predictions as substitutes for high-fidelity physics.*

3. Background

3.1. High-fidelity CFD data

In our prior study [42], a high-fidelity STAR-CCM+ model of the lower plenum was developed for a full-scale commercial PWR based on the BEAVRS benchmark. The computational domain, illustrated in Figure 1, incorporates a realistic pump swirl profile at the cold legs to capture flow asymmetries that potentially drive core power tilts. Through mesh sensitivity study, the axial evolution of assembly-level mass flow rates was quantified. The resulting mass flow distribution at the core inlet exhibits a persistent assembly-scale asymmetry of approximately 4.7%, arising in part from pump-induced swirl at the cold legs. This magnitude is comparable to the approximately 5% power tilt observed in the BEAVRS detector measurements and therefore provides quantitative motivation for incorporating the CFD spatial pattern into the inverse reconstruction. For the subsequent data-driven framework, the averaged mass flow rates at the core inlet plane, extracted from this high-fidelity simulation, serve as the baseline dataset. This CFD-informed flow rate pattern is injected into the core model as an equivalent inlet temperature

distribution: the assembly-wise mass flow variations are mapped to inlet temperature perturbations, so that the CFD-derived spatial structure, assumed to drive fuel bowing and therefore the power tilt, is preserved in the DLT field.




Figure 1: Visualization of 4-loop PWR geometry developed in StarCCM+ in our prior study [42]

3.2. Guided diffusion

Diffusion models are generative models that learn to transform samples from a simple noise distribution into samples from a target data distribution through an iterative denoising process [5]. In denoising diffusion probabilistic models (DDPMs), a forward process gradually adds Gaussian noise to a clean data sample $\mathbf{x}_0$, producing noisy states $\mathbf{x}_t$ over time steps $t = 1, \dots, T$. A neural network is then trained to approximate the reverse denoising process. A common DDPM objective is

$$\mathcal{L}_{\text{DDPM}} = \mathbb{E}_{\mathbf{x}_0, t, \epsilon} \left[\left\| \epsilon - \epsilon_\theta \left(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t \right) \right\|_2^2 \right] \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, I)$, $\bar{\alpha}_t$ is determined by the noise schedule, and ϵ_θ is the learned denoising model [5]. While diffusion models were originally developed for image generation, they have also been adapted to tabular data [43].

To adjust the reverse generation trajectory toward specific design criteria without altering the underlying data density prior, conditioning and guidance mechanisms are applied during inference. Classic classifier guidance changes the sampling process by injecting the spatial gradients of an auxiliary, separately trained noise-robust classifier network to bias the denoising path toward a target discrete label [44]. Classifier-free guidance (CFG) constructs a joint architecture that simultaneously adapts conditional and unconditional trajectories within a single network model, mixing their noise predictions at inference time through a tunable scaling factor [45].

For complex engineering applications where targets are governed by nonlinear systems, standard categorical guidance is insufficient. To bypass the non-differentiable or computationally prohibitive nature of high-fidelity nu-

merical solvers, recent frameworks deploy a differentiable surrogate network to execute Diffusion Posterior Sampling (DPS) [46]. By computing the gradient of an objective loss function directly through the surrogate model with respect to the intermediate state $\mathbf{x}_t$, the reverse sampling path can be explicitly steered to satisfy strict physical constraints or match downstream target signals, establishing a robust paradigm for black-box simulation-based optimization [46].

3.3. Black-box optimization

Many engineering inverse problems can be formulated as black-box optimization problems, where evaluating a candidate design requires an expensive simulator or experiment. Generative-model-based optimization addresses this setting by learning a distribution over promising and feasible designs rather than directly searching the design space using objective gradients typically leveraged by gradient-based [47], evolutionary [48], reinforcement learning [49], or hybrid neuroevolution [50, 51] algorithms.

An early representative approach is Conditioning by Adaptive Sampling [52], which was developed for robust design under black-box stochastic oracles. CbAS begins with a generative prior $p_0(\mathbf{x})$ representing realistic designs and seeks to approximate the prior distribution conditioned on a desired property event S , such as exceeding a prescribed performance threshold. The target distribution is expressed as

$$p(\mathbf{x} | S) \propto P(S | \mathbf{x})p_0(\mathbf{x}), \quad (2)$$

where $P(S | \mathbf{x})$ is estimated using a property oracle. Instead of directly maximizing the oracle prediction, CbAS iteratively updates a generative search distribution through importance-weighted maximum likelihood as in Algorithm 1. The importance weights favor designs that are likely to satisfy the target event while correcting for the distribution shift between the original prior and the current search distribution. This mechanism allows the search to move toward high-performing regions without departing excessively from the support of realistic training data.

Denosing Diffusion Optimization Models (DDOM) [7] extend the conditional-generation perspective to offline black-box optimization. In this setting, the objective is to optimize an expensive function using only a fixed dataset of previously evaluated designs without additional simulator or oracle evaluations during model training. DDOM trains a conditional diffusion model to approximate the inverse design distribution $p_\theta(\mathbf{x} | y)$ where y denotes the objective value associated with design $\mathbf{x}$. At inference time, the diffusion model is conditioned on a desired high objective value, and classifier-free guidance can be used to strengthen the dependence of generated samples on the target objective. By conditioning on values near or beyond the upper range of the offline dataset, DDOM attempts to generate improved candidate designs.

CbAS and DDOM therefore share the same general principle: optimization is reformulated as sampling from a design distribution conditioned on desirable performance. Their implementations, however, differ substantially. CbAS constructs the conditional distribution indirectly through adaptive importance weighting and repeated re-fitting of a generative model, whereas DDOM learns the inverse conditional distribution directly from paired design-objective data using a diffusion process. CbAS emphasizes maintaining proximity to a realistic prior under oracle-guided adaptation, while DDOM emphasizes offline conditional generation and guidance toward user-specified objective values. Both approaches are particularly attractive for high-dimensional inverse problems in which multiple feasible designs may yield similar objective values and the inverse mapping from performance to design is inherently one-to-many.

Algorithm 1 CbAS [52]

```
1: for  $t = 1, \dots, T$  do
2:   Sample  $M$  candidates  $\mathbf{x} \sim p_t(\mathbf{x})$ 
3:   Evaluate oracle  $f_\phi(\mathbf{x})$  (with uncertainty)
4:   Compute importance weights via prior ratio  $\times$  oracle probability  $P(f_\phi(\mathbf{x}) > \gamma_t)$ 
5:   Fit  $p_{t+1}$  by weighted MLE (retrain generative model, from scratch or fine-tune)
6:   Anneal threshold  $\gamma_t \rightarrow \gamma$ 
7: return best sample under  $p_T$ 
```

3.4. Tabular data generation model

Generative modeling of tabular data has advanced in recent years using diffusion- and flow-based methods that better capture the heterogeneous, mixed-type nature of tabular features. These newer approaches address challenges in tabular generation, including the joint modeling of continuous and categorical columns, sensitivity to feature scaling, and mode collapse under limited or imbalanced data. We review four representative tabular generative models spanning diffusion-based, flow-based, and non-parametric paradigms.

TabDDPM [43] adapts the denoising diffusion probabilistic model (DDPM) framework to mixed-type tabular data by applying Gaussian diffusion to continuous features and multinomial diffusion to categorical features within a unified reverse-denoising network. Its straightforward score-based formulation makes it a natural target for diffusion posterior sampling (DPS)-style guidance, since the learned noise-prediction network directly yields the estimate required at each denoising step.

TabDiff [53] extends the diffusion paradigm by introducing a unified continuous-time diffusion process for mixed-type tabular data, employing feature-specific noise schedules and a classifier-free-guidance (CFG)-based conditioning mechanism to better capture heterogeneous marginal distributions across continuous and categorical columns. While its native CFG offers strong unconditional sample quality, its guidance mechanism is coupled to the training objective, which trades away some of the per-round retraining flexibility that sampling-time guidance methods such as DPS provide in iterative, oracle-in-the-loop settings.

TabbyFlow [54] formulates tabular data generation as exponential-family variational flow matching, generalizing flow matching beyond the Gaussian probability-path assumption to distributions natively suited to each feature type. As a flow-based model, it produces straighter probability paths than diffusion-based counterparts, which can reduce the number of integration steps required at sampling time, at the potential cost of additional overhead when paired with optimal-transport-based minibatch coupling.

TabSDS [55] departs from parametric deep generative modeling entirely, instead proposing a lightweight, fully non-parametric and model-free procedure for synthesizing tabular data. Its non-parametric nature makes it well suited for broad, low-cost exploratory sampling in early optimization rounds, though its lack of an explicit conditional density model precludes direct gradient-based guidance and favors exact elite duplication over loss-reweighted fine-tuning when incorporating high-scoring candidates.

Table 1 summarizes the four models along the dimensions most relevant to their role within our proposed work.

4. Methodology

We develop two novel simulator-in-the-loop optimization frameworks. The first combines broad exploration with surrogate-guided refinement using verified elite candidates. The second replaces elite augmentation with rank-weighted generator adaptation. In both frameworks, the surrogate is used only to *guide* the generation process; all objective values used for model updating and final selection are evaluated with physics-based simulation tools such as SIMULATE3 [56] in this research.

Table 1: Comparison of surveyed tabular generative models

Model	Paradigm	Conditioning	Relevance to this work
TabDDPM [43]	DDPM (Gaussian + multinomial)	None (unconditional); amenable to sampling-time guidance	Direct $\hat{\mathbf{x}}_0$ estimation supports DPS guidance
TabDiff [53]	Continuous-time diffusion	Classifier-free guidance (CFG)	Strong unconditional quality; guidance coupled to training
TabbyFlow [54]	Exponential-family flow matching	None (unconditional)	Straighter paths; fewer sampling steps
TabSDS [55]	Non-parametric, model-free	Not applicable	Fast exploratory seeding; no gradient guidance

4.1. Reactor physics

SIMULATE3 is an advanced three-dimensional, two-group nodal reactor analysis code developed by Studsvik Scandpower for light-water reactor core analysis [56]. In this work, SIMULATE3 is used as the reactor physics solver to evaluate the neutronic response of candidate inlet-condition fields. The target response is the axially integrated detector reaction rate (RR) at the measured fuel-assembly locations. For each candidate distribution generated by the tabular generative model, a SIMULATE3 input deck is prepared by perturbing the assembly-wise inlet coolant temperature field through the *HYD.DLT* card. The resulting RR distribution is then compared with the measured BEAVRS detector reaction rates.

DLT card input in SIMULATE3 represents an additive inlet-temperature perturbation from the reference core inlet temperature. The FLO card input represents an assembly-wise normalized mass flow rate fraction with respect to the nominal value, where $F_i = 1.0$ denotes the unperturbed reference flow condition. Thus, DLT and FLO have different physical meanings: DLT is a temperature deviation, while FLO is a dimensionless flow multiplier.

The use of DLT instead of the FLO card is motivated by its direct impact on moderator density and, consequently, on the local moderation condition. A non-uniform inlet temperature field modifies the moderator mass in each fuel assembly channel and can therefore reproduce radial reaction-rate asymmetries. This treatment is consistent with the surrogate-variable approach in [11], where the assembly inlet temperature distribution is introduced through the SIMULATE3 DLT card to mimic the neutronic effect of fuel bowing. In that study, DLT is treated as a continuous surrogate variable, and the local detector response is shown to vary approximately linearly with DLT within the considered perturbation range. This indicates that the DLT field provides an effective control variable for modifying the reaction-rate distribution in a more physical way than what the FLO card offers in this context.

It is important to note that BEAVRS is considered at hot zero power, meaning the core is nominally isothermal at approximately 566.5 K. The non-zero DLT fields used here are not physical thermal-hydraulics perturbations: they are surrogate variables that act on local water density to reproduce the neutronic effect of geometric assembly bowing without explicitly modeling the deformed geometry. The physical chain is: bowing displaces fuel assemblies laterally, modifying inter-assembly water gaps; this redistributes local moderator volume; which shifts neutron thermalization and reaction rates assembly-by-assembly. Since SIMULATE3 does not offer a direct bowing model, DLT provides an equivalent handle on the same moderation effect.

Each DLT vector is inserted into the SIMULATE3 input deck and evaluated independently. The calculated reaction-rate vector is expressed as

$$\mathbf{r}_{\text{cal}} = \mathcal{S}(\mathbf{x}_{\text{DLT}}) \quad (3)$$

where $\mathbf{x}_{\text{DLT}}$ is the 193-dimensional vector of assembly-wise inlet-temperature perturbations, $\mathcal{S}(\cdot)$ denotes the SIM-

ULATE3 calculation, and $\mathbf{r}_{\text{cal}}$ is the resulting calculated detector reaction-rate vector. Because the BEAVRS benchmark provides measured detector signals for only a subset of assemblies, the objective function is evaluated at the measured detector locations:

$$\text{RMS} = \sqrt{\frac{1}{N_m} \sum_{k=1}^{N_m} \left(\frac{r_{\text{cal},k} - r_{\text{meas},k}}{r_{\text{meas},k}} \right)^2} \times 100 \quad (4)$$

where RMS is the percentage root mean squared mismatch, $N_m = 58$ is the number of measured detector locations, $r_{\text{cal},k}$ is the calculated reaction rate at detector location k , and $r_{\text{meas},k}$ is the corresponding measured BEAVRS reaction rate. A spatial smoothness metric is also computed for the generated DLT field using total variation:

$$\text{TV} = \sqrt{\frac{1}{2N} \sum_{i=1}^N \sum_{j \in \mathcal{N}(i)} (\Delta T_i - \Delta T_j)^2}, \quad N = 193 \quad (5)$$

where TV measures assembly-to-assembly spatial variation, $N = 193$ is the total number of fuel assemblies, ΔT_i and ΔT_j are the inlet-temperature perturbations of neighboring assemblies i and j , respectively, and $\mathcal{N}(i)$ denotes the set of assemblies neighboring assembly i .

Although both mass flow rate and temperature perturbations can be considered as assembly-wise inlet-condition variables, the present framework uses DLT as the primary SIMULATE3 control variable. The CFD-derived mass flow rate field is not directly optimized in SIMULATE3; instead, it is used to construct physically informed DLT perturbation fields that preserve the spatial redistribution pattern predicted by the high-fidelity CFD calculation.

4.2. Baseline CFD dataset

The baseline dataset was constructed from the transient CFD simulation reported in [42]. In the simulation, as shown in Figure 2, sensor planes were placed at the inlet of each fuel assembly to record the assembly-wise average mass flow rate. Although nine axial layers of sensor planes were generated in the CFD model, this study uses the mass flow rate distributions of the 193 fuel assemblies at the core inlet. The transient simulation was performed for 15 s with a time step of 0.0005 s. The first 10 s were excluded to allow flow development. The remaining interval from 10 to 15 s was used for data collection, yielding 10,000 snapshots of the 193-dimensional mass flow rate distribution.

Because these snapshots are sampled much more rapidly than the characteristic evolution of the assembly-level flow field, the nominal sample count substantially exceeds the amount of distinct temporal information in the trajectory. Temporal autocorrelation was evaluated for the assembly-wise histories after removing the instantaneous core-wide mean so that the diagnostic emphasizes persistence of the spatial redistribution pattern. The median integrated autocorrelation time was approximately $\tau_{\text{int}} = 0.137$ s, corresponding to roughly 274 stored CFD time steps. Thus, the 10,000 snapshots are highly redundant even though they densely resolve the 5 s transient. This motivates the use of a generative model to learn the spatial prior represented by the CFD trajectory and produce additional candidate states rather than treating every stored snapshot as an independent realization.

Running a CFD simulation at this scale is computationally expensive, requiring approximately 20 days of wall-clock time using 200 CPU cores in parallel. Rather than extrapolating the CFD solution forward in time, the generative models seek to sample plausible flow patterns consistent with the observed turbulent behavior, thereby extending the diversity of inlet conditions beyond what can be obtained through computationally prohibitive long-duration CFD simulations (e.g., minutes to hours).

To incorporate the CFD-derived flow distribution into the reactor physics calculation, each mass flow rate snapshot was transformed into a fuel-assembly inlet temperature perturbation field. The normalized flow perturbation was mapped to the inlet temperature deviation as

$$\Delta T_i = T_{\text{in}} (F_i - 1.0) f, \quad (6)$$

where F_i is the normalized mass flow rate of the i -th fuel assembly, T_{in} is the reference core inlet temperature, and f is a scaling factor controlling the magnitude of the temperature perturbation. The resulting 193-dimensional ΔT vectors represent assembly-wise inlet temperature perturbation fields derived from the high-fidelity CFD flow distributions. These transformed fields were used as the base dataset for training the tabular generative models.




Figure 2: Selected axial planes for monitoring coolant mass flow rate distribution (kg/s) across the fuel assemblies from the lower entrance of the core near Plane 1 to the upper exit of the core near Plane 9. [42]

4.3. Hybrid Exploration and Surrogate-Guided Optimization

The first inverse framework combines an exploration model with a surrogate-guided diffusion generator. TabSDS [55] is used for exploration and TabDDPM [43] for surrogate-guided refinement. At each iteration, the exploration model proposes candidates that are screened by a trend constraint and verified in SIMULATE3; these verified exploratory candidates form the training set for both the surrogate and the guided generator of that iteration. The guided generator then produces refined candidates, which are verified and admitted to the optimization pool. The procedure is summarized in Algorithm 2. Here, $\mathcal{P}_r$ denotes the optimization pool after round r , $\mathcal{A}_r$ is the cumulative archive of simulator-verified candidates, $\mathcal{E}_r$ is the verified exploratory set, $\mathcal{G}_r$ is the set of verified surrogate-guided candidates, $\mathcal{G}_r^*$ contains the top- n guided candidates selected by FOM, and R is the total number of optimization rounds.

The procedure begins from an initial verified pool $\mathcal{P}_0$. At optimization round r , TabSDS is trained on the current pool $\mathcal{P}_{r-1}$ and generates a new set of exploratory DLT fields. A centered Pearson-correlation constraint is applied to preserve the spatial trend represented by the CFD-informed base pool. Exploratory candidates are retained

only when their maximum centered Pearson correlation with the reference pool exceeds a prescribed threshold. All retained candidates are then evaluated using SIMULATE3, forming the verified exploratory set $\mathcal{E}_r$.

The candidate is measured using the normalized form of ?? as below;

$$\text{FOM}(\mathbf{x}) = \frac{\text{RMS}(\mathbf{x})}{\text{RMS}_{\text{init}}} + \lambda_{\text{TV}} \frac{\text{TV}(\mathbf{x})}{\text{TV}_{\text{limit}}} \quad (7)$$

Here, RMS_{init} is the RMS error of the unperturbed reference solution, and TV_{limit} is the prescribed upper scale for acceptable spatial variation. This normalization makes the two terms dimensionless and comparable in magnitude, avoiding direct competition between quantities with different units and numerical scales.

The verified exploratory candidates are subsequently used to train both the surrogate model and TabDDPM generator. The differentiable surrogate f_ϕ approximates the simulator-evaluated objective as

$$f_\phi(\mathbf{x}) \approx \text{FOM}(\mathbf{x}) \quad (8)$$

where $\mathbf{x}$ is a candidate DLT field, $\text{FOM}(\mathbf{x})$ is its verified SIMULATE3-based objective value, $f_\phi(\mathbf{x})$ is the corresponding surrogate prediction, and ϕ denotes the trainable parameters of the surrogate model.

During diffusion-based generation, the surrogate gradient is used to modify the sampling trajectory following [46],

$$\mathbf{g}_\phi(\mathbf{x}_t) = -\nabla_{\mathbf{x}_t} f_\phi(\mathbf{x}_t) \quad (9)$$

where $\mathbf{x}_t$ denotes the candidate state at diffusion step t , $f_\phi(\mathbf{x}_t)$ is the surrogate-predicted objective at that state, $\nabla_{\mathbf{x}_t}$ denotes the gradient with respect to $\mathbf{x}_t$, and $\mathbf{g}_\phi(\mathbf{x}_t)$ is the resulting guidance direction. The negative gradient biases the sampling trajectory toward candidates predicted to have lower FOM, with the magnitude of the correction controlled by the guidance coefficient.

Each round contains both guided and unguided proposals. Guided samples exploit regions predicted to have low objective values, whereas unguided samples preserve broader exploration. All proposals are subsequently evaluated using SIMULATE3. Therefore, the surrogate serves only as a proposal mechanism, while the optimization archive and final solution are determined exclusively from verified simulator evaluations.

Algorithm 2 Hybrid Exploration and Surrogate-Guided Optimization

Require: Initial verified pool $\mathcal{P}_0$

- 1: Initialize archive $\mathcal{A}_0 \leftarrow \mathcal{P}_0$
 - 2: **for** $r = 1, \dots, R$ **do**
 - 3: Train exploration model E_ψ on $\mathcal{P}_{r-1}$
 - 4: Generate exploratory candidates using E_ψ
 - 5: Apply the Pearson-correlation trend constraint
 - 6: Evaluate retained candidates using SIMULATE3
 - 7: Collect verified exploratory candidates $\mathcal{E}_r$; add them to $\mathcal{A}_r$
 - 8: Train surrogate f_ϕ on $\mathcal{E}_r$
 - 9: Train generator G_θ on $\mathcal{E}_r$
 - 10: Generate surrogate-guided candidates using G_θ
 - 11: Evaluate all guided candidates using SIMULATE3
 - 12: Add valid guided evaluations $\mathcal{G}_r$ to $\mathcal{A}_r$
 - 13: Select the top- n guided candidates $\mathcal{G}_r^* \subseteq \mathcal{G}_r$ by FOM
 - 14: $\mathcal{P}_r = \mathcal{P}_{r-1} + \mathcal{G}_r^*$
 - 15: **return** $\arg \min_{\mathbf{x} \in \mathcal{A}_R} \text{FOM}(\mathbf{x})$
-

4.4. Rank-Weighted Generative Optimization

The second inverse framework inherits some features from the first one and combines two components: rank-weighted adaptation of a tabular generator and gradient-based surrogate guidance. Unlike strictly offline black-box optimization, the verified dataset is expanded after each optimization round through new SIMULATE3 evaluations. As stated previously, the surrogate affects only candidate generation. It is not used to determine objective values, update rankings, or select the final design. Thus, every accepted objective value is verified by the black-box simulator. The complete procedure is summarized in Algorithm 3, and its workflow is illustrated in Figure 3.




Figure 3: CFD-informed rank-weighted generative optimization with surrogate-guided proposal generation and SIMULATE3 verification.

At round r , candidates in $\mathcal{P}_{r-1}$ are sorted in ascending order of $\text{FOM}(\mathbf{x})$ in Equation (7). Let $\rho_i \in \{0, \dots, N-1\}$ denote the rank of candidate i , where $\rho_i = 0$ corresponds to the best candidate. We use the normalized rank

$$\bar{\rho}_i = \frac{\rho_i}{N-1} \quad (10)$$

which removes the direct dependence of the weighting scale on the pool size.

The rank-based sampling probability is obtained from softmax function as below;

$$p_i(T) = \frac{\exp(-\bar{\rho}_i/T)}{\sum_{j=1}^N \exp(-\bar{\rho}_j/T)} \quad (11)$$

here $T > 0$ controls the selection pressure: large T approaches uniform sampling, small T concentrates probability on the best-ranked candidates. A bootstrap training set $\mathcal{T}_r$ of size N is drawn with replacement from $\mathcal{P}_{r-1}$ using $p_i(T)$, so better-ranked candidates appear more often and gain more influence on the refitted generator. In contrast to deterministic elite-only training, every candidate retains nonzero sampling probability, and the weighting depends only on the ordering of verified objective values.

Because the rank-weighting temperature T determines how strongly retraining concentrates on the highest-ranked candidates, excessive selection pressure can reduce the effective diversity of the bootstrap distribution. This concentration is monitored using the relative effective sample size,

$$\text{ESS}_{\text{rel}} = \frac{1}{N \sum_{i=1}^N p_i(T)^2}, \quad (12)$$

where ESS_{rel} is the effective fraction of the N pool members represented by the rank-weighted sampling distribution and $p_i(T)$ is given by Equation (11). For nearly uniform probabilities, ESS_{rel} approaches unity, indicating that most candidates contribute meaningfully to resampling. As the distribution becomes concentrated on a small number of high-ranked candidates, ESS_{rel} decreases toward zero. If the effective sample size falls below a prescribed floor, the temperature T is relaxed to reduce selection pressure and restore diversity. Conversely, when optimization progress plateaus, the rank weighting and surrogate-guidance strength can be annealed toward stronger exploitation. This adaptive mechanism provides a balance between concentrating the generator around promising regions and maintaining sufficient diversity to avoid premature collapse of the search distribution.

The remaining components follow the surrogate-guided generation strategy described in Section 4.3. At each round, the surrogate f_ϕ is fitted to the current verified pool $\mathcal{P}_{r-1}$, while the generator G_θ is fitted to the rank-weighted bootstrap set $\mathcal{T}_r$. The updated generator then produces two candidate streams. Surrogate-guided candidates $\mathcal{X}_r^{\text{exploit}}$ use the gradient of f_ϕ during sampling to favor regions predicted to have lower FOM, whereas unguided candidates $\mathcal{X}_r^{\text{explore}}$ are sampled without this correction to preserve broader exploration of the learned distribution. The two streams are combined to form the candidate set $\mathcal{X}_r$. In this way, exploration and exploitation are retained within a single generative framework rather than through separate explorer and guided-generator models.

All candidates in $\mathcal{X}_r$ are evaluated independently using SIMULATE3. If a simulation completes successfully, its RMS, TV, and FOM are computed from the simulator output and the result is added to the cumulative verified archive $\mathcal{A}_r$. The surrogate therefore affects only the proposal-generation stage; surrogate predictions are not treated as objective values and do not determine the final ranking of candidates. This simulator-in-the-loop verification prevents candidates from being accepted solely because they exploit inaccuracies in the surrogate model.

The resulting framework therefore combines three complementary mechanisms: rank-weighted retraining shifts the learned generative distribution toward simulator-verified high-performing candidates, surrogate gradients guide individual samples toward lower predicted FOM, and unguided generation together with ESS-based adaptation preserves exploration. Every objective value used for pool refinement and final selection remains independently verified by SIMULATE3. The complete iterative procedure is summarized in Algorithm 3.

Algorithm 3 Rank-Weighted Generative Optimization

Require: CFD-informed verified pool $\mathcal{P}_0$, generator G_θ , surrogate f_ϕ

```
1: Initialize archive  $\mathcal{A}_0 \leftarrow \mathcal{P}_0$ 
2: for  $r = 1, \dots, R$  do
3:   Rank  $\mathcal{P}_{r-1}$  using the verified objective in Equation (7)
4:   Compute  $p_i(T)$  using Equation (11)
5:   Sample bootstrap set  $\mathcal{T}_r$  of pool size  $N$ 
6:   Fit  $G_\theta$  to  $\mathcal{T}_r$ 
7:   Fit  $f_\phi$  to the fixed-size pool  $\mathcal{P}_{r-1}$ 
8:   Generate guided candidates  $\mathcal{X}_r^{\text{exploit}}$ 
9:   Generate unguided candidates  $\mathcal{X}_r^{\text{explore}}$ 
10:   $\mathcal{X}_r = \mathcal{X}_r^{\text{exploit}} + \mathcal{X}_r^{\text{explore}}$ 
11:  for each  $\mathbf{x} \in \mathcal{X}_r$  do
12:    Evaluate  $\mathbf{x}$  using SIMULATE3
13:    if the simulation is valid then
14:      Compute verified RMS( $\mathbf{x}$ ), TV( $\mathbf{x}$ ), and FOM( $\mathbf{x}$ )
15:      Add the result to  $\mathcal{A}_r$ 
16:  Reject valid candidates below the Pearson-correlation threshold from pool admission
17:  Form  $\mathcal{P}_r$  as the  $N$  lowest-FOM candidates among  $\mathcal{P}_{r-1}$  and the newly admitted candidates
18:  Compute ESSrel and adapt  $T$  and  $\lambda$  (anneal on plateau; relax  $T$  if ESS below floor)
19: return  $\arg \min_{\mathbf{x} \in \mathcal{A}_R} \text{FOM}(\mathbf{x})$ 
```

5. Experiments

5.1. Computational environment

All optimization runs reported in this section were executed on an internal server available at the University of Michigan whose configuration is summarized in Table 2. The two strategies compared below, the hybrid explorer-exploit generator (Section 5.3) and the rank-weighted generative framework (Section 5.4), were run on this identical hardware and software stack.

Table 2: Hardware and software environment used for optimization runs

Component	Specification
CPU	AMD Ryzen Threadripper PRO 7985WX
GPU	NVIDIA RTX PRO 6000 Blackwell
Neutronics verification	SIMULATE3, up to 20 concurrent workers

5.2. CFD Prior Redundancy and Data Sufficiency

To determine how much unique information is contained in the densely sampled CFD trajectory, the ordered 5 s dataset was systematically thinned at increasing stride while retaining coverage of the full analysis window. Each thinned empirical prior was compared with the full 10,000-snapshot dataset using correlation- and proper-orthogonal-decomposition (POD)-based diagnostics. TabDDPM was then retrained on the corresponding subsets using the same model configuration and training budget to separate loss of information in the CFD prior from loss of generative-model fidelity.

Figure 4 shows that the dominant spatial prior is preserved under substantial thinning. The affinity of the leading 20 POD modes remains approximately 0.998 with 625 retained snapshots and 0.994 with 400, compared with approximately 0.978 at 200 snapshots and larger degradation at 100 and 50 snapshots. The TabDDPM results show a similar transition. With the full dataset, the unconditional generator achieves a leading-POD affinity of

approximately 0.95 and a classifier two-sample test (C2ST) score of approximately 0.54. At 400 retained snapshots, these values remain approximately 0.93 and 0.62, respectively, while fidelity degrades more clearly below a few hundred snapshots. Thus, under the fixed training procedure used here, roughly 400–625 snapshots are sufficient to preserve the dominant CFD prior and support a representative TabDDPM model.

This result also clarifies the approximately 4.2% brute-force RMS floor of the CFD-informed initial pool. Because the prior is structurally stable under strong thinning and is represented well by the full-data TabDDPM, the later reduction in RMS cannot be attributed simply to a poorly resolved or poorly learned CFD prior. Instead, the declining Pearson similarity during optimization is consistent with purposeful exploration beyond a well-learned prior toward fields that better reproduce the measured detector response.




Figure 4: Redundancy and data-sufficiency analysis of the CFD-informed prior and TabDDPM generator. Temporal thinning is used to assess preservation of the dominant CFD structure, and TabDDPM is retrained on the corresponding subsets using a fixed model configuration and training budget. The prior remains well preserved with approximately 400–625 retained snapshots, while generative fidelity degrades more clearly below this range.

5.3. Results of the Hybrid Exploration and Surrogate-Guided Optimization Framework

To decouple exploration from exploitation, we evaluated a hybrid exploration framework of Algorithm 2 in which the two sampling streams of each round are produced by two different tabular generators. A non-parametric resampling synthesizer (TabSDS) serves as the broad explorer, while a surrogate-guided diffusion model (TabDDPM) serves as the exploitation generator. The detailed hyperparameters used in this framework are given in Table 3.

The unmodified inlet flow distribution yielded a baseline RMS of approximately 5.1% [11]. Brute-force evaluation of all 10,000 CFD-informed inlet flow distributions using SIMULATE3 only reduced the best verified RMS to approximately 4.2%. Figure 5c shows the verified best and top- k average RMS across iterations. The RMS of a

Table 3: Hyperparameters for the hybrid exploration and surrogate-guided optimization (TabDDPM guided generator, TabSDS explorer)

Group	Parameter	Value
Optimization loop	Explorer E_ψ	TabSDS
	Guided generator G_θ	TabDDPM
	Guidance mode	gradient
	Pool growth	increasing
	Initial pool size $ \mathcal{P}_0 $	9999
	Pool admission per round n (top-FOM)	200
Per-round generation	Exploratory candidates (E_ψ)	2000
	Guided candidates (G_θ)	2000
Figure of merit / evaluation	Selection metric	normalized FOM (score)
	λ_{TV}	0.5
	TV_{limit}	3.0
	RMS_{init}	0.0512 (5.12%, from SIMULATE3)
Surrogate f_ϕ (guidance-only)	Target	normalized FOM (score)
	Hidden width $\times$ depth	256×4
	Training steps	5000
Guidance gates	Min. R^2 / Spearman to activate	0.5 / 0.5
Explorer E_ψ : TabSDS	Levels	10
	Proportion	0.5
	Shuffle type	simple
	Clip to train range	true
Generator G_θ : TabDDPM	MLP hidden layers	512×4
	Diffusion timesteps	200
	Noise scheduler	cosine
	Loss	MSE
	Training steps	5000

candidate is the root-mean-square of the relative differences between the SIMULATE3-computed and the measured detector reaction rates over the 58 instrumented assembly locations and is obtained using Equation (4). The *best* curve is the RMS of the single lowest-error candidate produced in each iteration, and the *top-k average* curve is the mean RMS over that iteration’s verified candidates; together they track the iteration’s leading solution and the overall quality of the guided batch. Both decrease steadily: the best verified candidate falls below 2% RMS by iteration 24 and converges to $RMS \approx 1.85\%$, while the top- k average decreases from about 4.5 to 2.3%. The gap between the two curves stays roughly constant after the early iterations, indicating that the hybrid generative model continued to surface high-quality candidates rather than collapsing the entire guided batch onto a single mode.

The improvement is also evident spatially in Figure 5a, which compares the SIMULATE3 reaction-rate distribution obtained from the best reconstructed field with the measured detector response. The optimized distribution reproduces the principal spatial features of the measurements across the instrumented assemblies. The corresponding assembly-level DLT field is shown in Figure 5b. Rather than independently perturbing individual assemblies, the optimized field retains coherent spatial structures originating from the CFD-informed initialization while introducing the modifications required to improve the neutronic response.

This tendency is identified by the Pearson-similarity trajectory in Figure 5d: the mean row-max Pearson correlation, computed by matching each verified candidate to its nearest CFD-training pattern under centered

Pearson correlation, decreases from 0.97 to 0.55 by iteration 44. Thus, the optimizer improves detector agreement while moving away from simple nearest-neighbor reproduction of the CFD-informed initialization.




Figure 5: Results of the hybrid exploration optimization framework: (a) comparison of the power profiles predicted by SIMULATE3 following CFD coupling through the generative model; (b) temperature distribution corresponding to the best power profile identified in (a); (c) best and mean RMS values throughout the guided search process; and (d) mean row-wise maximum Pearson similarity throughout the search.

5.4. Results of the Rank-Weighted Generative Optimization Framework

We evaluated the rank-weighted generative framework in Algorithm 3 using two tabular diffusion backbones: TabDDPM and TabDiff. The rank-weighted framework has the same simulator-in-the-loop structure as the hybrid exploration-exploitation framework, but it uses a single generator and a smaller verification batch to reduce computational cost. In both runs the verified pool was kept at a fixed size, and each iteration generated 500 candidates split into 300 surrogate-guided exploitation samples and 200 unguided exploration samples. The retraining set for the next generator was drawn from the verified pool using the Boltzmann rank weights of Equation (11). The hyperparameters are summarized in Table 4. Thus, unlike the hybrid explorer-generator model in Section 5.3, which used two different generators and 4,000 candidates per iteration, this experiment tests whether a single adaptively re-weighted generator can approach the same detector agreement with fewer SIMULATE3 evaluations.

In Figures 6 and 7, the best RMS candidate for each generator is displayed, while the evolution of their RMS values across optimization rounds is compared in Figure 8a. Table 5 summarizes the two rank-weighted runs and compares them with the hybrid exploration and surrogate-guided framework. The rank-weighted TabDDPM run

completed all 60 configured rounds and reached its best verified candidate at round 60 with $\text{RMS} = 2.08\%$ and $\text{TV} = 2.00$. This is substantially more economical in SIMULATE3 calls than the hybrid TabSDS-TabDDPM run.

The spatial distributions of the best verified solutions further illustrate the differences between the two rank-weighted variants. The TabDDPM solution in Figure 6 reproduces the measured reaction-rate pattern with an RMS of 2.08% (Figure 6a), while its corresponding temperature field retains coherent spatial variations with $\text{TV} = 2.00$ (Figure 6b). In comparison, the TabDiff solution exhibits larger spatial deviations from the measured reaction rates, consistent with its higher RMS of 3.12% (Figure 7a). However, its reconstructed temperature field is smoother, with a lower $\text{TV} = 1.44$ (Figure 7b). These spatial results are consistent with the quantitative metrics and highlight the trade-off observed between improved detector agreement with TabDDPM and greater spatial smoothness with TabDiff.

Figure 8b further shows distinct exploration behavior between the two generators. The mean Pearson similarity of TabDDPM decreases substantially from approximately 0.99 to 0.60, indicating progressive exploration away from the CFD-informed prior, whereas TabDiff remains comparatively stable near 0.75–0.80 throughout the optimization, suggesting more limited departure from the CFD prior and potentially explaining its higher final RMS.

Table 4: Hyperparameters for the rank-weighted generative optimization

Group	Parameter	Value
Optimization loop	Candidates per round	500
	Pool size $ \mathcal{P} $	9999
	Exploit / explore	0.6 / 0.4
	SIMULATE3 workers	20
Rank-weighted resampling	Initial temperature T_0	$0.5 \mathcal{P} $
	Minimum temperature $T_{\min}$	$0.05 \mathcal{P} $
Figure of merit / evaluation	λ_{TV}	0.5
	TV_{limit}	2.0
Surrogate f_ϕ (guidance-only)	Target	verified objective
	Hidden width $\times$ depth	256×4
	Training steps	5000
Guidance	Guidance strength λ_g	0.05
	Min. R^2 / Spearman to activate	0.5 / 0.5
Similarity gate	Metric	Pearson
	Minimum similarity	0.8
Generator G_θ : TabDDPM	MLP hidden layers	512×4
	Diffusion timesteps	200
	Noise scheduler	cosine
	Loss	MSE
	Training steps	5000
Generator G_θ : TabDiff	Transformer layers	2
	Token dimension / heads	4 / 1
	Diffusion timesteps	200
	Noise scheduler	power mean per column
	Sigma range / ρ	0.002–80 / 7
	Training steps	5000

Table 5: Comparison of the hybrid explorer-generator engine with the rank-weighted single-generator model

Method	Rounds	Best RMS (%)	TV
Hybrid TabSDS-TabDDPM	44	1.86	1.84
Rank-weighted TabDDPM	60	2.08	2.00
Rank-weighted TabDiff	59	3.12	1.44




Figure 6: Best verified rank-weighted TabDDPM candidate with Pearson threshold = 0.6




Figure 7: Best verified rank-weighted TabDiff candidate with Pearson threshold = 0.6




Figure 8: Comparison of rank-weighted TabDDPM and TabDiff generator

5.5. Ablation test

To isolate the contribution of spatial regularization in the rank-weighted framework, we performed two otherwise identical TabDDPM optimization runs using a centered-Pearson pool-admission threshold of 0.8. Both cases used the same condition described in Table 4. The only ablated parameter was the total-variation weight: the unregularized case used $\lambda_{TV} = 0$, whereas the regularized case used $\lambda_{TV} = 0.5$.


The best results are summarized in Table 6. Without TV regularization, the best candidate at round 71 achieved a detector mismatch of $RMS = 1.74\%$ and $TV = 6.03$. In comparison, the TV-regularized run produced a final-round best candidate at round 67 with $RMS = 3.87\%$ and $TV = 0.99$. The TV penalty reduced the spatial-variation metric while increasing the detector RMS.

Table 6: Ablation of TV regularization in the rank-weighted TabDDPM framework with a centered-Pearson pool-admission threshold of 0.8.


Configuration	Final round	RMS (%)	TV
Without TV regularization	71	1.74	6.03
With TV regularization	67	3.87	0.99

The reconstructed inlet-temperature fields in Figure 9 illustrate the physical consequence of this trade-off. When the TV term is omitted, the optimizer has sufficient freedom to introduce large and highly localized assembly-wise perturbations. These perturbations provide additional degrees of freedom for matching the measured detector reaction rates, resulting in the lower RMS shown in Figure 9a. However, the resulting field contains pronounced temperature differences between neighboring assemblies and lacks the coherent large-scale spatial structure expected from the CFD-informed inlet-flow distribution.

In contrast, the TV-regularized solution in Figure 9b exhibits gradual spatial variations across the core. Localized extrema are suppressed, and neighboring assemblies have substantially more consistent inlet-temperature perturbations. This improved spatial coherence reduces the flexibility available to match the detector measurements, explaining the higher RMS of the regularized solution. The comparison therefore demonstrates that minimizing detector RMS alone does not guarantee a spatially regular reconstruction, even when a strict Pearson threshold is used during pool admission.



(a) Without TV regularization: RMS = 1.74% and TV = 6.03.



(b) With TV regularization: RMS = 3.87% and TV = 0.99.

Figure 9: Final-round inlet-temperature perturbation fields obtained by rank-weighted TabDDPM optimization with a centered-Pearson pool-admission threshold of 0.8: (a) without TV regularization and (b) with TV regularization.

The ablation results indicate that the Pearson and TV constraints preserve different aspects of the CFD-informed prior. The Pearson criterion controls global pattern similarity during optimization-pool admission, whereas TV regularization directly suppresses local assembly-to-assembly variation. Consequently, a high Pearson threshold alone does not prevent locally oscillatory temperature fields. Explicit spatial regularization is necessary when the inferred inlet condition is required to remain smooth, although this requirement introduces a measurable compromise in detector agreement. The data-sufficiency study in Section 5.2 also indicates that unconditional TabDDPM samples are somewhat rougher than the original CFD prior, consistent with the need for an explicit local-smoothness constraint during optimization.

5.6. Discussions and Implications

Starting from the same CFD-informed prior, the hybrid explorer-generator framework achieved the lowest detector mismatch, reaching a best RMS of approximately 1.86%, compared with 2.08% for the rank-weighted TabDDPM framework. As shown in Figure 10, both methods exhibit similar reductions in the average RMS during the early optimization rounds, while the hybrid framework ultimately reaches a lower best RMS. This improvement comes at a higher computational cost because the hybrid approach employs separate exploration and guided-generation stages, an expanding candidate pool, and substantially more SIMULATE3 evaluations per round. In contrast, the rank-weighted framework maintains a fixed-size pool and uses rank-weighted resampling to progressively concentrate a single generator toward promising regions, providing a more computationally economical search at the expense of a slightly higher final RMS.

The computational cost of the two TabDDPM-based optimization strategies shown in Figure 10 is compared in Table 7. The completed hybrid archive therefore required about 3.5 $\times$ more elapsed time (4 h 40 m versus 1 h 20 m) and roughly 6 $\times$ more SIMULATE3 evaluations ($44 \times 4,000$ versus 60×500 candidates) than the rank-weighted TabDDPM run. Its larger verification batches produced a richer verified archive, but the total wall time was dominated by the much larger per-round candidate budget.

The generator backbone matters as well for the rank-weighted framework: TabDDPM, whose noise-prediction network directly admits sampling-time gradient guidance, adapted quickly in both frameworks, whereas TabDiff plateaued above 3% RMS but retained a nearly constant Pearson similarity to the CFD prior (Figure 8b), consistent with its guidance mechanism being coupled to training rather than to inference.




Figure 10: Convergence comparison between the hybrid explorer-generator and rank-weighted TabDDPM frameworks. Solid lines show the best verified detector RMS at each optimization round, while dashed lines show the corresponding average RMS. The hybrid framework reaches a lower final RMS, whereas the rank-weighted approach achieves comparable early convergence using a more computationally economical fixed-size optimization pool.

Table 7: Computational-time comparison between the hybrid explorer-generator run and the rank-weighted TabDDPM run.

Method	Completed rounds	Wall time
Hybrid TabSDS-TabDDPM	44	4h 40 m
Rank-weighted TabDDPM	60	1h 20 m

To isolate the effect of the CFD-manifold similarity constraint, we compared three otherwise similar rank-weighted TabDDPM trajectories using centered-Pearson admission thresholds of 0.6, 0.7, and 0.8. At each iteration, 2,000 synthetic fields were generated, with the similarity diagnostics evaluated using 200 sampled fields. As shown in Figure 11a, all three cases substantially reduce the detector RMS from the initial value; however, the achievable improvement depends on the imposed similarity constraint. The least restrictive threshold of 0.6 produces the lowest RMS, approaching approximately 2%, while the 0.7 case converges to approximately 2.2%. In contrast, the more restrictive 0.8 threshold remains at a higher RMS of approximately 2.3–2.4% over the evaluated iterations. The corresponding mean RMS trajectories show the same general ordering, indicating that the effect is reflected in the broader generated population rather than only in the single best candidate.

The associated similarity trajectories in Figure 11b clarify this behavior. As optimization proceeds, the mean row-maximum Pearson similarity progressively decreases toward the prescribed admission threshold (e.g., 0.6, 0.7, and 0.8). Thus, a higher similarity threshold constrains the generated fields to remain closer to the CFD-informed prior, whereas a lower threshold permits greater departure from the original CFD patterns. When considered together, Figures 11a and 11b show a clear trade-off between preserving similarity to the CFD prior and reducing the detector mismatch: allowing greater exploration away from the prior leads to lower RMS, while stronger preservation of CFD-like spatial patterns limits the attainable improvement.

The thinning study also suggests that future CFD-informed generative workflows need not retain every CFD time step for downstream model training. For the present trajectory, several hundred snapshots distributed across the transient recover nearly the same dominant spatial prior, so sparse export or post-processing can reduce data storage, downstream processing, and training costs without materially changing that prior. Sparse sampling may also reduce the computational overhead associated with recording assembly-wise mass flow rates, since each recorded state requires integration over the full core-inlet cross section.

This interpretation is subject to three caveats specific to the thinning and data-sufficiency analysis. The 5 s




Figure 11: Effect of the centered-Pearson similarity threshold on the rank-weighted TabDDPM optimization. (a) Evolution of the best (solid lines) and mean (dashed lines) detector RMS for Pearson admission thresholds of 0.6, 0.7, and 0.8. (b) Corresponding mean row-maximum Pearson similarity between the generated candidate fields and the CFD-informed prior dataset; horizontal dotted lines indicate the prescribed similarity thresholds. Lower thresholds permit greater departure from the CFD prior and achieve lower detector RMS, illustrating the trade-off between CFD-prior fidelity and objective-driven exploration.

analysis window is short relative to the slowest temporal variation represented by the POD modes, so the 400–625-state range is specific to the present trajectory. The relative effective sample size ESS_{rel} in the rank-weighted framework measures concentration of the resampling probabilities rather than geometric diversity of the candidate fields and can therefore overstate early-round diversity. Finally, the data-sufficiency threshold was obtained with a fixed TabDDPM architecture and training budget and should not be interpreted as a universal minimum sample count. Beyond these analysis-specific caveats, several broader limitations of the present study should also be considered.

First, fuel assembly bowing is represented through a surrogate variable rather than an explicit bowing model, not included in SIMULATE3. Instead, following [11], the HYD.DLT card is used as a proxy. Bowing displaces assemblies laterally, modifying inter-assembly water gaps and therefore the local moderator volume; DLT reproduces this effect indirectly by perturbing channel-wise inlet temperature, which acts on moderator density through the water equation of state. Since BEAVRS operates at hot zero power (the core is nominally isothermal), the DLT perturbations used here carry no physical thermal-hydraulic meaning: they are purely a mathematical handle on moderator density redistribution. This limits the physical interpretability of the reconstructed fields, though it does not affect the validity of the optimization and inference methodology itself. The proposed framework is directly extensible to a more explicit bowing model once it becomes available.

Second, the study is restricted to a single reactor operating condition at hot zero power (BEAVRS), one detector measurement set, and a fixed objective formulation. The relative performance of the two optimization frameworks may therefore change under other operating conditions, measurement configurations, or inverse problems with different neutronic sensitivities and inlet-flow structures. The reported computational comparison is also specific to the present SIMULATE3 verification cost. Because simulator evaluations constitute a major portion of the optimization expense, the relative importance of generative-model training and broad candidate exploration would diminish as the cost of the underlying physics calculation increases. Extending the proposed methodology to more computationally demanding multiphysics applications, such as those developed within the Multiphysics Object Oriented Simulation Environment (MOOSE) framework [1], would therefore provide an important test of its scalability and generalization beyond the present inverse problem.

Third, several parameters controlling the balance between physical-prior preservation and objective-driven exploration were prescribed for each run, including the Pearson similarity threshold, TV regularization weight, and

surrogate-guidance strength. The results indicate that these choices can directly influence both the achievable detector mismatch and the degree of departure from the CFD-informed prior. In this work, these parameters were intentionally kept fixed to limit the number of hyperparameters requiring manual tuning, as extensive hyperparameter optimization would introduce additional computational overhead and potentially require long execution times. Future work should therefore investigate adaptive strategies in which similarity constraints, guidance strength, and exploration–exploitation behavior evolve according to optimization progress, candidate diversity, or surrogate reliability rather than remaining fixed. Such extensions would be particularly valuable when coupling the framework to higher-fidelity or multiphysics solvers, where the number of affordable verification calculations may be substantially smaller and each generated candidate must be selected more efficiently.

6. Conclusion

This work developed a CFD-informed generative optimization framework for reactor core inverse reconstruction, in which assembly-level DLT fields are inferred from measured reaction-rate distributions. High-fidelity lower-plenum CFD simulations provide a physically informed prior, tabular generative models explore candidate inlet-condition fields, and a differentiable surrogate guides diffusion sampling toward improved solutions. Importantly, the surrogate is used only for sampling-time guidance, while every candidate is independently evaluated by the black-box neutronics solver SIMULATE3. Two optimization strategies were investigated: a hybrid explorer–generator framework combining the non-parametric TabSDS explorer with surrogate-guided TabDDPM, and a rank-weighted framework that iteratively adapts a single generator through rank-based resampling of a fixed-size verified pool.

Starting from a baseline root mean square error mismatch of approximately 5.1%, brute-force evaluation of the CFD-informed initial pool alone reached approximately 4.2%. Temporal-thinning and data-sufficiency analyses showed that the dominant CFD prior is preserved with roughly 400–625 retained snapshots and that the full-data TabDDPM represents this prior well. The hybrid framework further reduced the best verified RMS to approximately 1.86%, approaching the reported detector uncertainty [11]. The more computationally economical rank-weighted TabDDPM framework achieved 2.08%, while rank-weighted TabDiff reached 3.12% but produced a smoother final field with $TV = 1.44$. The decreasing Pearson similarity observed as RMS improved therefore indicates controlled exploration beyond a well-resolved CFD prior rather than compensation for a poorly learned prior. Collectively, these results demonstrate a trade-off between reconstruction accuracy, preservation of the CFD-informed structure, and computational cost: broader hybrid exploration produced the lowest mismatch, whereas rank-weighted adaptation provided competitive reconstruction with substantially fewer SIMULATE3 evaluations.

The present study is limited to a single reactor design, a single hot-zero-power condition, and a fixed objective formulation. Future work will extend the framework to additional operating conditions and higher-fidelity multiphysics problems, while developing adaptive strategies for controlling prior similarity, surrogate guidance, and exploration.

Acknowledgment

The authors gratefully acknowledge Robert A. Brewster and Emilio Baglietto for their valuable insights and contributions to this work through regular monthly meetings and discussions on the CFD data and model used in this study. This research is sponsored by the Department of Energy Office of Nuclear Energy through the Nuclear Energy University Programs (Award number: DE-NE0009506). The first author received sponsorship through the Korea Atomic Energy Research Institute Visiting Scholar Fellowship Program.

Data Availability

The authors have all the datasets and codes to reproduce all the results in this work currently in a private GitHub repository. To ensure the confidentiality of this research, the authors will make this repository public during an advanced stage of the review process, which will be listed under our research group’s public GitHub page (<https://github.com/aims-umich/>). Given the significant storage demands and export control issues, the raw STAR-CCM+ models and SIMULATE3 raw input files will be provided by the authors upon request. Nevertheless, the geometry CAD files will be made available on that GitHub repository.

References

- [1] C. J. Permann, D. R. Gaston, D. Andrš, R. W. Carlsen, F. Kong, A. D. Lindsay, J. M. Miller, J. W. Peterson, A. E. Slaughter, R. H. Stogner, et al., MOOSE: Enabling massively parallel multiphysics simulation, *SoftwareX* 11 (2020) 100430.
- [2] D. Price, N. Roskoff, M. I. Radaideh, B. Kochunas, Multiphysics modeling of heat pipe microreactor with critical control drum position search, *Nuclear Science and Engineering* (2024) 1–20.
- [3] S. L. Brunton, B. R. Noack, P. Koumoutsakos, Machine learning for fluid mechanics, *Annual review of fluid mechanics* 52 (1) (2020) 477–508.
- [4] D. Price, N. Roskoff, M. I. Radaideh, B. Kochunas, Thermal modeling of an evinci™-like heat pipe microreactor using openfoam, *Nuclear Engineering and Design* 415 (2023) 112709.
- [5] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: *Advances in Neural Information Processing Systems*, Vol. 33, 2020, pp. 6840–6851.
- [6] D. H. Brookes, H. Park, J. Listgarten, Conditioning by adaptive sampling for robust design, in: *Proceedings of the 36th International Conference on Machine Learning*, Vol. 97 of *Proceedings of Machine Learning Research*, 2019, pp. 773–782.
- [7] S. Krishnamoorthy, S. M. Mashkaria, A. Grover, Diffusion models for black-box optimization, in: *Proceedings of the 40th International Conference on Machine Learning*, Vol. 202 of *Proceedings of Machine Learning Research*, 2023, pp. 17842–17857.
- [8] G. Daras, H. Chung, C.-H. Lai, Y. Mitsufuji, J. C. Ye, P. Milanfar, A. G. Dimakis, M. Delbracio, A survey on diffusion models for inverse problems, *arXiv preprint arXiv:2410.00083*[doi:10.48550/arXiv.2410.00083](https://doi.org/10.48550/arXiv.2410.00083).
- [9] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, J. C. Ye, Diffusion posterior sampling for general noisy inverse problems, in: *International Conference on Learning Representations*, 2023.
- [10] N. Horelik, B. Herman, B. Forget, K. Smith, Benchmark for evaluation and validation of reactor simulations (BEAVRS), v1.0.1, in: *Proc. Int. Conf. Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, American Nuclear Society, Sun Valley, Idaho, 2013.
- [11] A. Persico, B. Forget, Preliminary Optimization of Fuel Bowing Surrogate Variables for the BEAVRS Benchmark, in: *The International Conference on Physics of Reactors (PHYSOR 2026)*, American Nuclear Society, Torino, Italy, 2026. [doi:https://doi.org/10.5281/zenodo.20803205](https://doi.org/10.5281/zenodo.20803205).
- [12] T. V. Santosh, G. Vinod, R. K. Saraf, A. K. Ghosh, H. S. Kushwaha, Application of artificial neural networks to nuclear power plant transient diagnosis, *Reliability Engineering & System Safety* 92 (10) (2007) 1468–1472. [doi:10.1016/j.res.2006.10.009](https://doi.org/10.1016/j.res.2006.10.009).
- [13] Y. Yao, D. Ge, J. Yu, M. Xie, Model-based deep transfer learning method to fault detection and diagnosis in nuclear power plants, *Frontiers in Energy Research* 10 (2022) 823395. [doi:10.3389/fenrg.2022.823395](https://doi.org/10.3389/fenrg.2022.823395).

- [14] G. Qian, J. Liu, A comparative study of deep learning-based fault diagnosis methods for rotating machines in nuclear power plants, *Annals of Nuclear Energy* 178 (2022) 109334. [doi:10.1016/j.anucene.2022.109334](https://doi.org/10.1016/j.anucene.2022.109334).
- [15] M. I. Radaideh, T. Kozlowski, Combining simulations and data with deep learning and uncertainty quantification for advanced energy modeling, *International Journal of Energy Research* 43 (14) (2019) 7866–7890.
- [16] M. I. Radaideh, T. Kozlowski, Analyzing nuclear reactor simulation data and uncertainty with the group method of data handling, *Nuclear Engineering and Technology* 52 (2) (2020) 287–295.
- [17] M. I. Radaideh, C. Pigg, T. Kozlowski, Y. Deng, A. Qu, Neural-based time series forecasting of loss of coolant accidents in nuclear power plants, *Expert Systems with Applications* 160 (2020) 113699.
- [18] R. A. Saleem, M. I. Radaideh, T. Kozlowski, Application of deep neural networks for high-dimensional large bwr core neutronics, *Nuclear Engineering and Technology* 52 (12) (2020) 2709–2716.
- [19] H. Gong, S. Cheng, Z. Chen, Q. Li, Data-enabled physics-informed machine learning for reduced-order modeling digital twin: Application to nuclear reactor physics, *Nuclear Science and Engineering* 196 (6) (2022) 668–693. [doi:10.1080/00295639.2021.2014752](https://doi.org/10.1080/00295639.2021.2014752).
- [20] H. Song, M. Song, X. Liu, Online autonomous calibration of digital twins using machine learning with application to nuclear power plants, *Applied Energy* 326 (2022) 119995. [doi:10.1016/j.apenergy.2022.119995](https://doi.org/10.1016/j.apenergy.2022.119995).
- [21] L. Tunkle, K. Abdulraheem, L. Lin, M. I. Radaideh, Nuclear microreactor transient and load-following control with deep reinforcement learning, *Energy Conversion and Management: X* 27 (2025) 101090.
- [22] M. I. Radaideh, L. Tunkle, D. Price, K. Abdulraheem, L. Lin, M. Elias, Multistep criticality search and power shaping in nuclear microreactors with deep reinforcement learning, *Nuclear Science and Engineering* (2025) 1–13.
- [23] L. A. Burnett, U. M. Nabila, M. I. Radaideh, Variational digital twins, *Energy and AI* (2026) 100756.
- [24] N. R. Panczyk, O. F. Erdem, M. I. Radaideh, Opening the ai black-box: Symbolic regression with kolmogorov-arnold networks for advanced energy applications, *Energy and AI* (2025) 100595.
- [25] C. Jendoubi, A. Asad, A survey of artificial intelligence applications in nuclear power plants, *IoT* 5 (4) (2024) 666–691. [doi:10.3390/iot5040030](https://doi.org/10.3390/iot5040030).
- [26] D. A. Ejigu, Y. Tuo, X. Liu, Application of artificial intelligence technologies and big data computing for nuclear power plants control: a review, *Frontiers in Nuclear Engineering* 3 (2024) 1355630. [doi:10.3389/fnuen.2024.1355630](https://doi.org/10.3389/fnuen.2024.1355630).
- [27] D. R. Jones, M. Schonlau, W. J. Welch, Efficient global optimization of expensive black-box functions, *Journal of Global Optimization* 13 (4) (1998) 455–492. [doi:10.1023/A:1008306431147](https://doi.org/10.1023/A:1008306431147).
- [28] N. V. Queipo, R. T. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, P. K. Tucker, Surrogate-based analysis and optimization, *Progress in Aerospace Sciences* 41 (1) (2005) 1–28. [doi:10.1016/j.paerosci.2005.02.001](https://doi.org/10.1016/j.paerosci.2005.02.001).
- [29] M. I. Radaideh, T. Kozlowski, Surrogate modeling of advanced computer simulations using deep gaussian processes, *Reliability Engineering & System Safety* 195 (2020) 106731.
- [30] A. Figueroa, M. Götsche, Gaussian processes for surrogate modeling of discharged fuel nuclide compositions, *Annals of Nuclear Energy* 156 (2021) 108085. [doi:10.1016/j.anucene.2020.108085](https://doi.org/10.1016/j.anucene.2020.108085).
- [31] X. Bei, Y. Dai, K. Yu, M. Cheng, Three-dimensional surrogate model based on back-propagation neural network for key neutronics parameters prediction in molten salt reactor, *Energies* 16 (10) (2023) 4044. [doi:10.3390/en16104044](https://doi.org/10.3390/en16104044).
- [32] P. Chakroborty, S. L. N. Dhulipala, Y. Che, W. Jiang, B. W. Spencer, J. D. Hales, M. D. Shields, General

- multifidelity surrogate models: Framework and active-learning strategies for efficient rare event simulation, *Journal of Engineering Mechanics* 149 (12) (2023) 04023096. doi:[10.1061/JENMDT.EMENG-7111](https://doi.org/10.1061/JENMDT.EMENG-7111).
- [33] X. Zheng, H. Tamaki, T. Sugiyama, Y. Maruyama, Dynamic probabilistic risk assessment of nuclear power plants using multi-fidelity simulations, *Reliability Engineering & System Safety* 223 (2022) 108503. doi:[10.1016/j.resss.2022.108503](https://doi.org/10.1016/j.resss.2022.108503).
- [34] M. Eaheart, M. I. Radaideh, Multifidelity surrogate modeling of onset of natural circulation following a depressurized loss of forced cooling in high-temperature gas reactors, *Applied Thermal Engineering* 303 (5) (2026) 132489.
- [35] O. F. Erdem, D. P. Broughton, J. Svoboda, C. Huang, M. I. Radaideh, Surrogate-driven design optimization with uncertainty constraints in monte carlo simulations, *Energy and AI* (2025) 100655.
- [36] A. Kumar, P. V. Tsvetkov, A new approach to nuclear reactor design optimization using genetic algorithms and regression analysis, *Annals of Nuclear Energy* 85 (2015) 27–35.
- [37] M. I. Radaideh, B. Forget, K. Shirvan, Large-scale design optimisation of boiling water reactor bundles with neuroevolution, *Annals of Nuclear Energy* 160 (2021) 108355.
- [38] X. Gu, M. I. Radaideh, J. Liang, Openneomc: A framework for design optimization in particle transport simulations based on openmc and neorl, *Annals of Nuclear Energy* 180 (2023) 109450.
- [39] P. I. Frazier, A tutorial on bayesian optimization, arXiv preprint arXiv:1807.02811doi:[10.48550/arXiv.1807.02811](https://doi.org/10.48550/arXiv.1807.02811).
- [40] B. Trabucco, X. Geng, A. Kumar, S. Levine, Design-bench: Benchmarks for data-driven offline model-based optimization, in: *Proceedings of the 39th International Conference on Machine Learning*, Vol. 162 of *Proceedings of Machine Learning Research*, PMLR, 2022, pp. 21658–21676.
- [41] B. Trabucco, A. Kumar, X. Geng, S. Levine, Conservative objective models for effective offline model-based optimization, in: *Proceedings of the 38th International Conference on Machine Learning*, Vol. 139 of *Proceedings of Machine Learning Research*, PMLR, 2021, pp. 10358–10368.
- [42] L. A. Burnett, H. Kim, H.-C. Chou, A. Witoelar, R. A. Brewster, B. Forget, E. Baglietto, M. I. Radaideh, High-fidelity modeling of pressurized water reactor flow fields for data-driven applications, *Energy* 362 (2026) 142104.
- [43] A. Kotelnikov, D. Baranchuk, I. Rubachev, A. Babenko, TabDDPM: Modelling tabular data with diffusion models, in: *Proceedings of the 40th International Conference on Machine Learning*, Vol. 202 of *Proceedings of Machine Learning Research*, PMLR, 2023, pp. 17564–17579.
- [44] P. Dhariwal, A. Nichol, Diffusion models beat gans on image synthesis, in: *Advances in Neural Information Processing Systems*, Vol. 34, 2021, pp. 8780–8794.
- [45] J. Ho, T. Salimans, Classifier-free diffusion guidance (2022). [arXiv:2207.12598](https://arxiv.org/abs/2207.12598).
- [46] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, J. C. Ye, Diffusion posterior sampling for general noisy inverse problems (2024). [arXiv:2209.14687](https://arxiv.org/abs/2209.14687).
- [47] J. Pevey, B. Hiscox, A. Williams, O. Chvála, V. Sobes, J. W. Hines, Gradient-informed design optimization of select nuclear systems, *Nuclear Science and Engineering* 196 (12) (2022) 1559–1571.
- [48] M. I. Radaideh, K. Shirvan, Pesa: Prioritized experience replay for parallel hybrid evolutionary and swarm algorithms-application to nuclear fuel, *Nuclear Engineering and Technology* 54 (10) (2022) 3864–3877.
- [49] M. I. Radaideh, I. Wolverton, J. Joseph, J. J. Tusar, U. Otgonbaatar, N. Roy, B. Forget, K. Shirvan, Physics-informed reinforcement learning optimization of nuclear assembly design, *Nuclear Engineering and Design* 372 (2021) 110966.

- [50] M. I. Radaideh, K. Du, P. Seurin, D. Seyler, X. Gu, H. Wang, K. Shirvan, Neorl: Neuroevolution optimization with reinforcement learning, arXiv preprint arXiv:2112.07057.
- [51] M. I. Radaideh, K. Du, P. Seurin, D. Seyler, X. Gu, H. Wang, K. Shirvan, Neorl: Neuroevolution optimization with reinforcement learning—applications to carbon-free energy systems, Nuclear Engineering and Design 412 (2023) 112423.
- [52] D. H. Brookes, H. Park, J. Listgarten, Conditioning by adaptive sampling for robust design, in: Proceedings of the 36th International Conference on Machine Learning, PMLR, 2019, pp. 773–782.
- [53] J. Shi, M. Xu, H. Hua, H. Zhang, S. Ermon, J. Leskovec, Tabdiff: a mixed-type diffusion model for tabular data generation (2025). [arXiv:2410.20626](https://arxiv.org/abs/2410.20626).
- [54] A. Guzmán-Cordero, F. Eijkelboom, J.-W. van de Meent, Exponential family variational flow matching for tabular data generation (2025). [arXiv:2506.05940](https://arxiv.org/abs/2506.05940).
- [55] E. C. Neto, TabSDS: a lightweight, fully non-parametric, and model free approach for generating synthetic tabular data, in: Forty-second International Conference on Machine Learning, 2025.
- [56] K. Rempe, K. S. Smith, A. Henry, Simulate-3 pin power reconstruction: methodology and benchmarking, Nuclear Science and Engineering 103 (4) (1989) 334–342.