Active Learning for Physics-Informed Digital Twins of Integrated Thermal Energy Distribution Systems

Umme Mahbuba Nabila[a], Paul Seurin[c], Linyu Lin[c], Majdi I. Radaideh[a,b,*]

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

[b] *Department of Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109, United States*

[c] *Nuclear Science & Technology Division, Idaho National Laboratory, Idaho Falls, ID 83415, United States*

Abstract

The Thermal Energy Distribution System (TEDS) at Idaho National Laboratory (INL) provides a unique experimental platform for testing advanced supervisory control strategies in hybrid energy systems that combine renewable, nuclear, and thermal energy storage (TES) resources. Real-time control of such systems requires surrogate models that are accurate, interpretable, and uncertainty-aware. This work presents an active learning (AL) framework that integrates high-fidelity Modelica simulations with physics-informed and data-driven surrogates to construct a digital twin (DT) of the TEDS glycol heat exchanger (GHX) subsystem. Four surrogate variants are examined: the deterministic Sparse Identification of Nonlinear Dynamics with Control (SINDyC), its probabilistic multivariate-Gaussian extension (MvG-SINDyC), a feedforward neural network (FNN), and a gated recurrent unit (GRU) network. These models are trained to reproduce the GHX transient dynamics and are compared in terms of predictive accuracy, interpretability, and computational efficiency. The AL loop iteratively selects the most informative simulation trajectories, accelerating convergence and reducing the training demand relative to random sampling. Two model-specific query strategies underpin the framework: Mahalanobis-distance sampling in the coefficient space for MvG-SINDyC and error-based sampling in the prediction space for SINDyC, FNN, and GRU. Across both GHX outputs—the bypass mass flow rate $\dot{m}_{\mathrm{GHX}}$ and heat transfer rate Q_{GHX}—AL substantially improves data efficiency, achieving comparable accuracy with as few as one-fifth of the trajectories required by random sampling. Among the evaluated surrogates, the GRU network achieves the highest predictive fidelity, with root mean square errors (RMSE) below 0.003 kg/s and 1 W. The deterministic SINDyC model remains the lightest and fastest to train, while its probabilistic extension (MvG-SINDyC) provides uncertainty quantification through multivariate-Gaussian inference and exhibits the largest computational gains under AL. The FNN surrogate shows overfitting tendencies, particularly without experimental supervision. Overall, the proposed AL-driven workflow offers a scalable pathway for constructing adaptive, interpretable, and uncertainty-aware digital twins for real-time supervisory control of complex energy systems.

Keywords: Digital Twin, Integrated Energy Systems, Sparse Identification of Nonlinear Dynamics with Control, Active Learning, Surrogate Modeling

*Corresponding Authors: Umme Nabila (unabila@umich.edu), Majdi I. Radaideh (radaideh@umich.edu)

Nomenclature

AL	active learning
DT	digital twin
FNN	feedforward neural network
GHX	glycol heat exchanger
GRU	gated recurrent unit
IES	integrated energy system
ML	machine learning
MPC	model predictive control
MvG-SINDyC	multivariate-Gaussian SINDyC
RMSE	root mean squared error
SINDyC	Sparse Identification of Nonlinear Dynamics with Control
TES	thermal energy storage
UQ	uncertainty quantification

1. Introduction

The Thermal Energy Distribution System (TEDS), located in the Dynamic Energy Transport and Integration Laboratory at Idaho National Laboratory (INL), serves as an experimental platform for studying hybrid energy systems that co-generate electricity and non-electric products [1]. Globally, integrated energy systems (IES) have emerged as a key pathway toward deep decarbonization and energy resilience, coupling nuclear heat and electricity with renewables, storage, and multi-product offtakes such as hydrogen, water, and process heat. The United States programmatic efforts being led by the U.S. Department of Energy and INL outline the vision, technical gaps, and demonstration roadmap pertaining to nuclear-driven IES [2], while recent studies have analyzed the related architectures, markets, and deployment barriers, including the roles of thermal storage, supervisory control, and verification/validation [3, 4, 5]. On the international scene, modular and object-oriented modeling frameworks (e.g., Modelica) have enabled plug-and-play nuclear-hydrogen-thermal energy storage (TES) configurations and scenario testing under realistic operating conditions [6]. At the planning scale, technology-rich optimization frameworks such as the TIMES-Europe model integrate power, industrial, and transport sectors across the EU-27+UK, leading to policy-relevant decarbonization scenarios that complement component-level analyses [7]. Taken together, these efforts highlight the need for digital twins (DTs) and advanced control strategies capable of managing tightly coupled assets under dynamic, uncertain conditions.

A substantial body of modeling and simulation research underpins this momentum. At INL, TEDS was first developed as a high-fidelity Modelica model to explore operating modes and component interactions, thereby providing a benchmark for subsequent experiments [1]. Follow-on work analyzed nuclear-TES arbitrage and supervisory control by using dynamic models [8] and this work was then extended to nuclear-renewable-hydrogen systems through hierarchical Modelica architectures [9]. Complementary studies using OpenModelica demonstrated hybrid desalination concepts that integrate TES and improve grid flexibility [10]. At the system level, optimization tools have been cross-validated with dynamic simulations for islanded microgrids that involve molten-salt TES [11], while real-time, hardware-in-the-loop experiments evaluated power-quality and control issues regarding nuclear-renewable-hydrogen integration [12]. Collectively, these multi-institutional efforts established a robust modeling and verification/validation foundation that the present work extends via a data-efficient, uncertainty-calibrated surrogate and an active learning (AL) strategy tailored for DT and model predictive control (MPC) applications in TEDS.

Traditionally, TEDS components have been operated using decentralized proportional-integral controllers. However, this type of control fails to capture cross-component coupling and system constraints. To overcome these limitations, supervisory controllers are being developed to coordinate multi-component signals and optimize system-wide performance [13]. These strategies improve efficiency and responsiveness, offering faster, more reliable control

than that afforded by manual operation [14]. Of the various advanced control methods, MPC stands out for its ability to explicitly incorporate system dynamics and constraints into an optimization problem solved over a prediction horizon [15]. At each time step, MPC applies only the first optimized input, then re-solves the problem by applying updated measurements, thus enabling adaptive, real-time control under uncertainty. This predictive capability makes MPC particularly suitable for complex, multi-input/multi-output systems such as TEDS [16].

To support such advanced control, a detailed physics-based model of TEDS was developed in Modelica (executed in Dymola) during its pre-experimental phase [17]. While accurate, these high-fidelity models are computationally prohibitive for real-time MPC, due to solver latency and numerical stiffness. Effective MPC therefore requires surrogate models that are fast, accurate, interpretable. Surrogates approximate system dynamics by using data from high-fidelity simulations, thus reducing system dimensionality and enabling orders-of-magnitude faster computation.

Of the different approaches available, Sparse Identification of Nonlinear Dynamics with Control (SINDyC) has gained attention for its simplicity, interpretability, and ability to recover governing equations directly from data [18, 19]. SINDyC yields analytical expressions that facilitate physical insights and direct embedding within optimization routines. Furthermore, its differentiable form supports gradient-based optimization within MPC loops. However, robust control requires not only speed and interpretability, but also reliable uncertainty quantification (UQ). To address this, a probabilistic extension—*multivariate Gaussian SINDyC (MvG-SINDyC)*—was proposed, aggregating thousands of deterministic SINDyC realizations into a single probabilistic surrogate [16]. In earlier implementations, trajectory selection relied on random sampling, which is simple yet insensitive to information content. Following the PySINDy convention, each simulation dataset is referred to herein as a *trajectory* [20].

Building on the probabilistic MvG-SINDyC foundation, we introduced an AL strategy to achieve data-efficient surrogate modeling for TEDS—and more broadly, for IES. Rather than relying on random or exhaustive sampling, the proposed AL framework adaptively identifies the most informative simulation trajectories by using two complementary query strategies: (1) *Mahalanobis-distance sampling* in the coefficient space for MvG-SINDyC, and (2) *error-reduction sampling* in the prediction space for neural surrogates. The Mahalanobis distance quantifies model coefficient deviations relative to the experimental baseline, accounting for both scale and correlation [21]. This targeted selection accelerates, using fewer simulations, convergence toward accurate, uncertainty-calibrated surrogates, thus enabling scalable DTs for MPC but without compromising robustness or interpretability.

While physics-informed surrogates such as SINDyC provide interpretability and computational efficiency, they are limited in capturing nonlinear and temporal dynamics when key states are unobserved or only partially measurable. In contrast, data-driven models based on machine learning (ML) technology such as a feedforward neural network (FNN) or gated recurrent unit (GRU) based on recurrent neural works offer superior predictive capabilities for nonlinear and sequential systems alike [22, 23, 24], and have been successfully applied as surrogate models in a range of studies [25, 26]. However, these black-box models typically require large datasets in order to generalize reliably, and this represents a challenge in data-limited regimes involving costly experiments or computationally intensive simulations, and lack of interpretability.

This study therefore extends the AL framework beyond SINDyC so as to include FNN and GRU surrogates, enabling a unified comparison between interpretable, uncertainty-calibrated, and high-fidelity black-box models. The comparison reveals key trade-offs among interpretability, UQ, and predictive fidelity, and highlights how targeted AL sampling improves accuracy, data efficiency, and robustness across modeling paradigms. Ultimately, this work lays the foundation for surrogate modeling approaches that combine physics-based transparency with the adaptability of ML—an essential step toward creating trustworthy, self-improving DTs for complex energy systems. This comparative study contributes to the advancement of DTs for IES in three key ways:

1. It demonstrates that AL can substantially accelerate surrogate-model convergence across diverse modeling

approaches—including interpretable physics-informed surrogates (MvG-SINDyC) and black-box neural networks (FNN and GRU)—thereby improving both data efficiency and predictive accuracy.

2. It clarifies the trade-offs among *interpretability, UQ, and predictive fidelity*, offering a framework for selecting or hybridizing surrogates based on the operational requirements of DTs and MPC controllers.
3. It validates the trained surrogate models by comparing their predictions against actual experimental measurements.

The remainder of this paper is organized as follows. Section 3 describes the experimental data and simulation models used for model training and validation. Section 4 details the theoretical foundations and methodologies employed for surrogate model construction, including MvG-SINDyC, FNN, and GRU, and introduces the conceptual framework for applying AL to DT development. Section 5 presents a comparative analysis of AL versus random sampling across all the modeling approaches explored, and discusses the implications of these findings for scalable, reliable DTs in IES applications. Finally, Section 6 summarizes the key contributions and future research directions.

2. Related Work

Recent research at the intersection of machine learning, optimization, and low-carbon energy systems spans surrogate modeling, reinforcement learning (RL), microreactor design and control, digital twins, and fault detection across integrated multi-energy infrastructures. In the area of surrogate modeling, deep Gaussian processes have been developed to emulate highly nonlinear nuclear simulations with improved uncertainty characterization [27], and surrogate-driven variance-based sensitivity analysis has been applied to thermal storage tanks in integrated energy systems [28]. Additional reduced-order techniques such as simplified matching pursuits enable fast reconstruction of three-dimensional temperature fields from sparse sensor measurements [29]. Complementary work in the broader energy-systems community includes the use of machine-learning surrogates for the EnergyPLAN framework to speed up country-level low-carbon energy system optimization [30], as well as systematic studies of surrogate model architecture, sampling, and scaling choices for multi-energy system design [31]. Data-driven distributionally robust formulations further leverage historical data to construct ambiguity sets for stochastic scheduling and dispatch in community and integrated energy systems [32, 33].

Reinforcement learning and advanced optimization have also seen rapid progress in low-carbon energy systems. Physics-informed RL has been applied to optimize nuclear assembly design under reactor-physics constraints [34], while prioritized experience replay mechanisms accelerate hybrid evolutionary–swarm search algorithms for nuclear fuel optimization [35]. These developments culminated in an optimization framework, which unifies evolutionary algorithms, neuroevolution, and RL for scalable optimization of carbon-free energy systems [36]. In parallel, deep RL and multi-agent RL have been employed to realize low-carbon economic dispatch and optimal scheduling for a variety of integrated energy systems, including electricity–gas coupled systems [37], community integrated energy systems under multiple uncertainties [38], transient, load-following control, and criticality search for nuclear microreactors [39, 40], integrated electricity–heat–hydrogen systems [41], multi-integrated energy systems with hierarchical control and demand response [42], and integrated electricity–heat–gas–hydrogen systems with demand-side management [43]. These data-driven optimization and RL-based control strategies provide a natural extension of nuclear-focused RL frameworks to broader low-carbon multi-energy contexts. At the system level, related work in integrated energy systems investigates data-driven optimal dispatch for zero-carbon building systems that explicitly consider occupant comfort and uncertainty [44], as well as distributionally robust and low-carbon economic scheduling strategies for community and regional integrated energy systems [32, 38, 33].

Digital-twin research has advanced through variational formulations that integrate latent-space inference with physics-based priors for real-time state estimation and uncertainty tracking [45], complemented by fast surrogate-based reconstruction and reduced-order modeling that support online monitoring and control. In parallel, significant contributions have been made in fault detection and predictive diagnostics for safety-critical energy components. Neural time-series forecasting has been used to predict the evolution of loss-of-coolant accidents from plant signals [46], while convolutional and feedforward neural networks have been demonstrated for fault detection in accelerator power systems [47]. Recurrent and ConvLSTM autoencoder models enable high-fidelity anomaly detection in power-electronics signals [48], and ensemble learning approaches further improve robustness across diverse fault types and operating regimes [49]. Moreover, multi-module conditional variational autoencoders have been applied to accurately predict high-voltage converter modulator faults in particle accelerators [50]. Collectively, this body of work establishes a coherent methodological foundation in which surrogate-assisted modeling, RL-based optimization, and intelligent fault diagnostics are key enablers for future low-carbon integrated energy systems, including nuclear power.

3. Experiment and Simulation Data Collection

3.1. TEDS Experimental Facility

TEDS, a flagship testbed within INL's Dynamic Energy Transport and Integration Laboratory, was designed to investigate the dynamics of coupled heat transfer, energy storage, and distribution components in support of IES applications [13]. It serves as a controllable, instrumented platform for validating reduced-order models, developing supervisory control strategies, and benchmarking DT methodologies.

Fig. 1 shows a schematic of the TEDS facility, illustrating the heater, TES tank, thermocline, fluid storage tanks, and control valves [51]. TEDS is comprised of a 200 kW Chromalox electrical heater, a single-tank packed-bed TES, and a glycol heat exchanger (GHX), along with the associated piping, five control valves, and distributed temperature and flow sensors. The TES consists of a vertical cylindrical tank filled with high-conductivity alumina beads (0.125 in.) that act as a filler medium, reducing the required fluid inventory and enhancing thermal stratification. Charging occurs when hot Therminol-66 is introduced at the top of the tank while cooler fluid is withdrawn from the bottom, gradually forming a thermocline. During discharging, the flow is reversed, with colder fluid being sent to the bottom and warmer fluid being extracted from the top, enabling efficient energy release. The GHX, constructed as a shell-and-tube heat exchanger, transfers heat between the Therminol loop and a secondary ethylene glycol loop, which emulates industrial or grid-connected demand. The Chromalox heater supplies high-temperature fluid to initiate charging cycles and regulate operating conditions.

The experiment cycles through five primary operating phases: startup (warm-up), charging, standby, discharging, and cooldown. For modeling and control, the TES state vector includes the inlet mass flow rate $\dot{m}_{\mathrm{tes,in}}$, outlet temperature $T_{\mathrm{tes,out}}$, and three internal node temperatures T_{top}, T_{mid}, T_{bot} that resolve the thermocline. The GHX state vector is represented by the bypass flow $\dot{m}_{\mathrm{ghx,bypass}}$ and the extracted heat rate Q_{ghx}. Actuation is effected by four controllable variables: valve position PV_{006}, pump outlet flow $\dot{m}_{\mathrm{pump,out}}$, inlet-loop temperature $T_{\mathrm{pump,in}}$, and heater outlet temperature $T_{\mathrm{heater,out}}$. Together, these inputs govern the coupled TES-GHX dynamics studied herein.

The present study emphasizes the discharging process ($\approx 9,180$–$14,640$ seconds) via which the packed-bed TES delivers stored heat to the GHX loop. The supervisory controller actuates PV_{006} and the circulation pump to route fluid from the TES through the GHX, as illustrated in Fig. 2. The discharge behavior (Experiment) is observed in Fig. 4. In these plots, $\mathrm{PV}_{006} = 1$ indicates that the valve is open to the GHX path, while $\mathrm{PV}_{006} = 0$ indicates that it is closed and that the flow is redirected through the bypass line. When PV_{006} is open, hot fluid from the upper TES

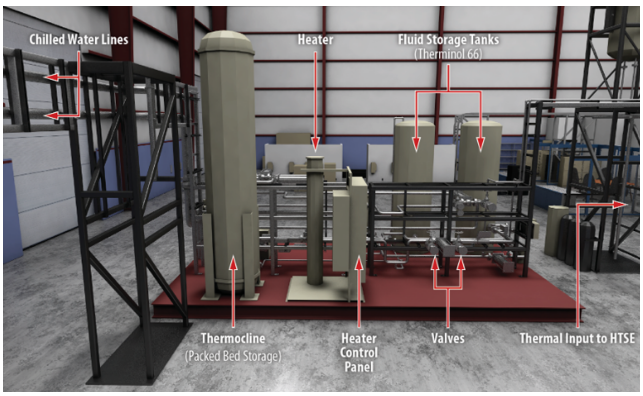

Figure 1: Experimental TEDS facility, showing the heater, TES tank, thermocline, fluid storage tanks, and control valves [51].

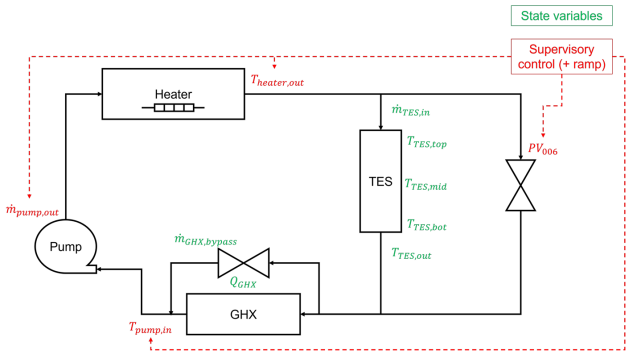

Figure 2: Simplified schematic of the TEDS loop and supervisory control architecture [16].

layer flows through the GHX, transferring heat to the glycol loop. As a result, the GHX heat rate Q_{GHX} increases rapidly, while the bypass flow $\dot{m}_{\mathrm{GHX,bypass}}$ decreases toward zero. The pump maintains an almost constant outlet flow, ensuring stable loop circulation. Over time, the TES gradually cools from the top downward as the hot region (thermocline) moves lower in the bed. This reduces the temperature difference driving the heat transfer, leading to a slow decline in Q_{GHX}. At around 2,000 seconds, PV_{006} switches from open to closed ($1 \rightarrow 0$). The flow is then redirected through the bypass, causing $\dot{m}_{\mathrm{GHX,bypass}}$ to rise sharply while Q_{GHX} drops nearly to zero. The present analysis specifically considers the GHX states $\dot{m}_{\mathrm{ghx,bypass}}$ and Q_{ghx}, which are critical for characterizing TEDS's transient heat delivery capability and for evaluating surrogate model predictions. This configuration allows simultaneous exploration of trade-offs between model fidelity, interpretability, and computational tractability—all while being anchored to a high-fidelity experimental dataset.

3.2. Modelica Simulation Framework

While the TEDS experiment provided high-fidelity measurements of coupled TES-GHX dynamics, only a single discharge trajectory was available for each quantity of interest. To overcome this data scarcity and enable robust surrogate training, we leveraged a physics-based DT of TEDS in Modelica, as implemented through the commercial platform Dymola [17]. As shown in the TEDS Modelica framework in Fig. 3, the simulation environment mirrors the experimental configuration, including the TES packed bed and the GHX loop, valves, heater, and pump, thus providing a virtual platform for controlled trajectory generation.

To enrich the training set, synthetic actuator trajectories were generated around the experimental operating points. Specifically, the ranges of the four actuators—valve position PV_{006}, pump outlet flow $\dot{m}_{\mathrm{pump,out}}$, pump inlet temperature $T_{\mathrm{pump,in}}$, and heater outlet temperature $T_{\mathrm{heater,out}}$—were expanded beyond the nominal discharge setpoints. A Sobol sequence sampler was then employed to draw 500 candidate time series across this expanded operating envelope. This quasi-random sampling ensured good coverage of the high-dimensional control space while also avoiding clustering effects typical of purely random draws.

Each actuator signal was refined to a common resolution of 5,251 points over a simulation window of 5,460 seconds, as is consistent with the experimental discharge duration. The experimental and simulated trajectories are compared in Fig. 4, with the black curves denoting the 500 Modelica-generated signals and the blue curves representing the de-noised experimental trajectories, filtered via a Savitzky-Golay filter. The figure highlights how the synthetic trajectories explore a broader range of dynamics while still remaining consistent with the experimental observations.

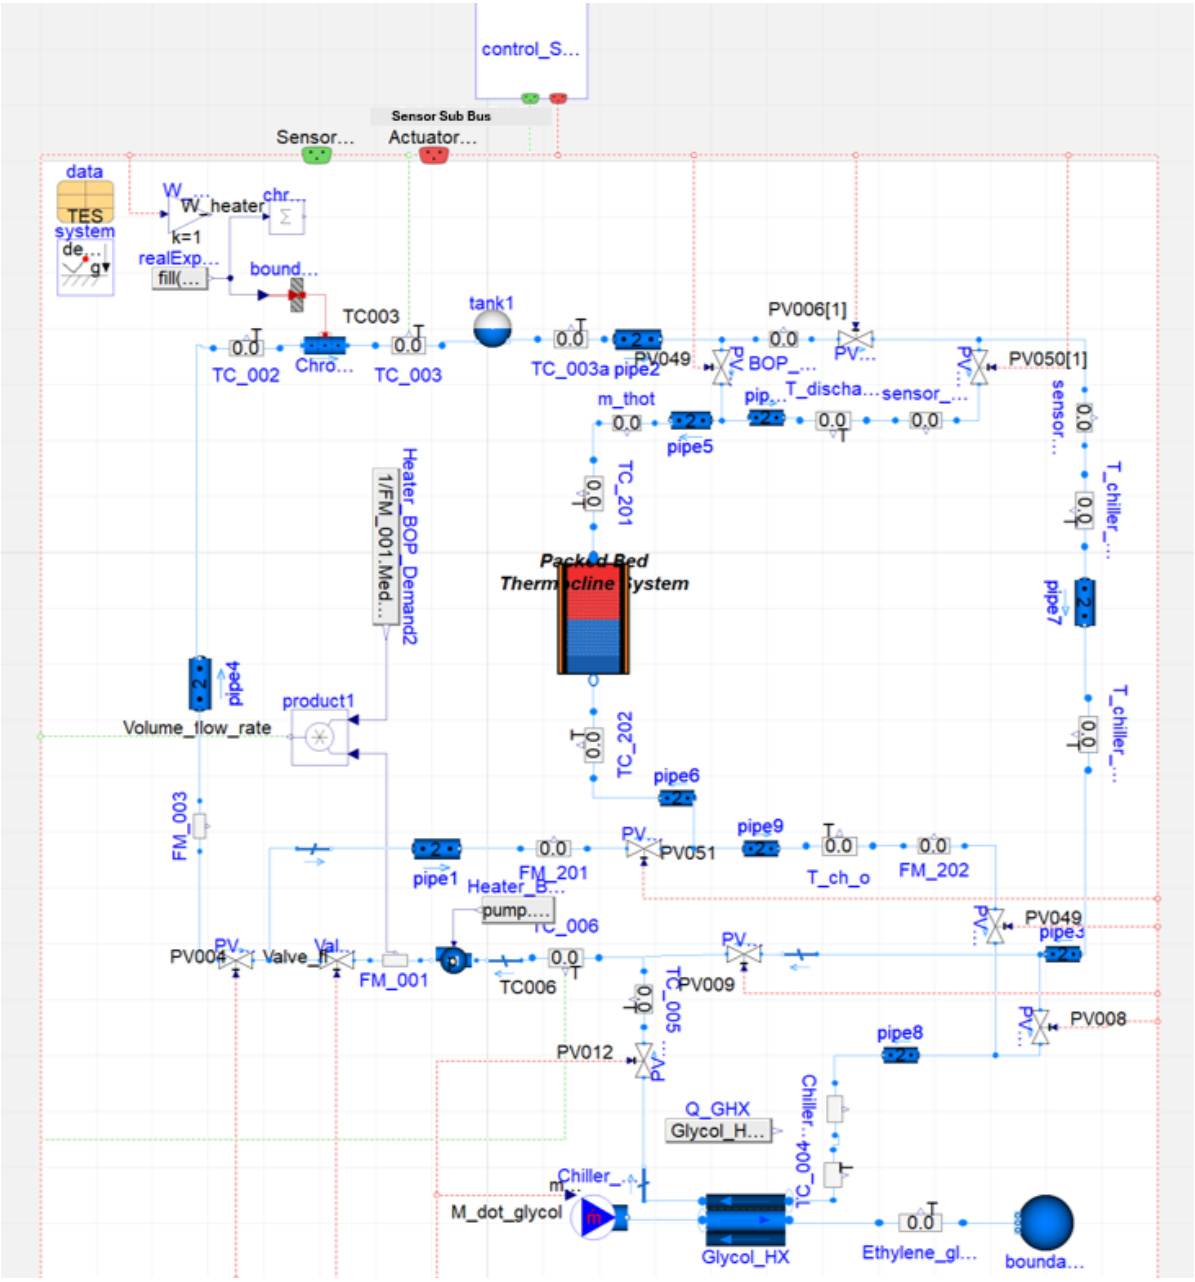

Figure 3: TEDS Modelica framework.

4. Methodology

4.1. Digital Twin Framework

Integration of experimental data and physics-based simulations is achieved through a DT framework for supervisory control of TEDS. As illustrated in Fig. 5, the DT connects the physical system, high-fidelity Modelica simulations, and real-time surrogate models.

In the ideal route, surrogate models would be trained and updated directly from continuous experimental measurements of temperature, flow, and heat rate $(T, \dot{m}, Q)$, thereby closing the DT loop with physical data. However, as the experimental cost and availability constraints make this infeasible, the current route is adopted: physics-based simulations are used to generate synthetic data, which could encompass a wide range of potential

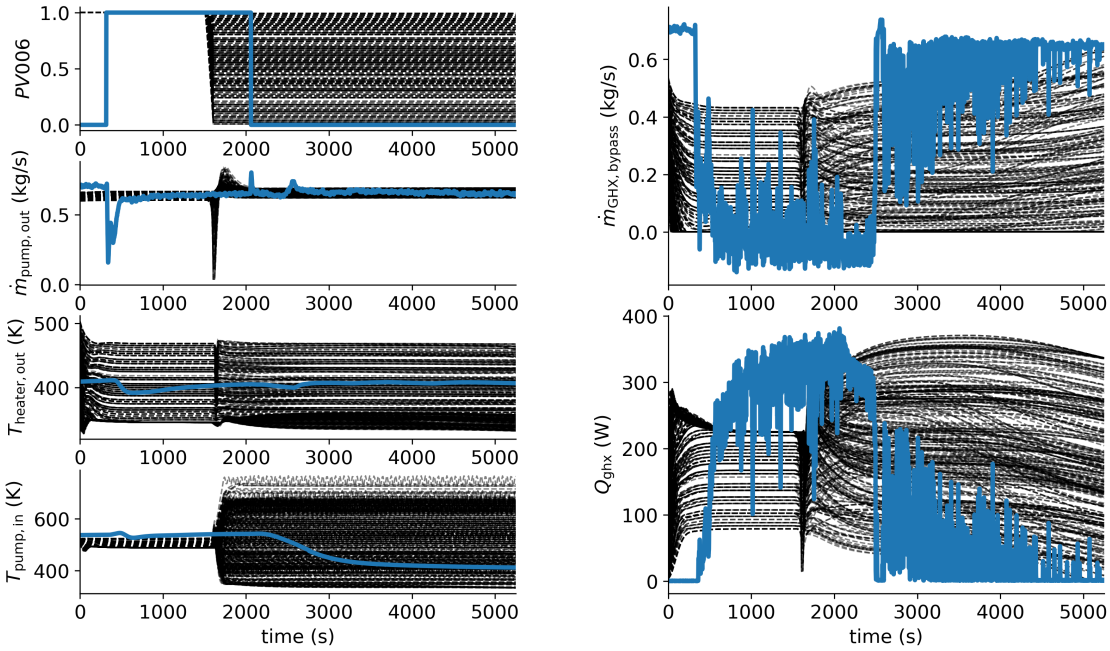

Figure 4: Experimental and Modelica actuators and states of GHX. The blue lines represent the experimental data and the black curves represent different Modelica simulations.

experimental scenarios, and a surrogate model (FNN, GRU, or SINDyC-based) is then trained to emulate the system dynamics in real-time, enabling rapid predictions and UQ.

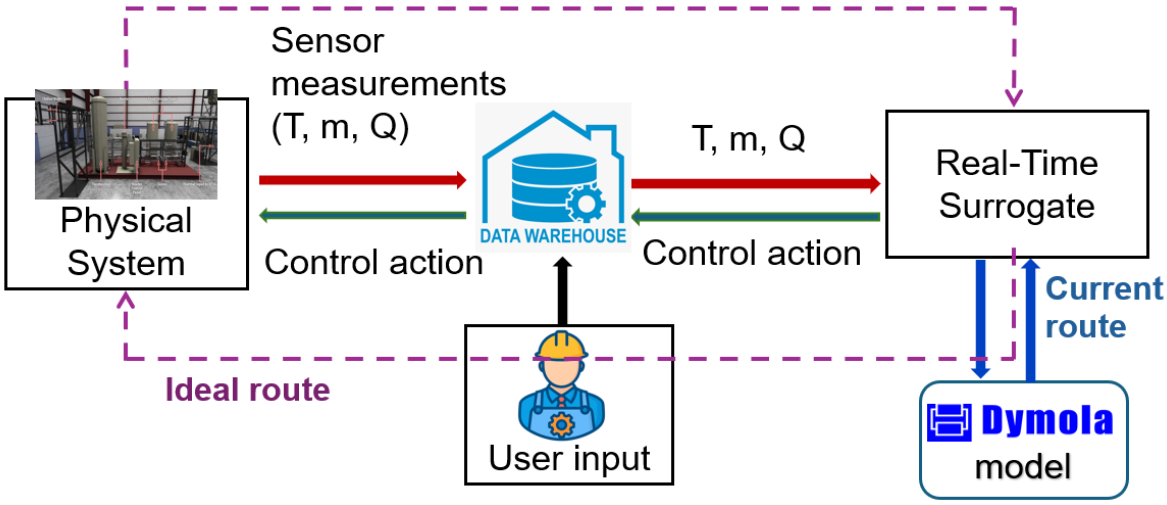

Figure 5: DT framework for autonomous control of TEDS. The ideal route (purple) corresponds to direct training on experimental data, whereas the current route (blue) uses simulation data for surrogate training.

4.2. Sparse Identification of Nonlinear Dynamics with Control (SINDyC)

The SINDyC framework discovers governing equations from time-series measurements in the presence of control inputs [18, 20]. In brief, the time derivative of the state vector $\mathbf{x}$ is expressed as a sparse linear combination of candidate functions of states and inputs:

$$\frac{d\mathbf{x}}{dt} \approx \Xi\,\Theta(\mathbf{x},\mathbf{u}), \tag{1}$$

where $\Theta(\mathbf{x}, \mathbf{u})$ is a library of basis functions and Ξ is a coefficient matrix identified via sparsity-promoting regression (e.g., sequential thresholding or LASSO). This yields a parsimonious model suitable for forecasting and control.

In this work, we restricted the candidate library to linear terms in states and controls, thus producing linear state-space surrogates that are easy to interpret and integrate into the control system [3]. For TEDS—when using the states and actuators shown in Fig. 4—the surrogates read:

$$\dot{\mathbf{x}}_{\mathrm{TES}} = \mathbf{A}_{\mathrm{TES}}\mathbf{x}_{\mathrm{TES}} + \mathbf{B}_{\mathrm{TES}}\mathbf{u}_{\mathrm{TES}} + \mathbf{d}_{\mathrm{TES}}, \tag{2}$$

$$\dot{\mathbf{x}}_{\mathrm{GHX}} = \mathbf{A}_{\mathrm{GHX}}\mathbf{x}_{\mathrm{GHX}} + \mathbf{B}_{\mathrm{GHX}}\mathbf{u}_{\mathrm{GHX}} + \mathbf{d}_{\mathrm{GHX}}, \tag{3}$$

with

$$\mathbf{x}_{\mathrm{TES}} = \left[\dot{m}_{\mathrm{TES,in}}, T_{\mathrm{TES,out}}, T_{\mathrm{top}}, T_{\mathrm{mid}}, T_{\mathrm{bot}}\right]^{\top}, \quad \mathbf{x}_{\mathrm{GHX}} = \left[\dot{m}_{\mathrm{GHX,bypass}}, Q_{\mathrm{GHX}}\right]^{\top},$$

$$\mathbf{u}_{\mathrm{TES}} = \mathbf{u}_{\mathrm{GHX}} = \left[PV_{006}, \dot{m}_{\mathrm{pump,out}}, T_{\mathrm{pump,in}}, T_{\mathrm{heater,out}}\right]^{\top}.$$

The matrices $\mathbf{A}$ and $\mathbf{B}$ are identified directly from the data. $\mathbf{d}_{\mathrm{TES}} \in \mathbb{R}^5$ and $\mathbf{d}_{\mathrm{GHX}} \in \mathbb{R}^2$ are learned constant offset (bias) vectors—the intercepts of each state equation arising from the constant (all-ones) column in the SINDyC library. They capture steady-state drift/unmodeled effects and sensor biases.

A single deterministic fit can be sensitive to noise and limited trajectories. To stabilize the surrogate and obtain calibrated uncertainty, we built a pool of linear SINDyC models in the following manner. From the full library of simulated trajectories (size S), we uniformly sampled 500 distinct subsets of size $p = 4$ where p = the number of trajectories used per SINDyC fit. Each 4-group was used to fit one linear SINDyC, yielding 500 coefficient vectors. We then fit a multivariate Gaussian to these coefficient vectors, defining the MvG-SINDyC. This preserved interpretability (explicit linear state-space form) while providing the uncertainty needed for robust DT/MPC.

In an associated AL preliminary study (not covered here, for the sake of brevity), we repeated the same procedure with larger subset sizes $p \in \{8, 12, 16\}$. For a fixed computational budget, $p = 4$ delivered the most favorable error-coverage trade-off (a lower RMSE with well-calibrated 95% coverage); therefore, we adopted $p = 4$ throughout this paper. *The authors are willing to expand further on this analysis during the review process if the reviewers prefer to do so (this sentence will be removed later).*

4.3. Machine Learning Surrogates: Feedforward Neural Network and Gated Recurrent Unit

To complement the physics-informed SINDyC surrogates, two data-driven ML models were developed: an FNN and a GRU network. These models served as flexible approximates capable of capturing nonlinear relationships and temporal dependencies directly from the data, without requiring an explicit governing equation form.

The FNN represents the simplest form of a deep learning surrogate, and consists of sequential layers of neurons that transform input vectors through nonlinear activation functions [23]. Each hidden layer performs an affine transformation followed by a nonlinear mapping, enabling the network to approximate complex input-output relationships. For this study, the FNN was trained to predict the future GHX states—mass flow rate and heat transfer rate—given actuator signals as inputs. Its simplicity and differentiability made it computationally efficient and suitable for mapping TEDS behavior.

Two variants of the FNN were trained to assess the impact of experimental data on model generalization. The first, denoted as FNN (with EXP), incorporated both simulation and experimental data during training, providing the network with direct exposure to real system dynamics and noise characteristics. The second, FNN (w/o EXP), was trained exclusively on simulation-generated trajectories, representing a purely synthetic surrogate. This distinction, as will be shown later, helps us assess whether or not FNN has overfitting tendencies.

The GRU extends this capability to temporal sequences, enabling the model to learn dynamic behaviors and dependencies over time. Unlike traditional recurrent neural networks, GRUs introduce gating mechanisms that regulate information flow between successive time steps, allowing the network to retain long-term dependencies while also mitigating vanishing or exploding gradient issues [24]. The GRU's hidden state acts as a compact memory of past observations, updated through reset and update gates that determine how much past information should be retained and/or overwritten. This structure enables GRUs to model the transient responses of the GHX and TES systems more effectively than can be done by purely feedforward architectures. Note that since AL requires repeated model retraining, the authors employed a GRU model instead of the more complex long short-term memory network. Although long short-term memory (LSTM) is a widely used recurrent architecture, the GRU proved sufficient in this study, affording a favorable balance between computational efficiency and predictive accuracy.

4.4. Active Learning for Trajectory Sampling

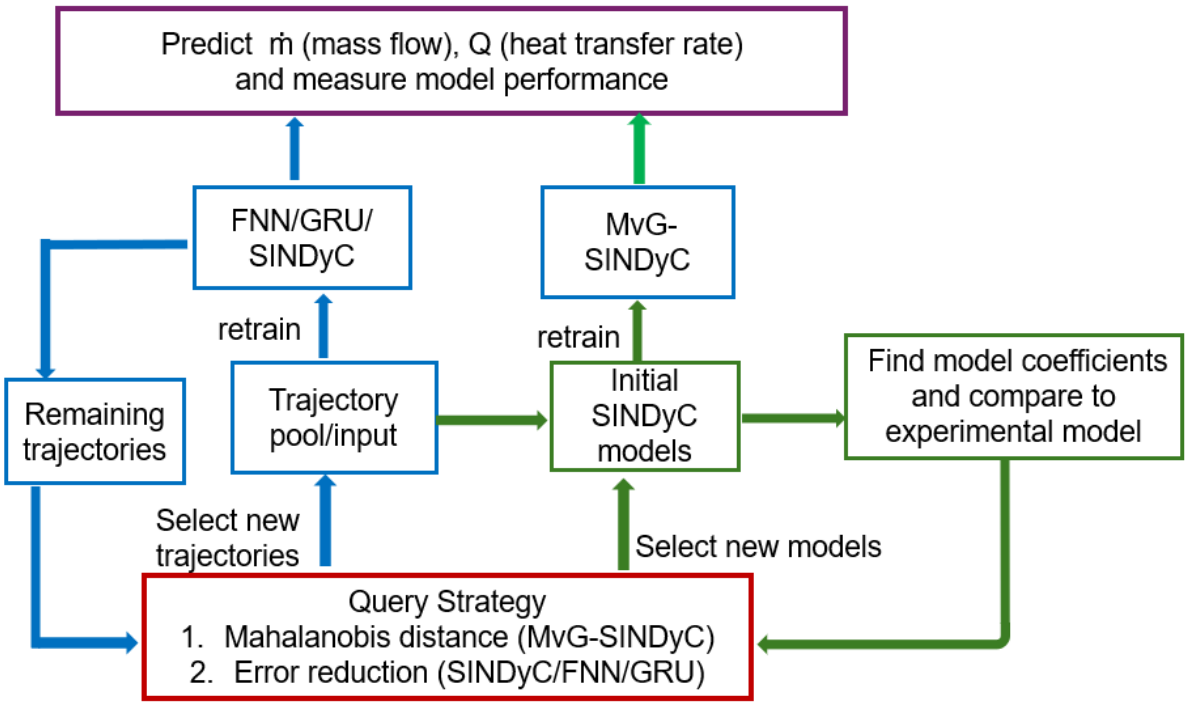

Figure 6: AL workflow for physics-informed (SINDyC/MvG-SINDyC) and data-driven (FNN/GRU) surrogates.

The proposed methodology integrates AL into the development of surrogate models for the TEDS DT. Fig. 6 illustrates the overall workflow, which consists of two complementary branches: a physics-informed surrogate (MvG-SINDyC) and purely data-driven neural networks (FNN/GRU). Both branches rely on a unique AL query strategy that adaptively selects the most informative training trajectories.

The right-hand branch of the framework employs MvG-SINDyC to approximate the governing dynamics of TES-GHX states. Initially, a subset of trajectories is drawn from the simulation pool to fit the SINDyC models. Their learned coefficient vectors are aggregated to form an initial multivariate Gaussian distribution, which serves as the probabilistic surrogate. Model coefficients are then compared against the experimental trajectory so as to quantify the discrepancy. AL iteratively identifies new models that minimize the Mahalanobis distance between the simulation model and experiment model in the coefficient space. The loop proceeds until the surrogate achieves

stable predictive performance with calibrated uncertainty. Figure 7 shows comparison among SINDyC model coefficient vectors from candidate simulation–trained models (thin curves) against the experimentally derived coefficient vector (bold red). The horizontal axis enumerates SINDyC basis terms; the vertical axis is the learned coefficient value. To rank candidates, we measure how close a model's coefficient vector $\mathbf{a} \in \mathbb{R}^P$ is to the experimental vector $\mathbf{a}_{\text{exp}} \in \mathbb{R}^P$ using the Mahalanobis distance

$$d_M(\mathbf{a}, \mathbf{a}_{\text{exp}}) = \sqrt{(\mathbf{a} - \mathbf{a}_{\text{exp}})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{a} - \mathbf{a}_{\text{exp}})}, \tag{4}$$

where $\boldsymbol{\Sigma} \in \mathbb{R}^{P \times P}$ is the empirical covariance matrix of the model coefficients, estimated from the current pool of candidates or selected models. This metric scales each direction by its variance and captures correlations among coefficients, so parameters that vary widely across models contribute less to the distance, while more stable, informative directions are emphasized. In the figure, the clustering of thin curves around the red line indicates that active selection based on (4) drives the surrogate toward experimental dynamics with fewer training iterations.

The left-hand branch leverages black-box ML models (FNN and GRU) and deterministic SINDyC. A small set of initial trajectories is used to train the models, which are then applied to predict the remaining trajectories. Model errors are computed relative to the experimental trajectory, and the AL query strategy selects the next most informative samples, based on prediction error. Newly added trajectories are incorporated into the training set, and the models are retrained. This iterative cycle improves predictive accuracy while also reducing reliance on large, randomly sampled datasets. At the core of the framework lies the query strategy, which unifies the two branches under a shared principle of information-driven sampling. Two types of query strategies are used:

- For the MvG-SINDyC branch, the criterion is the Mahalanobis distance in coefficient space. The models that yield the least deviation from experimental models are chosen.
- For the SINDyC/FNN/GRU branch, the criterion is prediction error reduction. The trajectories that produce the highest discrepancies are prioritized.

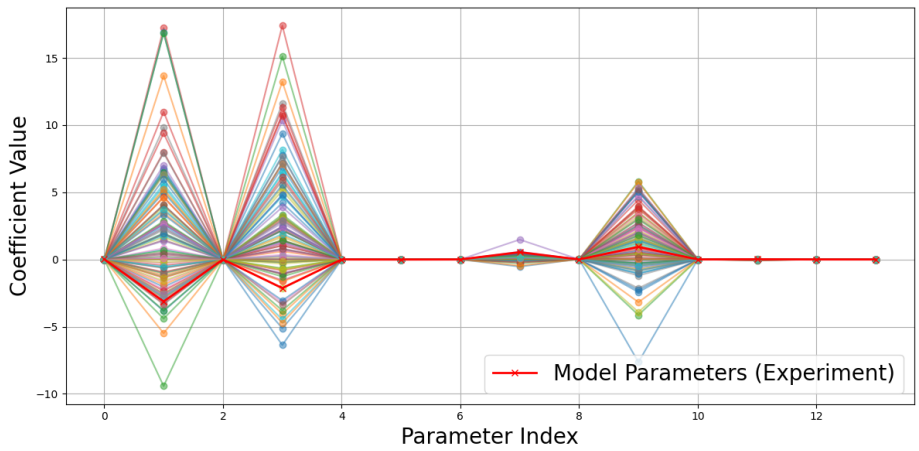

Figure 7: SINDyC model coefficient comparison. Thin lines: candidate models trained on individual set of four simulation trajectories; bold red line: model coefficients based on experimental data.

For MvG-SINDyC, we also tested an error-based sampling strategy analogous to the approach used for the neural surrogates. However, this method did not yield measurable improvements in terms of convergence or robustness. Because the MvG-SINDyC framework represents model variability through the mean and covariance of the identified coefficient distributions, its informativeness is governed by structural diversity in parameter space,

Table 1: Model architectures and training hyperparameters used in this study.

Model type	Architecture	Hyperparameters
SINDyC/ MvG-SINDyC	Linear candidate library (Polynomial Library, degree = 1); optimizer: Sequential Thresholded Least Squares.	$\lambda_{\text{TES}} = 10^{-6}$, $\alpha_{\text{TES}} = 10^{-3}$; $\lambda_{\text{GHX}} = 10^{-8}$, $\alpha_{\text{GHX}} = 10^{-6}$; integration via LSODA (rtol = 10^{-12}, atol = 10^{-12}).
FNN	Feedforward network with four layers: $4 \rightarrow 128 \rightarrow 128 \rightarrow 2$ nodes with ReLU between linear layers.	Optimizer: Adam; learning rate = 10^{-3}; loss: MSE; epochs = 40.
GRU	GRU network: 6 (features: past states + controls) $\rightarrow$ 128 $\rightarrow$ 128 $\rightarrow$ 2 nodes; lookback = 120; ReLU applied in output layer.	Optimizer: Adam; learning rate = 10^{-3}; loss: MSE; epochs = 20.

rather than pointwise prediction error. Consequently, error-based sampling provided no additional benefit, as the output residuals mainly reflected local noise rather than new dynamical information. In contrast, the Mahalanobis-distance criterion—defined in the coefficient space—proved more effective for identifying distinct dynamical regimes and was therefore adopted as the default query strategy for the MvG-SINDyC branch. By focusing on those trajectories that contribute maximal information, AL ensures that both interpretable and black-box surrogates converge faster, even with fewer simulation or experimental runs.

5. Results and Discussions

The architectures and training hyperparameters for each surrogate model—SINDyC, FNN, and GRU—are summarized in Table 1. The SINDyC models employed the Sequential Thresholded Least Squares algorithm—where the threshold parameter λ controlled sparsity by eliminating small coefficients, and the ridge coefficient α provided L2 regularization to stabilize the least-squares fit. The reported rtol and atol values denote the relative and absolute tolerances of the LSODA integrator used for the model rollout. In the neural surrogates, the FNN followed a fully connected architecture with ReLU activations, and was optimized using the Adam algorithm, with a learning rate of 10^{-3}. FNN (with EXP) and FNN (w/o EXP) shared the same architecture, and differed only in regard to whether experimental data were included during training. The GRU network processed temporal sequences by using a lookback window of 120 time steps, with each input vector being comprised of past state variables and control signals. This lookback length defined the temporal context available to the model. All the models were trained with mean-squared-error loss and normalized input-output features so as to ensure stable convergence and consistent comparison across methods.

Figs. 8 and 9 compare the surrogate model predictions against three unseen simulation trajectories and the experimental data, respectively. In the simulation case (Fig. 8), the GRU and FNN surrogates reproduce the reference trajectory with high fidelity, whereas SINDyC fails to capture the nonlinear transitions. This outcome highlights the trade-off between interpretability and accuracy: although SINDyC provides an analytical linear surrogate, it cannot resolve complex nonlinear behaviors inherent in the simulation data. Also, it is worth mentioning that GRU generally shows better predictive performance than FNN in Fig. 8, taking advantage of its temporal features.

In the experimental case (Fig. 9), the GRU and FNN trained with experimental supervision achieve the best alignment with the measured GHX mass flow and heat transfer, accurately capturing the temporal dynamics. The FNN trained without experimental guidance shows larger variance and deviation, particularly in the steady-state region, underscoring the importance of experimental calibration. By contrast, the deterministic SINDyC surrogate systematically underestimates the long-term dynamics owing to its linear coefficient formulation, which limits its

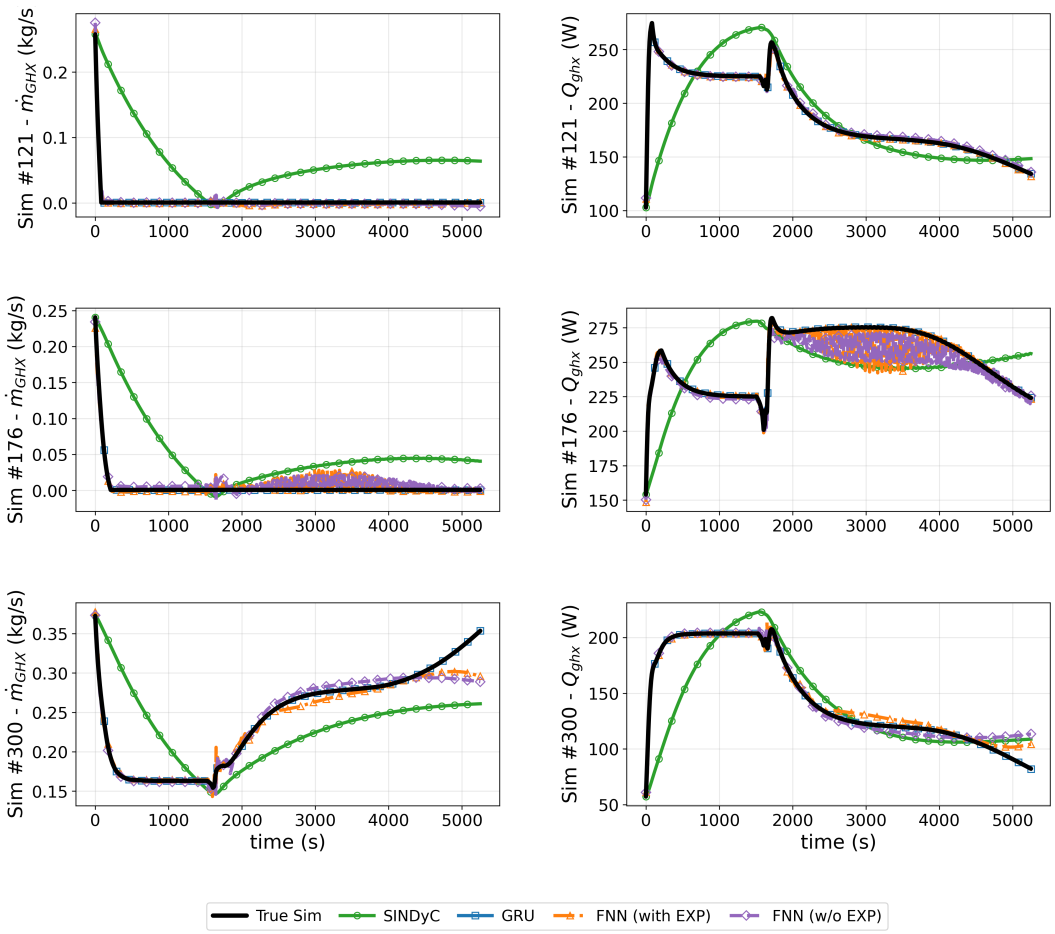

Figure 8: Comparison of the surrogate model predictions for three unseen simulation trajectories.

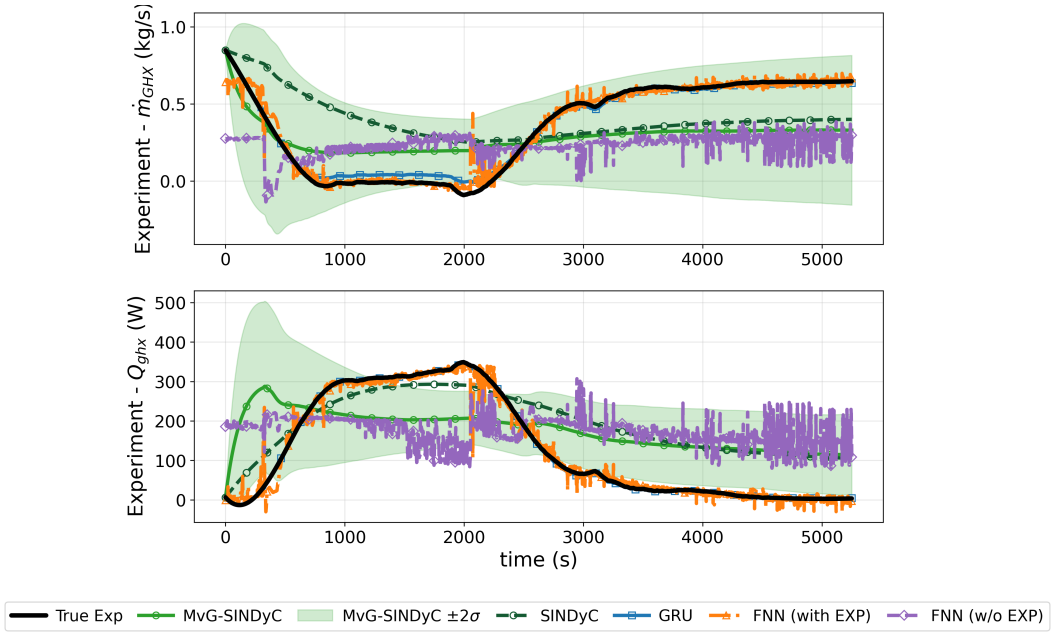

Figure 9: Comparison of surrogate model for experiment trajectory prediction.

Table 2: Prediction error (RMSE) for three representative unseen simulations and the experiment, for $\dot{m}_{\mathrm{GHX,bypass}}$ and Q_{GHX}.

Model	Simulation RMSE						Experiment RMSE	
	Sim #121		Sim #176		Sim #300		(Exp.)	
	$\dot{m}$ (kg/s)	Q (W)	$\dot{m}$ (kg/s)	Q (W)	$\dot{m}$ (kg/s)	Q (W)	$\dot{m}$ (kg/s)	Q (W)
SINDyC	0.052	25.055	0.064	31.344	0.061	28.001	0.291	98.673
MvG-SINDyC	—	—	—	—	—	—	0.257	121.471
GRU	0.000	0.118	0.000	0.311	0.000	0.094	0.033	2.630
FNN (with EXP)	0.003	1.595	0.009	10.153	0.012	5.474	0.147	65.005
FNN (w/o EXP)	0.003	2.159	0.008	8.730	0.015	7.511	0.288	133.847

capacity to represent strongly nonlinear transient–steady interactions. Note that in this case, the FNN model generally fails to generalize beyond the simulation data, exhibiting clear signs of overfitting and poor performance as compared to the experimental results. The proposed MvG-SINDyC framework extends the classical deterministic SINDyC approach by incorporating a probabilistic multivariate Gaussian representation, yielding both the mean trajectory and a time-varying uncertainty band. As shown in Fig. 9, the mean prediction closely follows the experimental response while the shaded 95% credible interval realistically bounds most of the experimental data. The predictive uncertainty of the MvG-SINDyC model was further quantified by evaluating the fraction of the experimental trajectory lying within its 95% predictive interval. This metric measures the uncertainty calibration of the probabilistic surrogate, i.e., the consistency between predicted credible regions and observed experimental variability. The model achieved coverage values of 0.89 for the GHX mass-flow rate and 0.96 for the heat-transfer rate.

Table 2 provides quantitative support to these observations, reporting the RMSEs for the three representative simulation trajectories (i.e., 121, 176, and 300) as well as the experimental trajectory. Across all simulations, the GRU is consistently the most accurate: the $\dot{m}$ errors are essentially zero and the Q errors remain small (about 0.094–0.311 W). The FNN is competitive but less precise. With experimental supervision, the FNN attains an $\dot{m}$ RMSE of 0.003–0.012 kg/s and a Q RMSE of 1.595–10.153 W. Without supervision, it yields 0.003–0.015 kg/s and 2.159–8.730 W. In contrast, the linear SINDyC surrogate is roughly an order of magnitude less accurate on these nonlinear trajectories (e.g., $\dot{m}$ RMSE 0.052–0.064 kg/s and Q RMSE 25–31 W), highlighting that linear surrogates such as SINDyC are less suited for capturing nonlinear simulation dynamics.

In the experimental case, performance differences are amplified. GRU remains the most accurate, with RMSE values of 0.033 kg/s and 2.630 W, demonstrating strong generalization from simulation-trained models to physical system data. The FNN with experimental supervision improves significantly over SINDyC, but still exhibits moderate errors (0.147 kg/s, 65.005 W). The FNN without experimental data, however, performs worst of all, with an RMSE exceeding that of SINDyC and MvG-SINDyC, reinforcing the necessity of experimental calibration. SINDyC models offer physically interpretable surrogates but underperforms in nonlinear regimes; GRU achieves the best accuracy and robustness, particularly in generalizing from simulations to experiments; and FNN occupies an intermediate position, with access to experimental data substantially improving performance.

Figure 10 compares the prediction error trends of the four surrogate models (SINDyC, MvG-SINDyC, FNN, GRU) under both AL and random sampling. For MvG-SINDyC (top row), AL consistently outperforms random sampling by producing a smoother convergence and lower RMSE for both quantities. The early-stage advantage (50–150 models) is particularly clear, with AL reducing the RMSE of $\dot{m}_{\mathrm{GHX}}$ from approximately 0.06 to 0.04 kg/s, and that of Q_{GHX} from about 30 to 25 W. Although the long-run difference narrows as the number of models

Figure 10: Prediction error trends (RMSE) pertaining to the mass-flow and power output for different models.

increases, AL achieves the same final accuracy with roughly half the data, thus confirming its benefit in terms of data efficiency. For the deterministic SINDyC (second row), both AL and random sampling exhibit nearly identical convergence behavior. Because this linear surrogate cannot capture the strongly nonlinear dynamics of the GHX subsystem, the overall RMSE remains flat and relatively high, similar to the FNN trained without experimental data. Although AL employed an error-based sampling strategy analogous to that used for the neural surrogates, its

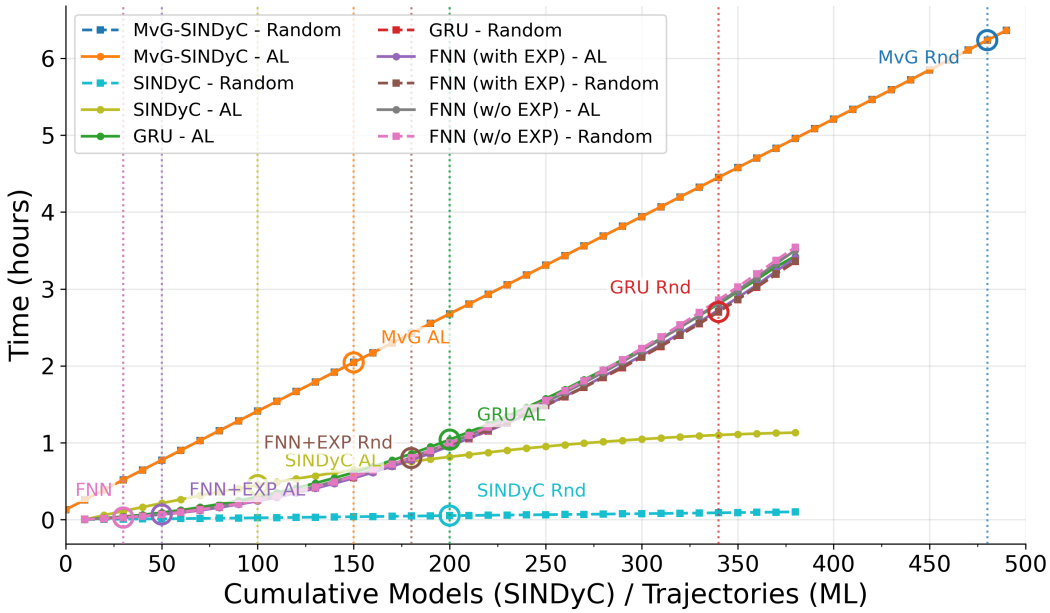

Figure 11: Computational time comparison of each model.

impact was minimal—reflecting the limited information gain. The RMSE of $\dot{m}_{\mathrm{GHX}}$ stabilizes around 0.30 kg/s, and that of Q_{GHX} around 100 W, with both metrics showing only marginal reduction under AL. These results highlight that, for purely deterministic and linearized formulations such as SINDyC, the benefits of AL are inherently limited by model capacity rather than data coverage.

For the FNN surrogates (third row), AL yields the most dramatic improvement. When experimental data are included during training, the RMSE for $\dot{m}_{\mathrm{GHX}}$ drops from ~ 0.01 to below 0.005 kg/s, and that for Q_{GHX} drops from ~ 10 to nearly 1 W after only ~ 200 trajectories, whereas random sampling remains higher and noisier. In contrast, the FNN without experimental conditioning stays at around 0.02 kg/s and 8–10 W, showing that exposure to real-system signals is crucial for stable black-box learning. As shown in Figure 10 (middle row), both models—FNN with EXP-AL, and FNN with EXP-Random—exhibit a sharp reduction in error after approximately 50 and 175 iterations, respectively. These points correspond to when the AL process requested that the experimental trajectory be sampled and incorporated into the training dataset, leading to significant improvement in model performance.

The GRU surrogates (bottom row) also benefit from AL, showing a consistently lower RMSE than for random sampling, across all stages. AL drives $\dot{m}_{\mathrm{GHX}}$ below 0.003 kg/s and Q_{GHX} below 1 W within the first 100–150 models, whereas random sampling converges more slowly. Notably, GRUs converge more quickly than SINDyC and with less noise than FNNs, reflecting their strength in capturing sequential dependencies in time-series data. This suggests that GRUs are particularly well-suited for dynamical systems (e.g., TEDS) when sufficient simulation trajectories are available.

When compared with the runtime analysis in Figure 11, these results reveal a clear trade-off between model interpretability and computational efficiency. The MvG-SINDyC models required the longest cumulative training time, reaching over 6 hours for 500 models under random sampling. The training time reduces to approximately 2 hours when guided by AL, representing a near- threefold reduction in total runtime. This cost stems from the repeated sparse regression and integration steps in each model fit, which scale linearly with the number of trajectories. Despite the higher cost, the interpretability of MvG-SINDyC makes this gain from AL particularly valuable, as fewer but more informative models can be used to achieve comparable accuracy. On the other hand, the neural surrogates (FNN and GRU) trained substantially faster, completing their cumulative training at some

point between 2 and 3 hours. Both architectures benefited from AL, though the relative speedup was smaller than MvG-SINDyC. For example, the GRU reduced its training time from roughly 3.2 hours (random) to 2.8 hours (AL), while the FNN converged within 2.5 hours under random sampling and under 2 hours with AL. This efficiency stems from GPU-parallelized batch optimization and simpler loss evaluation in comparison to the MvG-SINDyC. In contrast, the deterministic SINDyC fits are comparatively cheap computationally—well under 0.2 hours even near 200 models with random sampling, and about 1.1 hours along the AL path.

Interestingly, while the total wall-clock time of neural surrogates is lower, MvG-SINDyC's AL pipeline achieves a larger relative improvement per sample. AL therefore plays a dual role: reducing the number of training trajectories needed, and lowering the computational cost by avoiding uninformative samples. Overall, these comparisons demonstrate that AL yields measurable efficiency gains across all model classes, with the largest absolute savings being observed for computationally intensive interpretable models (SINDyC). In addition, consistent—though smaller—reductions were seen for data-driven neural surrogates.

Together, Figures 10 and 11 demonstrate that AL accelerates convergence in both error and runtime dimensions, enabling scalable surrogate development for IES DTs.

6. Conclusions

This study introduced an AL framework for constructing a physics-informed DT of TEDS, integrating interpretable and data-driven surrogate modeling approaches. Four surrogate architectures were evaluated—deterministic SINDyC, probabilistic MvG-SINDyC, FNN, and GRU—to explore trade-offs among interpretability, data efficiency, predictive accuracy, and computational cost. By combining high-fidelity Modelica simulations with experimental measurements, the proposed workflow enables automated model refinement, UQ, and cross-validation between physics-based and ML paradigms. The framework integrated high-fidelity Modelica simulations with experimental measurements, enabling systematic model updates and validation across both physics-based and machine-learning paradigms.

Among the evaluated surrogates, the GRU model achieved the lowest prediction errors and strongest temporal generalization, with simulation RMSE values of 0.094–0.311 W and an experimental RMSE of 2.63 W. The FNN trained on experimental data achieved comparable accuracy despite its lighter architecture, making it attractive for real-time deployment. In contrast, the deterministic SINDyC captured basic linearized dynamics with minimal computational overhead, making it the lightest and fastest to train. However, due to its limited representational capacity and reliance on error-based sampling, AL offered only marginal gains in predictive accuracy. The probabilistic MvG-SINDyC incorporated multivariate Gaussian inference to quantify uncertainty, producing credible intervals that captured experimental behavior and improving stability through sampling in coefficient space.

AL consistently outperformed random sampling across all surrogate families, improving both prediction accuracy and computational efficiency. Across the MvG-SINDyC, FNN, and GRU models, AL achieved up to an order-of-magnitude gain in data efficiency and nearly a threefold reduction in total training time. The runtime analysis revealed a complementary trade-off between interpretability and computational efficiency. MvG-SINDyC incurred the highest cumulative cost, requiring approximately 6 hours to train 500 models under random sampling. This was reduced to about 2 hours with AL—a threefold speedup. Neural surrogates were trained substantially faster: the GRU completed training in $\sim$3.2 hours (random) versus 2.8 hours (AL), while the FNN required $\sim$2.5 hours (random) and only 2 hours (AL). These results confirm that AL not only reduces the number of required trajectories but also shortens the total compute time by avoiding redundant or low-information samples. The largest relative gains were observed for the computationally intensive SINDyC model, while neural networks benefited from smaller-yet-consistent runtime improvements.

Overall, the results highlight the complementary strengths of the four paradigms: SINDyC and MvG-SINDyC for interpretability and UQ, FNN for computational simplicity, and GRU for temporal fidelity. The proposed AL-DT framework unifies these advantages, providing a scalable pathway toward adaptive, trustworthy, physics-grounded DTs for complex thermal-hydraulic systems and IES.

The next stage of this research will extend AL beyond model training to encompass the DT's entire supervisory control loop. This integration will enable autonomous supervisory control capable of balancing exploration (data acquisition) and exploitation (control performance) while also ensuring physical consistency and operational safety. By fusing physics-based interpretability with AL, the extended AL-DT framework represents a step toward self-improving, trustworthy AI for resilient and decarbonized energy infrastructures.

Data Availability

Currently, the authors have, in a private GitHub repository, all the datasets and codes necessary to reproduce all the results of this work. To ensure confidentiality of this research, the authors will make this repository public during an advanced stage of the review process, and it will be listed under their research group's public Github page: https://github.com/aims-umich.

Acknowledgment

This work was supported through Idaho National Laboratory (INL) Laboratory Directed Research and Development (LDRD) Program Award 24A1081-116FP, under U.S. Department of Energy Idaho Operations Office Contract DE-AC07-05ID14517. The authors acknowledge the use of INL's high-performance computing resources, which significantly contributed to the modeling and analysis efforts presented in this work.

CRediT Author Statement

- **Umme Mahbuba Nabila**: Conceptualization, Methodology, Software, Validation, Formal Analysis, Visualization, Investigation, Data Curation, Writing - Original Draft.
- **Paul Seurin**: Conceptualization, Methodology, Resources, Funding Acquisition, Supervision, Project Administration, Writing - Original Draft.
- **Linyu Lin**: Conceptualization, Methodology, Resources, Funding Acquisition, Supervision, Project Administration, Writing - Original Draft.
- **Majdi I. Radaideh**: Conceptualization, Methodology, Resources, Funding Acquisition, Supervision, Project Administration, Writing - Original Draft.

References

[1] K. L. Frick, S. M. Bragg-Sitton, C. Rabiti, Development of the inl thermal energy distribution system (teds) in the modelica eco-system for validation and verification, Tech. rep., Idaho National Laboratory (INL), Idaho Falls, ID (United States) (2020).

[2] S. M. Bragg-Sitton, R. Boardman, C. Rabiti, J. O'Brien, Reimagining future energy systems: Overview of the us program to maximize energy utilization via integrated nuclear-renewable energy systems, International Journal of Energy Research 44 (10) (2020) 8156–8169.

[3] P. Seurin, L. Lin, Control under uncertainty for a physics-informed model of a thermal energy distribution system: Qualitative analysis, Available at SSRN 5667550.

[4] A. I. Arvanitidis, V. Agarwal, M. Alamaniotis, Nuclear-driven integrated energy systems: A state-of-the-art review, Energies 16 (11) (2023) 4293.

[5] R. S. El-Emam, A. Constantin, R. Bhattacharyya, H. Ishaq, M. E. Ricotti, Nuclear and renewables in multi-purpose integrated energy systems: A critical review, Renewable and Sustainable Energy Reviews 192 (2024) 114157.

[6] G. C. Masotti, A. Cammi, S. Lorenzi, M. E. Ricotti, Modeling and simulation of nuclear hybrid energy systems architectures, Energy Conversion and Management 298 (2023) 117684.

[7] S. L. Luxembourg, S. S. Salim, K. Smekens, F. D. Longa, B. van der Zwaan, Times-europe: An integrated energy system model for analyzing europe's energy and climate challenges, Environmental Modeling & Assessment 30 (1) (2025) 1–19.

[8] D. Mikkelson, K. Frick, Analysis of controls for integrated energy storage system in energy arbitrage configuration with concrete thermal energy storage, Applied Energy 313 (2022) 118800.

[9] R. A. Jacob, J. Zhang, Modeling and control of nuclear–renewable integrated energy systems: Dynamic system model for green electricity and hydrogen production, Journal of Renewable and Sustainable Energy 15 (4) (2023) 046302.

[10] S. Hills, S. Dana, H. Wang, Dynamic modeling and simulation of nuclear hybrid energy systems using freeze desalination and reverse osmosis for clean water production, Energy Conversion and Management 247 (2021) 114724.

[11] L. Williams, J. M. Doster, D. Mikkelson, Modeling and optimization of a nuclear integrated energy system for the remote microgrid on el hierro, Energies 17 (23) (2024) 5826.

[12] S. Gautam, A. Szczublewski, A. Fox, S. Mahmud, A. Javaid, T. O. Olowu, T. Westover, R. Khanna, Digital real-time simulation and power quality analysis of a hydrogen-generating nuclear-renewable integrated energy system, Energies 18 (4) (2025) 937.

[13] L. Lin, Development of supervisory control system for thermal energy distribution system, in: 2024 Pacific Basin Nuclear Conference, PBNC 2024, American Nuclear Society, 2024, pp. 435–444.

[14] L. Lin, J. Oncken, V. Agarwal, Autonomous control for heat-pipe microreactor using data-driven model predictive control, Annals of Nuclear Energy 200 (2024) 110399.

[15] B. Kouvaritakis, M. Cannon, Model predictive control, Switzerland: Springer International Publishing 38 (13-56) (2016) 7.

[16] P. Seurin, L. Lin, Uncertainty quantification of a physics-informed model based on sparse identification of a thermal energy distribution system, Annals of Nuclear Energy 226 (2026) 111865.

[17] K. L. Frick, S. M. Bragg-Sitton, M. Garrouste, Validation and verification for inl modelica-based teds models via experimental results, Tech. rep., Idaho National Lab.(INL), Idaho Falls, ID (United States) (2021).

[18] S. L. Brunton, J. L. Proctor, J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, Proceedings of the national academy of sciences 113 (15) (2016) 3932–3937.

[19] E. Kaiser, J. N. Kutz, S. L. Brunton, Sparse identification of nonlinear dynamics for model predictive control in the low-data limit, Proceedings of the Royal Society A 474 (2219) (2018) 20180335.

[20] A. A. Kaptanoglu, B. M. de Silva, U. Fasel, K. Kaheman, A. J. Goldschmidt, J. L. Callaham, C. B. Delahunt, Z. G. Nicolaou, K. Champion, J.-C. Loiseau, et al., Pysindy: A comprehensive python package for robust sparse system identification, arXiv preprint arXiv:2111.08481.

[21] R. De Maesschalck, D. Jouan-Rimbaud, D. L. Massart, The mahalanobis distance, Chemometrics and intelligent laboratory systems 50 (1) (2000) 1–18.

[22] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, nature 521 (7553) (2015) 436–444.

[23] I. Goodfellow, Y. Bengio, A. Courville, Y. Bengio, Deep learning, Vol. 1, MIT press Cambridge, 2016.

[24] J. Chung, C. Gulcehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, arXiv preprint arXiv:1412.3555.

[25] 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.

[26] 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.

[27] M. I. Radaideh, T. Kozlowski, Surrogate modeling of advanced computer simulations using deep gaussian processes, Reliability Engineering & System Safety 195 (2020) 106731.

[28] S. Sene, L. Lin, J. Kim, M. I. Radaideh, Surrogate-driven variance-based sensitivity analysis of thermal storage tanks in integrated energy systems, *In: Nuclear Plant Instrumentation and Control & Human-Machine Interface Technology (NPIC&HMIT 2025)*, Chicago, Illinois, United States, June 15–18, 2025.

[29] D. Price, M. I. Radaideh, B. Kochunas, Simplified matching pursuits applied to 3d nuclear reactor temperature distribution construction, Applied Mathematical Modelling 131 (2024) 134–158.

[30] M. G. Prina, M. Dallapiccola, D. Moser, W. Sparber, Machine learning as a surrogate model for energyplan: Speeding up energy system optimization at the country level, Energy 307 (2024) 132735.

[31] F. Lédée, C. Crawford, R. Evins, Improved surrogate modeling for multi-energy system design: Model architecture, sampling and scaling choices, Applied Energy 390 (2025) 125812.

[32] Y. Li, M. Han, M. Shahidehpour, J. Li, C. Long, Data-driven distributionally robust scheduling of community integrated energy systems with uncertain renewable generations considering integrated demand response, Applied Energy 335 (2023) 120749.

[33] Y. Zhou, H. Hou, H. Yan, X. Wang, R. Zhou, Data-driven distributionally robust stochastic optimal dispatching method of integrated energy system considering multiple uncertainties, Energy 325 (2025) 136104.

[34] 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.

[35] 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.

[36] 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.

[37] B. Zhang, W. Hu, X. Xu, Z. Zhang, Z. Chen, Hybrid data-driven method for low-carbon economic energy management strategy in electricity-gas coupled energy systems based on transformer network and deep reinforcement learning, Energy 273 (2023) 127183.

[38] M. Mo, X. Xiong, Y. Wu, Z. Yu, Deep-reinforcement-learning-based low-carbon economic dispatch for community-integrated energy system under multiple uncertainties, Energies 16 (22) (2023) 7669.

[39] 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 (2025) 101090.

[40] 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.

[41] T. Liang, X. Zhang, J. Tan, Y. Jing, L. Liangnian, Deep reinforcement learning-based optimal scheduling of integrated energy systems for electricity, heat, and hydrogen storage, Electric Power Systems Research 233 (2024) 110480.

[42] A. Khodadadi, S. Adinehpour, R. Sepehrzad, A. Al-Durra, A. Anvari-Moghaddam, Data-driven hierarchical energy management in multi-integrated energy systems considering integrated demand response programs and energy storage system participation based on madrl approach, Sustainable Cities and Society 103 (2024) 105264.

[43] J. Liu, X. Meng, J. Wu, Data-driven optimal scheduling for integrated electricity-heat-gas-hydrogen energy system considering demand-side management: A deep reinforcement learning approach, International Journal of Hydrogen Energy 103 (2025) 147–165.

[44] K. Hua, Q. Xu, S. Li, L. Fang, Data-driven optimal dispatch of integrated energy in zero-carbon building system considering occupant comfort and uncertainty, Applied Sciences 15 (8) (2025) 4408.

[45] L. A. Burnett, U. M. Nabila, M. I. Radaideh, Variational digital twins, arXiv preprint arXiv:2507.01047.

[46] 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.

[47] M. Radaideh, C. Pappas, P. Ramuhalli, S. Cousineau, Application of convolutional and feedforward neural networks for fault detection in particle accelerator power systems, in: Annual Conference of the PHM Society, Vol. 14, 2022.

[48] M. I. Radaideh, C. Pappas, J. Walden, D. Lu, L. Vidyaratne, T. Britton, K. Rajput, M. Schram, S. Cousineau, Time series anomaly detection in power electronics signals with recurrent and convlstm autoencoders, Digital Signal Processing 130 (2022) 103704.

[49] M. I. Radaideh, C. Pappas, M. Wezensky, P. Ramuhalli, S. Cousineau, Early fault detection in particle accelerator power electronics using ensemble learning, International Journal of Prognostics and Health Management 14 (1).

[50] Y. Alanazi, M. Schram, K. Rajput, S. Goldenberg, L. Vidyaratne, C. Pappas, M. I. Radaideh, D. Lu, P. Ramuhalli, S. Cousineau, Multi-module-based cvae to predict hvcm faults in the sns accelerator, Machine Learning with Applications 13 (2023) 100484.

[51] T. J. Morton, Thermal energy distribution system (teds) startup, Tech. rep., Idaho National Laboratory (INL), Idaho Falls, ID (United States) (2020).