buildings

Article

Physics-Guided Neural ODEs for Building Thermal Prediction and Model Predictive Control

Yagang Wang [1], Lexue Chang [2], Enzhan Zhang [2,*], Kejun Jia [2], Xia Zhang [2] and Yonghao Li [1]

[1] Gansu Natural Energy Research Institute, Lanzhou 730046, China
[2] College of Computer and Artificial Intelligence, Lanzhou University of Technology, Lanzhou 730050, China
* Correspondence: zhangenzhan@lut.edu.cn

Abstract

Building supervisory control requires accurate, prior-consistent, and computationally tractable thermal models. This study proposes a physics-guided neural ordinary differential equation (PG-NODE) framework for thermal prediction and model predictive control (MPC). It combines a resistance–capacitance (RC)-inspired reference, a neural residual, RC-prior directional regularization, and validation-based checkpoint selection; the selected predictor remains fixed during MPC operation. Using 30 min EnergyPlus data from five zones, the framework was evaluated through held-out prediction, Gaussian noise and control-input-mismatch tests, ablation, and surrogate-based closed-loop experiments. Long short-term memory achieved the lowest prediction root mean square error (RMSE), whereas PG-NODE achieved the lowest RMSE among the evaluated neural ODE models and the lowest directional inconsistency rate among learned nonlinear models. In five-seed $\times$ five-window BACK SPACE experiments using independently trained, fixed PG-NODE surrogate environments rather than direct EnergyPlus interaction, fixed-block analysis supported lower reference-tracking RMSE for PG-NODE-MPC than for rule-based control and lower tariff-weighted normalized control effort than for the extended Kalman filter-based RC-MPC benchmark. This benchmark achieved the lowest descriptive mean reference-tracking RMSE and total objective. Mean PG-NODE-MPC optimization time was 0.86 s. Results suggest the potential for low-frequency building management system supervisory decision support, subject to physical command mapping and staged field validation.

Keywords: physics-guided neural networks; neural ordinary differential equations; model predictive control; building thermal dynamics; RC model; directional consistency; robustness analysis

Academic Editor: Apple L.S. Chan

Received: 29 June 2026
Revised: 27 July 2026
Accepted: 28 July 2026
Published: 29 July 2026

1. Introduction

Building operations represent a substantial component of life-cycle energy use, and improving heating, ventilation, and air-conditioning (HVAC) operation is therefore an important pathway for reducing operational demand [1,2]. Model predictive control (MPC) has been extensively studied for building applications because it uses predictive models and forecasts to optimize future control actions over a receding horizon [3,4]. Its practical performance, however, depends strongly on the predictive accuracy, robustness, and control-oriented reliability of the underlying building thermal model [5,6].

Control-oriented building thermal models span low-order resistance–capacitance (RC) grey-box formulations and fully data-driven predictors [7,8]. Recent physics-informed learning studies have sought to combine structured thermal knowledge with neural-network flexibility through architecture design, hybrid residual formulations, physical constraints, and physics-guided loss functions [9–17]. Recent studies have also shown that the selection and strength of physical regularization can materially affect predictive performance and consistency, motivating separate evaluation of these objectives rather than assuming that stronger incorporation of physical knowledge necessarily leads to better prediction accuracy [18–20].

Neural ordinary differential equations (neural ODEs) provide a continuous-time framework for learning dynamical systems [21]. In building applications, neural ODEs have been embedded in MPC for energy flexibility control [22,23], extended to physics-informed multi-zone thermal modeling using measured residential data [24], and combined with continuous-time physics-informed thermal models for demand-response MPC [25]. In parallel, recent real-life data-driven MPC studies [26–29] and uncertainty-aware thermal-modeling research [30] have emphasized the importance of measured-data validation and deployment-oriented assessment beyond simulation-only benchmarks [31].

Two lines of work are particularly close to the present study. Gokhale et al. [12] developed physics-informed neural models for control-oriented building thermal prediction, demonstrating the value of embedding thermal knowledge into learned dynamics; however, their formulation does not use the RC-reference–neural-residual decomposition considered here. He et al. [22] have demonstrated neural ODE-based MPC for building energy management, establishing the feasibility of continuous-time learned dynamics within predictive control. The contribution of the present work therefore does not lie in introducing neural ODEs into MPC itself. Rather, it lies in combining an identified RC-inspired reference component with a neural ODE residual, activity-thresholded directional regularization relative to the RC prior, and validation-based control-oriented checkpoint screening. Recent studies on physical-regularization sensitivity and explicit physical-consistency assessment further support the need to distinguish predictive accuracy from prior or physical consistency [18–20].

Building on these elements, this study develops a physics-guided neural ODE (PG-NODE) model for building thermal prediction and embeds the validation-selected predictor in a PG-NODE-based MPC controller (PG-NODE-MPC). The RC-inspired component provides the reference thermal trend used by the directional regularization, while the neural residual captures dynamics not represented by the low-order prior. For the main closed-loop controller, the selected PG-NODE remains fixed after offline training; the latest observed indoor temperature is used only for state-feedback reinitialization. Restricted final-layer parameter adaptation is investigated separately through ablation variants and is not part of the main PG-NODE-MPC controller. For closed-loop benchmarking, PG-NODE-MPC is compared with a rule-based controller (RBC) and an adaptive RC-based MPC controller using extended Kalman filter (EKF) parameter estimation, denoted RC-EKF-MPC.

The framework is evaluated through repeated prediction and robustness experiments across five thermal zones, controlled HVAC-input mismatch tests, directional regularization sensitivity analysis, ablation experiments, and repeated multi-window closed-loop controller comparisons. The multi-zone experiments characterize predictive behavior across zones within the same prototype building, whereas the main closed-loop evaluation is restricted to the BACK SPACE zone. Accordingly, the closed-loop results are interpreted as evidence specific to the evaluated simulation case rather than as general evidence of year-round, cross-building, or cross-climate superiority. Table 1 positions the proposed framework relative to representative related studies.

Table 1. Comparison of the present framework with representative related building thermal modeling and control studies.

Study	Core Approach	Evaluation Scope	Key Distinction Relative to the Present Study
Drgoňa et al. [11]	Physics-constrained structured recurrent thermal model	Multi-zone thermal modeling using real-building data	Uses a discrete-time structured neural formulation rather than an RC-reference–neural-ODE-residual architecture with RC-prior directional regularization
Gokhale et al. [12]	Physics-informed neural thermal model	Control-oriented thermal prediction using simulated and real data	Focuses on physics-informed thermal prediction rather than an RC-reference–neural-residual neural ODE predictor integrated into the receding-horizon MPC framework considered here
Xiao and You [13]	Physically consistent deep-learning thermal model with MPC	Building thermal modeling and MPC-based energy optimization	Enforces physical consistency through model structure, whereas the present study uses a soft RC-prior directional regularization together with validation-based deployment checkpoint selection
Jiang et al. [19]	Physics-informed modularized neural network for building control	Physical-consistency assessment and deep-reinforcement-learning-based control	Uses a modularized neural representation and reinforcement-learning-based control rather than a continuous-time RC-reference–neural-residual model embedded in receding-horizon MPC
He et al. [22]	Neural ODE thermal model with MPC	Building energy-flexibility control	Demonstrates neural ODE-based MPC, whereas the present study additionally introduces RC-reference–neural-residual decomposition, RC-prior directional regularization, and validation-based deployment checkpoint screening
Sabbagh et al. [24]	Physics-informed neural ODE with structured multi-zone dynamics	Multi-zone thermal modeling using measured residential data	Focuses on structured and interpretable neural ODE thermal modeling rather than integration with the RC-prior regularization and closed-loop MPC evaluation framework considered here
Wang et al. [25]	Continuous-time physics-informed thermal model with MPC	BOPTEST-based control and sensitivity analysis	Embeds explicit physics-informed continuous-time thermal modeling in MPC, whereas the present study adopts an identified RC reference plus learnable neural ODE residual and validation-based deployment screening
Present study	RC-inspired reference with neural ODE residual and RC-prior directional regularization	Five-zone prediction and robustness evaluation; single-zone repeated multi-window surrogate-based MPC evaluation	Integrates RC-reference–neural-residual modeling, activity-thresholded directional regularization, validation-based checkpoint screening, and matched closed-loop benchmarking against RBC and adaptive RC-EKF-MPC

The main contributions of this work are summarized as follows:

(1) An RC-reference–neural-residual neural ODE architecture is developed for control-oriented building thermal prediction, combining a low-order RC-inspired reference component with a learnable neural residual for nonlinear and unmodeled thermal dynamics.

(2) An activity-thresholded RC-prior directional regularization term is introduced to penalize disagreement between the learned local temperature change trend and the selected RC-inspired prior while excluding near-equilibrium samples with weak directional information.

(3) A validation-based deployment-model selection procedure is introduced to screen candidate PG-NODE checkpoints using local HVAC-response behavior and long-horizon rollout performance. The resulting validation-selected predictor is subsequently kept fixed in the main PG-NODE-MPC controller, with measured indoor temperature used only for state-feedback reinitialization.

(4) The framework is evaluated through five-zone prediction and robustness experiments, controlled HVAC-input mismatch and ablation analyses, and single-zone closed-loop comparison across repeated evaluation windows against RBC and adaptive RC-EKF-MPC, together with MPC-weight sensitivity and runtime assessments.

2. Overview of the Proposed Framework

Figure 1 summarizes the overall workflow connecting offline model development with surrogate-based closed-loop evaluation. In Stage 1, annual data generated using EnergyPlus version 22.1.0 are chronologically partitioned into training, validation, and test subsets. PG-NODE is trained offline, and the deployment checkpoint is selected exclusively using the validation subset according to the procedure described in Section 3.4. In Stage 2, the selected predictor remains fixed and is incorporated into the receding-horizon control loop, with each prediction horizon reinitialized using the latest observed indoor temperature. All controllers are evaluated under matched conditions using the same fixed seed-specific surrogate environment within each seed–window evaluation. EnergyPlus is used to generate the offline dataset and is not queried during the reported closed-loop rollouts. Detailed model formulation, feedback and update mechanisms, controller formulation, and evaluation protocols are presented in Sections 3–6.

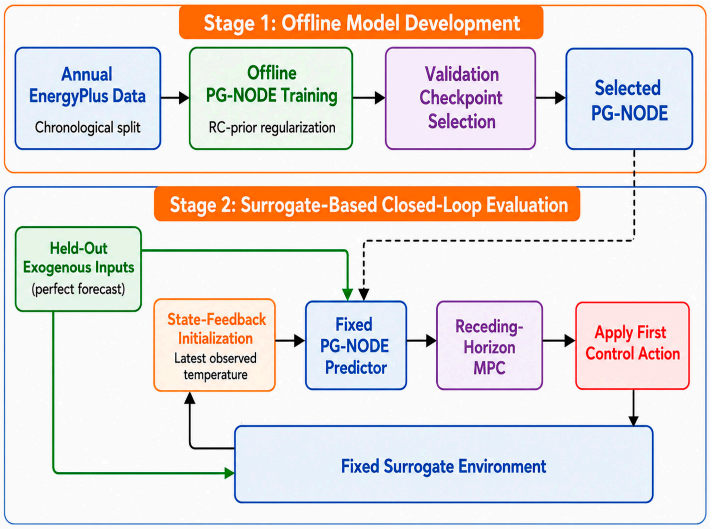

Figure 1. Overview of the proposed two-stage framework. PG-NODE is trained offline, selected using validation data, and then kept fixed during the main surrogate-based closed-loop evaluation. At each control step, the latest indoor temperature observation reinitializes the prediction horizon, while all controllers are evaluated using the same fixed seed-specific surrogate environment within each matched seed–window evaluation. Solid black arrows indicate the sequential model-development and closed-loop control flow, the dashed black arrow indicates transfer of the validation-selected PG-NODE to the fixed predictor used in Stage 2, and green arrows indicate the held-out exogenous inputs supplied to the predictor and surrogate environment.

3. Physics-Guided Neural ODE Thermal Prediction Model

3.1. PG-NODE Architecture

Neural ODEs parameterize continuous-time state derivatives using neural networks, providing a flexible framework for learning dynamical systems from temporal observations [21]. Building on the two-stage framework summarized in Figure 1, this section formulates the PG-NODE predictor used for offline model development and subsequent MPC deployment. The proposed formulation retains the continuous-time representation of neural ODEs while introducing structured guidance from a low-order RC-inspired thermal prior. The hybrid model architecture is first defined below, followed by the RC-prior directional regularization and numerical discretization in Sections 3.2 and 3.3, respectively.

3.1.1. Hybrid RC-Reference–Neural-Residual Formulation

To construct a control-oriented thermal predictor, the building thermal dynamics are decomposed into an RC-inspired reference component and a learnable neural residual. The former represents the dominant low-order thermal trend, whereas the latter captures residual nonlinearities and unmodeled dynamics beyond the selected prior. This semi-structured formulation reduces the burden on the neural network to learn the basic thermal response entirely from data and provides the reference trend used by the directional regularization introduced in Section 3.2. The continuous-time model is expressed as follows:

$$\frac{dh(t)}{dt} = f_{\mathrm{RC}}(h(t),\, u_{\mathrm{ctrl}}(t),\, \boldsymbol{d}(t);\, \boldsymbol{\theta}_{\mathrm{RC}}) + f_{\mathrm{NN}}(h(t),\, u_{\mathrm{ctrl}}(t),\, \boldsymbol{d}(t);\, \boldsymbol{\theta}_{\mathrm{NN}}) \tag{1}$$

where $h(t)$ denotes the normalized indoor thermal state, $u_{\mathrm{ctrl}}(t)$ denotes the normalized equivalent HVAC control input, and $\boldsymbol{d}(t) = [T_{\mathrm{out}}(t),\, S(t)]^{\mathrm{T}}$ denotes the exogenous input vector comprising outdoor temperature and solar radiation. The functions $f_{\mathrm{RC}}(\cdot)$ and $f_{\mathrm{NN}}(\cdot)$ denote the RC-inspired reference component and the neural residual component, parameterized by $\boldsymbol{\theta}_{\mathrm{RC}}$ and $\boldsymbol{\theta}_{\mathrm{NN}}$, respectively. Accordingly, the neural residual receives the four-dimensional input vector, as follows:

$$\boldsymbol{z}(t) = [h(t),\, u_{\mathrm{ctrl}}(t),\, T_{\mathrm{out}}(t),\, S(t)]^{\mathrm{T}} \tag{2}$$

which is consistent with the 4–32–32–1 architecture used in the implementation. The explicit low-order form and identification procedure of the RC-inspired reference component are described in Section 3.1.2.

3.1.2. RC-Inspired Prior Formulation and Identification

For each thermal zone, the RC-inspired reference component f_{RC} in Equation (1) was implemented as a low-order linear thermal prior rather than as a detailed multi-state RC network. Let $T_k \equiv h(t_k)$ denote the normalized indoor thermal state sampled at step k. For the transition from T_k to T_{k+1}, the sample-wise local temperature change trend $g_{\mathrm{RC},k}$ is expressed as follows:

$$g_{\mathrm{RC},k} = \boldsymbol{\theta}_{\mathrm{RC}}^{\mathrm{T}}[T_k,\, Q_{k+1},\, T_{\mathrm{out},k+1},\, S_{k+1},\, 1]^{\mathrm{T}} \tag{3}$$

where Q_{k+1}, $T_{\mathrm{out},k+1}$, and S_{k+1} denote the normalized interval-associated HVAC input, outdoor temperature, and solar radiation corresponding to the transition ending at T_{k+1}, respectively, and the final constant represents the intercept. The parameter vector $\boldsymbol{\theta}_{\mathrm{RC}} \in \mathbb{R}^5$ was identified separately for each thermal zone using ridge regression on the chronological training subset, with a regularization coefficient of $\alpha = 10^{-4}$. Accordingly, the RC-inspired reference function can be written compactly as $f_{\mathrm{RC}}(h, u, \boldsymbol{d}; \boldsymbol{\theta}_{\mathrm{RC}}) = \boldsymbol{\theta}_{\mathrm{RC}}^{\mathrm{T}}\left[h, u, \boldsymbol{d}^{\mathrm{T}}, 1\right]^{\mathrm{T}}$, while $g_{\mathrm{RC},k}$ denotes its sample-wise evaluation under the interval-aligned inputs in Equation (3). The resulting fixed RC-inspired prior provides the reference trend used by the directional regularization introduced in Section 3.2 and should not be interpreted as a complete first-principles thermal model.

3.2. RC-Prior Directional Regularization

The RC-inspired prior defined in Section 3.1.2 provides a low-order reference for the local temperature change direction. To encourage the learned PG-NODE dynamics to remain directionally consistent with this reference without imposing it as a hard physical constraint, an RC-prior directional regularization term is incorporated into the training objective. The total loss is defined as follows:

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{MSE}} + \lambda_{\mathrm{dir}}\mathcal{L}_{\mathrm{dir}} \tag{4}$$

where $\mathcal{L}_{\mathrm{MSE}}$ is the mean-squared-error prediction term, $\mathcal{L}_{\mathrm{dir}}$ is the RC-prior directional regularization term, and $\lambda_{\mathrm{dir}} \geq 0$ controls their relative weighting.

The prediction-error term is defined as follows:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N_b} \sum_{i=1}^{N_b} \left\| \widehat{h}(t_i) - h(t_i) \right\|_2^2 \tag{5}$$

where N_b denotes the number of samples or time steps in the current training batch, t_i denotes the i-th sampling instant, and $\hat{h}(t_i)$ and $h(t_i)$ denote the predicted and observed indoor thermal states, respectively.

For each training sample i, $g_{\text{RC},i}$ denotes the corresponding RC-prior temperature change trend defined by Equation (3). The local temperature change trend predicted by the PG-NODE is defined as follows:

$$\widehat{g}_i \equiv \left. \frac{d\widehat{h}(t)}{dt} \right|_{t=t_i} \tag{6}$$

A positive $g_{\text{RC},i}$ indicates a heating-oriented reference trend, whereas a negative value indicates a cooling- or heat-loss-oriented reference trend. The directional regularization term is then formulated as follows:

$$\mathcal{L}_{\text{dir}} = \frac{1}{N_b} \sum_{i=1}^{N_b} \mathbb{I}(|g_{\text{RC},i}| > \epsilon)[\text{ReLU}(m - \text{sign}(g_{\text{RC},i})\widehat{g}_i)]^2 \tag{7}$$

where $\mathbb{I}(\cdot)$ is the indicator function; $\text{sign}(\cdot)$ returns the sign of the RC-prior trend; $\text{ReLU}(\cdot)$ denotes the rectified linear unit, with $\text{ReLU}(x) = \max(0, x)$; m is the directional margin; and ϵ is the activity threshold used to exclude samples with weak RC-prior trends. In this study, $m = 0$. For each thermal zone, ϵ was determined exclusively from its chronological training subset as follows:

$$\epsilon = \max[0.015,\ P_{65}(|g_{\text{RC}}|_{\text{train}})] \tag{8}$$

where $P_{65}(\cdot)$ denotes the 65th percentile. The resulting threshold was fixed for subsequent validation and test evaluation. With $m = 0$, the penalty is nonzero only when an active sample has a predicted local trend opposite to the direction indicated by the selected RC-inspired prior. Accordingly, this term should be interpreted as soft prior-consistency regularization rather than as a strict energy-balance constraint or an independent guarantee of general physical validity.

Because regularization strength can materially affect predictive performance and the balance between data fitting and physics-related guidance [18], λ_{dir} was evaluated on the validation subset over $\{0,\ 1,\ 10,\ 20,\ 50,\ 100\}$. Based on this validation-only sensitivity analysis, $\lambda_{\text{dir}} = 10$ was retained as the nominal accuracy–prior-consistency trade-off for subsequent experiments rather than as an RMSE-optimal or universally optimal setting. The corresponding quantitative results are presented in Section 7.3 and Supplementary Table S1.

3.3. Numerical Discretization and Temporal Alignment for MPC

Following the continuous-time neural ODE formulation introduced by Chen et al. [21], the PG-NODE dynamics defined in Equation (1) are evaluated numerically at the discrete sampling instants used for model training and MPC prediction. Given an initial predicted thermal state $\hat{h}(t_0)$, the predicted trajectory can be expressed as follows:

$$\widehat{h}(t) = \widehat{h}(t_0) + \int_{t_0}^{t} \left[f_{\text{RC}}\left(\widehat{h}(\tau),\ u_{\text{ctrl}}(\tau),\ \boldsymbol{d}(\tau); \boldsymbol{\theta}_{\text{RC}}\right) + f_{\text{NN}}\left(\widehat{h}(\tau),\ u_{\text{ctrl}}(\tau),\ \boldsymbol{d}(\tau); \boldsymbol{\theta}_{\text{NN}}\right) \right] d\tau \tag{9}$$

where t_0 denotes the initial time and τ is the integration variable. The remaining variables and model parameters follow the definitions introduced in Section 3.1.

To ensure consistent temporal alignment between the stored EnergyPlus data and the subsequent control formulation, each one-step transition pairs the indoor temperature at sample k with the interval-associated HVAC and exogenous inputs leading to the next temperature sample, as follows:

$$(T_k, Q_{k+1}, T_{\text{out},k+1}, S_{k+1}) \to T_{k+1} \tag{10}$$

For the subsequent control-oriented formulation, these interval-associated inputs are re-indexed by the corresponding control interval as follows:

$$u_k \equiv Q_{k+1},\ \boldsymbol{d}_k \equiv [T_{\text{out},k+1},\ S_{k+1}]^{\text{T}} \tag{11}$$

so that the same transition can be written compactly as follows:

$$(T_k, u_k, \boldsymbol{d}_k) \to T_{k+1} \tag{12}$$

This is only a notational re-indexing and does not alter the underlying temporal alignment. The same convention is used throughout model training and evaluation, RC-prior identification, RC-EKF estimation, and closed-loop prediction.

In the implemented experiments, all NODE-based models used the same fixed-step forward Euler scheme. Time was expressed in sampling-step units, such that $\Delta\tau = 1$ corresponds to one 30 min sampling interval. Accordingly, the discrete PG-NODE transition is written as follows:

$$\widehat{h}_{k+1} = \widehat{h}_k + \Delta\tau\Big[f_{RC}\Big(\widehat{h}_k,\ u_k,\ \boldsymbol{d}_k; \boldsymbol{\theta}_{RC}\Big) + f_{\text{NN}}\Big(\widehat{h}_k,\ u_k,\ \boldsymbol{d}_k; \boldsymbol{\theta}_{NN}\Big)\Big] \tag{13}$$

where $\hat{h}_k$ is the predicted thermal state at sampling step k, and $\Delta\tau$ is the numerical integration step. The fixed RC-inspired reference function f_{RC} follows the formulation in Section 3.1.2 and is evaluated using the current predicted state and the re-indexed interval-associated inputs. The same fixed-step discretization was used throughout training, recursive prediction, and MPC evaluation; therefore, no adaptive ODE solver or relative or absolute integration tolerances were used.

3.4. Validation-Based Deployment Checkpoint Selection

For deployment in the main PG-NODE-MPC controller, candidate PG-NODE checkpoints were evaluated exclusively on the validation subset using a two-stage procedure. Each candidate was first screened for positive local HVAC-response gain, after which the admissible checkpoint with the lowest 48-step free-running validation rollout RMSE was selected for each model-training seed. This procedure was used solely for deployment-model selection and was not included in the training objective; the selected PG-NODE remained fixed throughout the main closed-loop evaluation. Complete checkpoint intervals, finite-difference settings, screening criteria, fallback rules, and selected epochs are reported in Supplementary Table S2.

4. Online Feedback and Parameter-Update Mechanisms

This section distinguishes state-feedback reinitialization in the fixed main PG-NODE-MPC controller from the restricted final-layer adaptation examined only in auxiliary ablation variants and the recursive parameter adaptation used by the RC-EKF-MPC benchmark.

4.1. State-Feedback Reinitialization in the Main PG-NODE-MPC

The main PG-NODE-MPC controller uses the validation-selected predictor described in Section 3.4, and its model parameters remain fixed throughout closed-loop operation. At each control step k, the latest observed indoor temperature is used to initialize the current prediction horizon:

$$\widehat{h}_{0|k} = T_k^{\mathrm{obs}} \tag{14}$$

where $\hat{h}_{0|k}$ denotes the initial predicted thermal state for the MPC optimization performed at control step k, and T_k^{obs} is the latest observed indoor temperature.

Using the temporal-alignment convention and discrete transition defined in Section 3.3, the corresponding one-step PG-NODE prediction is written as follows:

$$\widehat{T}_{k+1|k}^{\mathrm{PG}} = \mathbf{\Phi}_{\Delta\tau}\left(\widehat{h}_{0|k},\, u_k,\, \boldsymbol{d}_k;\, \mathbf{\Theta}_{PG}^{*}\right) \tag{15}$$

where $\mathbf{\Phi}_{\Delta\tau}(\cdot)$ denotes the one-step fixed-step PG-NODE transition over one control interval, and $\mathbf{\Theta}_{PG}^{*}$ denotes the fixed parameter set of the validation-selected predictor, comprising the identified RC-reference parameters and the trained neural-residual parameters. The interval-associated control input u_k and exogenous input vector $\boldsymbol{d}_k$ follow the definitions in Section 3.3.

After the next indoor temperature observation becomes available, the same initialization rule in Equation (14) is applied at control step $k+1$. The PG-NODE parameter set remains fixed throughout closed-loop operation:

$$\mathbf{\Theta}_{PG,k} = \mathbf{\Theta}_{PG}^{*},\ \forall k \tag{16}$$

where $\mathbf{\Theta}_{PG,k}$ denotes the PG-NODE parameter set used at control step k. Thus, online feedback updates the initial state of each prediction horizon but does not update the deployed PG-NODE parameters.

4.2. Restricted Final-Layer Adaptation for Ablation Variants

To examine the effect of limited online parameter adaptation while preserving the offline-learned PG-NODE representation, two auxiliary ablation variants were considered. M2 updates only the final output layer of the neural residual using the most recently observed transition, corresponding to a buffer length of $B = 1$, whereas M3 uses the five most recent transitions, corresponding to $B = 5$. The neural network parameters are partitioned as follows:

$$\boldsymbol{\theta}_{NN} = (\boldsymbol{\theta}_h, \boldsymbol{\theta}_o) \tag{17}$$

where $\boldsymbol{\theta}_h$ and $\boldsymbol{\theta}_o$ denote the hidden-layer and final output-layer parameter sets, respectively. During restricted adaptation, $\boldsymbol{\theta}_h$ remains fixed at its offline-trained value, whereas only $\boldsymbol{\theta}_o$ is updated.

After a new transition becomes fully observed, a buffer containing the most recent B transitions is formed as follows:

$$\mathcal{B}_k^{(B)} = \left\{\left(T_i^{\mathrm{obs}}, u_i, \boldsymbol{d}_i, T_{i+1}^{\mathrm{obs}}\right)\right\}_{i=k-B+1}^{k} \tag{18}$$

where $\mathcal{B}_k^{(B)}$ denotes the adaptation buffer containing the most recent B fully observed transitions. The remaining variables follow the definitions introduced in Sections 3.3 and 4.1. The local adaptation loss is as follows:

$$\mathcal{L}_{\mathrm{ad},k} = \frac{1}{\left|\mathcal{B}_k^{(B)}\right|}\sum\nolimits_{i\in\mathcal{B}_k^{(B)}}\left[\left(\widehat{T}_{i+1\,|\,i} - T_{i+1}^{\mathrm{obs}}\right)^2 + \lambda_{\mathrm{dir}}\ell_{\mathrm{dir},i}\right] \tag{19}$$

where $\mid \mathcal{B}_k^{(B)} \mid = B$ is the number of transitions in the current buffer, $\hat{T}_{i+1|i}$ is the one-step prediction produced by the current model, and $\ell_{\mathrm{dir},i}$ denotes the indicator-weighted squared directional penalty inside the summation in Equation (7). The nominal λ_{dir} setting selected in Section 3.2 is retained during adaptation.

After each newly observed transition becomes available, the final output-layer parameters are updated by one optimizer step using the settings reported in Supplementary Table S2. State-feedback reinitialization, as defined in Section 4.1, is retained in both variants. M2 and M3 are auxiliary ablation variants and are not part of the main PG-NODE-MPC controller.

4.3. EKF-Based Adaptive RC Model for the RC-EKF-MPC Baseline

Kalman-filter-based online estimation of grey-box thermal models and adaptive predictive control have previously been investigated for building thermal modeling and control [32,33]. In the present study, RC-EKF-MPC serves as a separate adaptive grey-box benchmark based on recursive estimation of a five-parameter RC-inspired linear predictor. Unlike the fixed PG-NODE predictor used in the main controller, its parameter vector is updated recursively during closed-loop operation. The initial RC-EKF parameter vector was identified offline by ordinary least squares using a predefined initial segment of the chronological training subset. The exact initialization settings are provided in Supplementary Table S2.

Following this established recursive estimation framework and adopting the control-oriented notation defined in Section 3.3, the RC-EKF predictor and parameter-update equations are written as follows. The one-step predictor is given by the following:

$$\hat{T}^{\text{RC}}_{k+1|k} = \boldsymbol{\phi}_k^{\text{T}} \boldsymbol{\theta}_k^{-} \tag{20}$$

$$\boldsymbol{\phi}_k = \left[T_k^{\text{obs}},\, u_k,\, \boldsymbol{d}_k^{\text{T}},\, 1 \right]^{\text{T}} \tag{21}$$

where $\hat{T}^{\text{RC}}_{k+1|k}$ denotes the one-step-ahead indoor temperature prediction produced by the RC-EKF model, $\boldsymbol{\phi}_k \in \mathbb{R}^5$ is the RC-EKF regression vector, and $\boldsymbol{\theta}_k^{-} \in \mathbb{R}^5$ is the prior RC parameter estimate available before assimilating the next temperature observation. The superscripts $(-)$ and $(+)$ denote prior and posterior quantities, respectively. The quantities T_k^{obs}, u_k, and $\boldsymbol{d}_k$ follow the definitions introduced in Sections 3.3 and 4.1. Because the predictor is linear in the recursively estimated parameter vector, the corresponding EKF parameter update reduces to the standard Kalman-filter form used below.

The RC parameters are modeled as a random-walk process with zero-mean process noise. Accordingly, before each measurement update, the prior parameter mean and covariance are propagated as follows:

$$\boldsymbol{\theta}_k^{-} = \boldsymbol{\theta}_{k-1}^{+}, \;\; \boldsymbol{P}_k^{-} = \boldsymbol{P}_{k-1}^{+} + \boldsymbol{Q}_{\text{EKF}} \tag{22}$$

where $\boldsymbol{\theta}_{k-1}^{+}$ and $\boldsymbol{P}_{k-1}^{+}$ denote the posterior parameter estimate and covariance matrix from the preceding update, respectively; $\boldsymbol{P}_k^{-}$ is the prior parameter covariance matrix at step k; and $\boldsymbol{Q}_{\text{EKF}}$ is the process-noise covariance matrix.

After T_{k+1}^{obs} becomes available, the prediction innovation is calculated as follows:

$$\nu_{k+1} = T_{k+1}^{\text{obs}} - \boldsymbol{\phi}_k^{\text{T}} \boldsymbol{\theta}_k^{-} \tag{23}$$

where ν_{k+1} denotes the one-step temperature-prediction residual. The Kalman gain is then computed as follows:

$$\boldsymbol{K}_k = \frac{\boldsymbol{P}_k^{-} \boldsymbol{\phi}_k}{\boldsymbol{\phi}_k^{\text{T}} \boldsymbol{P}_k^{-} \boldsymbol{\phi}_k + R_{\text{EKF}}} \tag{24}$$

where $\boldsymbol{K}_k$ is the Kalman gain vector and R_{EKF} is the measurement-noise variance. The posterior parameter estimate and covariance are updated according to the following:

$$\boldsymbol{\theta}_k^+ = \boldsymbol{\theta}_k^- + \boldsymbol{K}_k \nu_{k+1} \tag{25}$$

$$\boldsymbol{P}_k^+ = \left(\boldsymbol{I}_5 - \boldsymbol{K}_k \boldsymbol{\phi}_k^{\mathrm{T}}\right) \boldsymbol{P}_k^- \tag{26}$$

where $\boldsymbol{I}_5$ denotes the 5×5 identity matrix.

For the controller-comparison experiments, $\boldsymbol{Q}_{\mathrm{EKF}} = 10^{-4} \boldsymbol{I}_5$ and $R_{\mathrm{EKF}} = 10^{-3}$ were selected through validation-only sensitivity analysis. The complete numerical covariance-sensitivity results are reported only in Supplementary Table S3. The RC-EKF parameters were updated after each newly observed 30 min transition, and the resulting predictor was used in the subsequent MPC optimization under the same objective function, prediction horizon, exogenous forecasts, and normalized control bounds as PG-NODE-MPC. For clarity, Table 2 summarizes the distinctions among the fixed main PG-NODE-MPC controller, the restricted M2/M3 adaptation variants, and the adaptive RC-EKF-MPC benchmark.

Table 2. Summary of the feedback and parameter-update mechanisms considered in this study.

Method	State-Feedback Reinitialization	Online Parameter Update	Updated Parameters	Feedback/Update Information	Update Frequency
Main PG-NODE-MPC	Yes	No	None	Latest observed indoor temperature for state reinitialization	At each control step; no parameter update
M2	Yes	Yes	Final neural-residual output layer	Most recent observed transition, $B = 1$	After each observed transition
M3	Yes	Yes	Final neural-residual output layer	Five most recent observed transitions, $B = 5$	After each observed transition
RC-EKF-MPC	Yes	Yes	Five-parameter RC vector	Prediction innovation from the latest temperature observation	After each newly observed transition

Note: M2 and M3 are used only in the ablation analysis and are not part of the main PG-NODE-MPC controller.

5. Model Predictive Control Formulation

The PG-NODE-MPC and RC-EKF-MPC controllers use the same receding-horizon optimization formulation but differ in their prediction models and online model update behavior, as described in Section 4: the PG-NODE predictor remains fixed, whereas the RC-EKF predictor is recursively updated. The common objective function, control constraints, predictive rollout, and receding-horizon implementation are described below.

5.1. Objective Function

Following the standard receding-horizon MPC principle [3,4] and the control-input convention defined in Section 3.3, the study-specific objective used in this work comprises tariff-weighted normalized control effort, tracking of the normalized reference state, and control smoothness. At each closed-loop control step t, the control sequence is optimized over an N-step prediction horizon according to the following:

$$\min_{\tilde{\boldsymbol{U}}_t} J_t = \sum_{j=0}^{N-1} \left[c_{t+j} |u_{t+j}| + \lambda_c \left(\hat{h}_{j+1\,|\,t} - h_{\mathrm{ref}} \right)^2 + \lambda_{\Delta u} \left(u_{t+j} - u_{t+j-1} \right)^2 \right] \tag{27}$$

where t denotes the current closed-loop control step and j indexes the prediction steps within the current MPC horizon. The decision vector $\tilde{\boldsymbol{U}}_t$ contains the unconstrained control variables over the horizon and is mapped to the bounded control sequence as described in Section 5.2. The quantity $N = 48$ denotes the prediction horizon, corresponding to 24 h at the 30 min control interval; u_{t+j} is the normalized equivalent HVAC control input; c_{t+j} is the normalized time-of-use weighting coefficient; $\hat{h}_{j+1|t}$ is the indoor thermal state predicted $j+1$ steps ahead from control step t; h_{ref} is the normalized reference thermal state; and λ_c and $\lambda_{\Delta u}$ are the reference-tracking and control-smoothness penalty weights, respectively. For the first smoothness term, corresponding to $j = 0$, u_{t-1} denotes the control action applied at the preceding closed-loop control step.

The reference state was fixed at $h_{\text{ref}} = 0$ for all controller-comparison experiments. The three terms in the objective respectively represent tariff-weighted normalized control effort, reference-tracking penalty, and control-smoothness penalty.

The nominal time-of-use weighting coefficient was set to 0.40 during the peak period from 08:00 to 20:00 and to 0.10 otherwise. Sensitivity to alternative tariff weightings was evaluated post hoc using the fixed closed-loop trajectories without controller re-optimization; the complete post hoc results for tariff-weighted normalized control effort and the resulting reweighted total objective are reported only in Supplementary Table S4. Both MPC-based controllers used the common nominal weights $\lambda_c = 100$ and $\lambda_{\Delta u} = 0.1$. A validation-only sensitivity analysis was conducted to determine whether the qualitative controller comparison depended strongly on these weights. The full numerical results are reported only in Supplementary Table S5 and were not used to retune the controllers on the held-out test windows. The nominal weights are therefore treated as fixed common settings rather than as globally optimal parameters.

Because u_{t+j} is a normalized equivalent thermal-control coordinate rather than measured electrical power, the tariff-weighted control-effort term and the resulting total objective are used only for relative algorithmic comparison and should not be interpreted as physical energy consumption, HVAC electricity use, or monetary cost.

5.2. Control-Input Constraints

To enforce the prescribed normalized control-input range, each element of the unconstrained decision vector $\tilde{\boldsymbol{U}}_t$ is mapped through a smooth saturation function:

$$u_{t+j} = u_{\max} \tanh\left(\widetilde{u}_{t+j}\right), \quad j = 0, \ldots, N-1 \tag{28}$$

where $\widetilde{u}_{t+j}$ is the unconstrained decision variable at prediction step j, and $u_{\max} = 2.0$ is the prescribed normalized control bound. Hence,

$$-u_{\max} \le u_{t+j} \le u_{\max}, \quad j = 0, \ldots, N-1 \tag{29}$$

The same control bound is applied to both PG-NODE-MPC and RC-EKF-MPC to ensure a consistent controller comparison. No hard ramp-rate constraint is imposed; control smoothness is instead encouraged through the penalty term $\lambda_{\Delta u}\left(u_{t+j} - u_{t+j-1}\right)^2$ in the objective defined in Section 5.1.

5.3. Predictive Rollout Within MPC

At each closed-loop control step t, the prediction horizon is reinitialized using the latest observed indoor temperature according to the state-feedback rule defined in Section 4.1. For PG-NODE-MPC, the fixed validation-selected predictor is then recursively propagated over the N-step horizon as follows:

$$\widehat{h}_{j+1|t} = \Phi_{\Delta\tau}\left(\widehat{h}_{j|t}, u_{t+j}, \boldsymbol{d}_{t+j}; \boldsymbol{\Theta}^*_{PG}\right), \quad j = 0, \ldots, N-1 \tag{30}$$

where $\hat{h}_{j|t}$ denotes the thermal state predicted j steps ahead from control step t, and $\boldsymbol{d}_{t+j}$ is the corresponding exogenous input vector over the prediction horizon. The transition operator and fixed PG-NODE parameter set are defined in Section 4.1, while the temporal-alignment convention and control-input constraints are defined in Sections 3.3 and 5.2, respectively.

For RC-EKF-MPC, the same receding-horizon rollout structure is used, with the PG-NODE transition replaced by the most recently updated RC-EKF predictor described in Section 4.3. The resulting predicted trajectories are evaluated using the common MPC objective defined in Section 5.1.

5.4. Receding-Horizon Implementation

At each closed-loop control step t, the MPC problem defined in Sections 5.1–5.3 is solved over the N-step prediction horizon, yielding the following optimized control sequence:

$$\boldsymbol{U}_t^* = \left[u_t^*, u_{t+1}^*, \ldots, u_{t+N-1}^*\right]^{\mathrm{T}} \tag{31}$$

where $\boldsymbol{U}_t^*$ denotes the optimal bounded control sequence at control step t, and u_{t+j}^* is the corresponding optimized control action at prediction step j.

Each MPC problem was solved using 50 iterations of the optimizer specified in Supplementary Table S2, with a learning rate of 0.1. The unconstrained decision vector $\widetilde{\boldsymbol{U}}_t$ was initialized to zero at each control step, with no warm start or tolerance-based early stopping. The same optimization settings were applied to both MPC-based controllers.

Following the receding-horizon principle, only the first optimized control action is applied:

$$u_t = u_t^* \tag{32}$$

After the system advances to the next control step and a new indoor temperature observation becomes available, the prediction horizon is shifted forward and the optimization is repeated using the controller-specific feedback or update mechanism described in Section 4.

6. Experimental Setup and Benchmark Design

6.1. Target Building and Simulation Setup

The case study was implemented in EnergyPlus [34] using the U.S. Department of Energy (DOE) Standalone Retail prototype from the DOE Prototype Building Models [35], under the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Standard 90.1-2022 reference configuration [36]. The annual simulation for New York (ASHRAE Climate Zone 4A) used the corresponding Typical Meteorological Year 3 (TMY3) weather data [37]. Five thermal zones were considered: BACK SPACE, CORE RETAIL, FRONT ENTRY, FRONT RETAIL, and POINT OF SALE.

The annual dataset contained 17,520 samples at a 30 min reporting interval. For each thermal zone, the data were divided chronologically, without shuffling, into 60% training, 20% validation, and 20% held-out test subsets, corresponding to 10,512, 3504, and 3504 samples, respectively. All normalization statistics were calculated exclusively from the training subset. The temporal alignment of the indoor temperature, HVAC, and exogenous variables follows the convention defined in Section 3.3.

All five thermal zones were included in the repeated prediction, Gaussian noise robustness, and controlled HVAC-input mismatch experiments. The final closed-loop controller comparison was restricted to the BACK SPACE zone. Accordingly, the multi-zone experiments characterize predictive behavior across zones within the same prototype building, whereas the closed-loop conclusions are limited to the evaluated BACK SPACE simulation case. The detailed benchmark definitions and closed-loop evaluation protocol are provided in Sections 6.2–6.5.

6.2. Evaluation Metrics and Gaussian Noise Robustness Protocol

To evaluate prediction robustness under noisy sensing and disturbance measurements, independent zero-mean Gaussian perturbations were added in normalized coordinates to indoor temperature, outdoor temperature, and solar radiation, whereas the normalized equivalent HVAC input remained unperturbed:

$$\tilde{x}_t = x_t + \eta_t,\ \eta_t \sim \mathcal{N}\left(0, \sigma^2\right),\ \ \sigma \in \{0, 0.3, 0.6\} \tag{33}$$

where $x_t \in \{T_t, T_{\text{out},t}, S_t\}$ denotes a normalized variable at time step t, $\tilde{x}_t$ is its perturbed counterpart, η_t denotes an independent Gaussian perturbation sampled for each perturbed variable and time step, and σ is the prescribed noise standard deviation in normalized coordinates. For each thermal zone, noise level, and model-training seed, the learned prediction models were retrained using the corresponding noisy training subset and evaluated on an independently perturbed held-out test subset. The validation-based deployment checkpoint-selection procedure described in Section 3.4 was not applied in this robustness experiment; instead, the final model obtained after the prescribed 150 training epochs was used directly for test evaluation. The validation subset was therefore not used. The RC-inspired prior identified from the clean chronological training subset remained fixed across all noise levels. Detailed noise generation, retraining, random seed, and model selection settings are provided in Supplementary Table S2.

Prediction accuracy was quantified using normalized prediction RMSE:

$$RMSE = \sqrt{\frac{1}{N_{eval}} \sum_{i=1}^{N_{eval}} \left(\hat{T}_i - T_i\right)^2} \tag{34}$$

where N_{eval} denotes the number of evaluated samples, and $\hat{T}_i$ and T_i are the predicted and corresponding observed normalized indoor temperatures, respectively. For the Gaussian noise experiments, the metrics were calculated on the corresponding perturbed held-out test sequence.

Directional agreement with the selected RC-inspired prior was evaluated using the RC-prior directional inconsistency rate (RDI). Using the activity threshold ϵ defined in Section 3.2, the active sample set is defined as $\mathcal{A} = \{i : |g_{\text{RC},i}| > \epsilon\}$. RDI is then calculated as follows:

$$RDI = \frac{1}{|\mathcal{A}|} \sum_{i \in \mathcal{A}} \mathbb{I}[\hat{g}_i g_{\text{RC},i} < 0] \times 100\% \tag{35}$$

where $|\ \mathcal{A}\ |$ is the number of active samples. A lower RDI indicates greater directional agreement with the selected RC-inspired prior. The RC baseline has zero RDI by construction because it defines the reference direction. Accordingly, RDI is used here as a prior-consistency metric rather than as an independent measure of general physical validity.

Closed-loop performance was summarized using normalized reference-tracking RMSE, the three objective components defined in Section 5.1, total objective, and cumulative control variation. In this closed-loop context, RMSE denotes the run-wise root mean square deviation of the realized normalized indoor temperature trajectory from the prescribed reference state $h_{\text{ref}} = 0$. This is used as a reference-tracking metric and should not be interpreted as a formal measure of occupant thermal comfort. Cumulative control variation was calculated as follows:

$$TV_u = \sum_{k=1}^{N_{\text{sim}}-1} |u_k - u_{k-1}| \tag{36}$$

where N_{sim} denotes the number of evaluated closed-loop control steps. Lower TV_u indicates smoother control trajectories.

Two additional temperature-scale reference-tracking diagnostics were calculated exclusively for supplementary trajectory interpretation: the mean absolute temperature deviation from the reference, expressed in degrees Celsius, and the percentage of samples outside the diagnostic temperature band. These metrics are reported only in Supplementary Table S6 and were not used in the MPC objective, controller tuning, or primary inferential comparisons. The normalized trajectories were inverse-transformed using the training-subset statistics. The diagnostic band is neither a hard MPC constraint nor a formal thermal-comfort standard.

6.3. Benchmark Models and Controllers

Based on the evaluation metrics defined in Section 6.2, five thermal-prediction models and three closed-loop controllers were considered. All prediction models used the same chronological data split, temporal-alignment convention, and normalization protocol described in Sections 3.3 and 6.1.

For thermal prediction, the following models were compared. The RC baseline is the fixed low-order linear RC-inspired prior defined in Section 3.1.2. The LSTM baseline uses a single recurrent layer with a six-step input window as a representative data-driven sequence model [38]. The vanilla NODE uses the same 4–32–32–1 neural architecture and fixed-step discretization as the other NODE-based models but is trained using only the prediction-error objective without RC-prior directional regularization [21]. The regularized NODE is a study-defined comparison baseline using the same architecture with a three-step recursive rollout loss to encourage multi-step predictive consistency, without RC-prior directional regularization. The proposed PG-NODE combines the RC-inspired reference component, neural residual, and RC-prior directional regularization described in Section 3; the validation-selected fixed checkpoint is used in the main PG-NODE-MPC controller.

For closed-loop evaluation, three controllers were compared. The RBC applies a deadband of ± 0.10 around the normalized reference, with $u = +2$ below the lower threshold, $u = -2$ above the upper threshold, and $u = 0$ within the deadband. The RC-EKF-MPC uses the adaptive RC predictor described in Section 4.3 within the common MPC formulation of Section 5, whereas PG-NODE-MPC uses the fixed validation-selected PG-NODE predictor with the state-feedback reinitialization described in Section 4.1. Within each matched seed–window evaluation, all three controllers were evaluated in the same fixed seed-specific surrogate environment. The complete pairing protocol and shared closed-loop operating conditions are specified in Section 6.4, and detailed implementation settings are provided in Supplementary Table S2.

6.4. Multi-Window Closed-Loop Evaluation Protocol

The final closed-loop comparison was conducted for the BACK SPACE zone using five predefined, mutually non-overlapping 48 h windows from the held-out test subset, spanning 24 October to 29 December. Five model-realization seeds were evaluated in each window, yielding 25 matched seed–window cells per controller. The exact window dates, offsets, and endpoint conventions are reported in Supplementary Table S2. Because the evaluation period covers only late autumn to early winter, these results should not be interpreted as year-round validation or as separate validation of representative heating and cooling seasons.

For each model-realization seed, the fixed surrogate environment and the validation-selected PG-NODE controller predictor were trained independently and shared no parameters. Within each matched seed–window cell, RBC, RC-EKF-MPC, and PG-NODE-MPC were evaluated under the same surrogate dynamics, exogenous sequence, initial condition, and common operating settings. The five final test windows were not used for checkpoint selection or controller tuning, and the initial 96-step test segment used for implementation verification was excluded from all reported performance analyses. Complete surrogate-training, seed-pairing, environment-reuse, and evaluation-window settings are provided in Supplementary Table S2.

Both MPC-based controllers followed the common receding-horizon formulation described in Section 5. Future outdoor temperature and solar radiation inputs were taken directly from the corresponding held-out sequence under a perfect exogenous forecast assumption. Occupancy-related and other unmeasured internal gains were not supplied as separately forecast variables and were represented only implicitly through the learned dynamics and surrogate environment. Accordingly, the present closed-loop evaluation does not assess robustness to exogenous forecast errors.

All reported closed-loop controller comparisons were therefore conducted in independently trained, fixed, seed-specific PG-NODE surrogate environments derived from EnergyPlus-generated data rather than through direct EnergyPlus co-simulation or actuation. Controller performance was summarized descriptively across the 25 matched seed–window cells as mean ± standard deviation. Because the same five evaluation windows were used across the five model-realization seeds, the cells were not treated as fully independent replicates for inferential analysis.

For each performance metric and each prespecified comparison of PG-NODE-MPC with RBC or RC-EKF-MPC, the matched controller difference $d_{s,w}$ within model-realization seed s and evaluation window w was analyzed using a two-way additive fixed-block linear model:

$$d_{s,w} = \mu + \alpha_s + \beta_w + e_{s,w} \tag{37}$$

subject to

$$\sum_{s=1}^{5} \alpha_s = 0, \sum_{w=1}^{5} \beta_w = 0$$

where μ denotes the overall mean matched difference, α_s and β_w denote the fixed blocking effects of model-realization seed and evaluation window, respectively, and $e_{s,w}$ denotes the residual term. The sum-to-zero constraints ensure that the model intercept μ represents the overall mean matched difference. Statistical inference was based on the two-sided null hypothesis $H_0 : \mu = 0$. Because only one matched difference was available for each seed–window cell, the seed-by-window interaction was not separately estimable and was therefore confounded with the residual term. With 25 matched cells and nine estimable model coefficients, the residual degrees of freedom were 16.

6.5. Controlled HVAC-Input Mismatch Protocol

In addition to the closed-loop evaluation described in Section 6.4, a separate fixed-model stress test was conducted to examine the sensitivity of vanilla NODE and PG-NODE to controlled mismatch in the assumed HVAC input–response relationship. This experiment used the held-out test subset and evaluated the predictors only; it was not intended to represent a specific equipment fault or an online adaptation scenario.

The mismatch was introduced by scaling only the normalized interval-associated HVAC input supplied to the predictor:

$$u_k^{\text{mis}} = \gamma_k u_k \tag{38}$$

with the following:

$$\gamma_k = \begin{cases} 1, & k < k_d, \\ \gamma, & k \geq k_d, \end{cases}, \; k_d = \left\lfloor \frac{N_{\text{test}}}{2} \right\rfloor \tag{39}$$

where u_k follows the interval-associated input definition in Section 3.3, u_k^{mis} is the perturbed input supplied to the predictor, N_{test} is the number of aligned transitions in the held-out test sequence, and k_d denotes the mismatch-onset index. Three gain factors were considered: $\gamma \in \{1.0, 0.8, 0.6\}$, where $\gamma = 1.0$ represents the matched nominal condition. Indoor temperature observations and exogenous inputs were left unchanged, so the experiment isolates the effect of a controlled perturbation in the HVAC input channel. The model parameters remained fixed throughout the test. At each prediction step, the latest observed indoor temperature was used as the current state, consistent with the state-feedback reinitialization convention described in Section 4.1.

The evaluation was repeated across five thermal zones and five model-training seeds. Prediction performance after the mismatch onset was quantified as follows:

$$\text{RMSE}_{\text{post},\gamma} = \sqrt{\frac{1}{N_{\text{post}}} \sum_{k=k_d}^{N_{\text{test}}-1} \left(\hat{T}_{k+1\,|\,k}^{(\gamma)} - T_{k+1}^{\text{obs}} \right)^2} \tag{40}$$

where $\hat{T}_{k+1|k}^{(\gamma)}$ denotes the one-step-ahead prediction under gain factor γ, and $N_{\text{post}} = N_{\text{test}} - k_d$ denotes the number of post-mismatch transitions.

To account for baseline differences across matched zone–seed–model runs, mismatch-induced degradation was evaluated relative to the matched nominal condition at $\gamma = 1$. The absolute RMSE increase was defined as follows:

$$\Delta RMSE_{\text{abs},\gamma} = RMSE_{\text{post},\gamma} - RMSE_{\text{post},1} \tag{41}$$

and the corresponding relative RMSE increase was defined as follows:

$$\Delta RMSE_{\text{rel},\gamma} = \frac{RMSE_{\text{post},\gamma} - RMSE_{\text{post},1}}{RMSE_{\text{post},1}} \times 100\% \tag{42}$$

Both prediction-RMSE change metrics were calculated within each matched zone–seed–model run before aggregation across the 25 matched zone–seed runs for each model and mismatch setting. Detailed post-mismatch prediction RMSE values are reported only in Supplementary Table S7.

7. Results and Discussion

7.1. Prediction Accuracy and RC-Prior Directional Consistency

Table 3 summarizes the prediction performance across five thermal zones and five model-training seeds. LSTM achieved the lowest overall RMSE of 0.1017 ± 0.0423, whereas PG-NODE achieved the lowest RMSE of 0.1314 ± 0.0360 among the NODE-based models. Thus, the proposed model did not provide the highest prediction accuracy among all evaluated models.

Figure 2 shows the trade-off between prediction accuracy and RC-prior directional consistency. PG-NODE achieved a mean RDI of 4.49%, compared with 17.66% for vanilla NODE, 18.63% for regularized NODE, and 27.18% for LSTM. Relative to vanilla NODE, PG-NODE reduced mean RMSE by approximately 2.4% and mean RDI by 74.6%, indicating lower mean RMSE and substantially lower mean RDI than vanilla NODE within the evaluated design. Because RDI is defined relative to the identified RC prior, the low

RDI of the PG-NODE reflects a stronger agreement with that prior rather than reflecting independent evidence of general physical validity.

Table 3. Prediction performance of the five benchmark models across five thermal zones and five model-training seeds.

Model	RMSE	RDI (%)
RC	0.1603 ± 0.0388	0.00 *
LSTM	0.1017 ± 0.0423	27.18 ± 9.10
Vanilla NODE	0.1346 ± 0.0349	17.66 ± 14.62
Regularized NODE	0.1332 ± 0.0332	18.63 ± 11.45
PG-NODE	0.1314 ± 0.0360	4.49 ± 5.71

Notes: Values are reported as mean ± standard deviation across five thermal zones and five model-training seeds ($n = 25$ per model). * The RC baseline has RDI = 0 by construction because it defines the reference direction. RMSE denotes the normalized prediction RMSE defined in Equation (34).

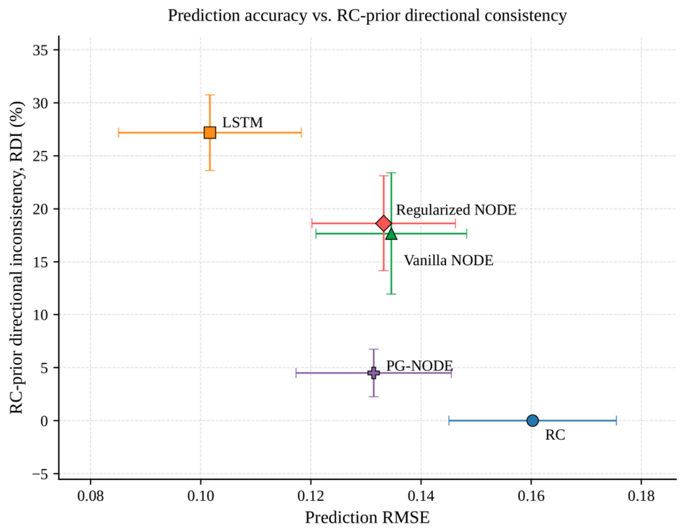

Figure 2. Trade-off between prediction accuracy and RC-prior directional consistency for the five benchmark models. Markers represent mean values across five thermal zones and five model-training seeds ($n = 25$ runs per model); the horizontal and vertical error bars indicate normal approximation 95% confidence intervals of the mean for prediction RMSE and RDI, respectively. PG-NODE achieved the lowest mean RDI among the learned nonlinear models and the lowest mean prediction RMSE among the NODE-based models.

The zone-wise results in Figure 3 show that prediction accuracy varies across thermal zones, while PG-NODE maintains the lowest mean RDI among the learned nonlinear models in all five evaluated zones. Detailed zone-wise numerical prediction RMSE and RDI values are reported only in Supplementary Table S8, while Supplementary Figure S1 shows the full distributions across the 25 zone–seed runs. These results indicate that the aggregate accuracy–prior-consistency pattern observed in Table 3 is broadly maintained across the evaluated zones.

7.2. Sensitivity to Controlled HVAC-Input Model Mismatch

Figure 4 compares the sensitivity of the fixed Vanilla NODE and PG-NODE predictors to the controlled HVAC-input mismatch defined in Section 6.5. As the input-gain factor decreased from the nominal value of $\gamma = 1.0$ to 0.8 and 0.6, the mean absolute RMSE increases were 0.01087 and 0.03709 for Vanilla NODE and 0.01294 and 0.03913 for PG-NODE, respectively. The corresponding relative RMSE increases were 7.82% and 26.69% for Vanilla NODE and 9.26% and 28.26% for PG-NODE. These results indicate that PG-NODE did not exhibit greater robustness to the controlled input mismatch, with the two fixed predictors showing broadly comparable degradation. Although the detailed results in Supplementary Table S7 show marginally lower post-mismatch prediction RMSE values for PG-NODE, the differences were small and should not be interpreted as evidence of superior mismatch robustness.

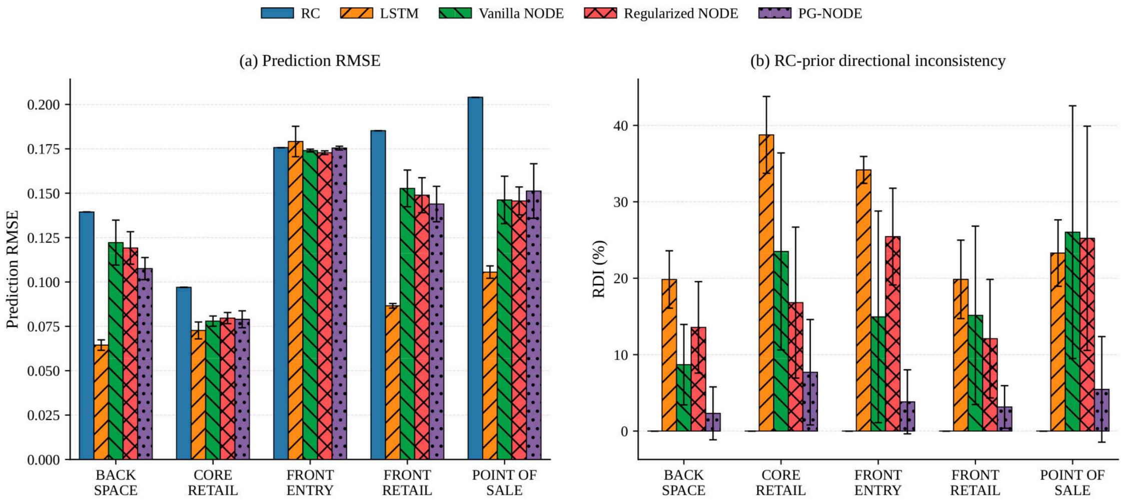

Figure 3. Zone-wise prediction performance of the five benchmark models: (**a**) prediction RMSE and (**b**) RDI. For each zone–model combination, bars represent the mean across five model-training seeds ($n = 5$), and error bars indicate normal approximation 95% confidence intervals of the mean. Across all five evaluated zones, PG-NODE maintained the lowest mean RDI among the learned nonlinear models, whereas prediction-accuracy rankings varied across zones.

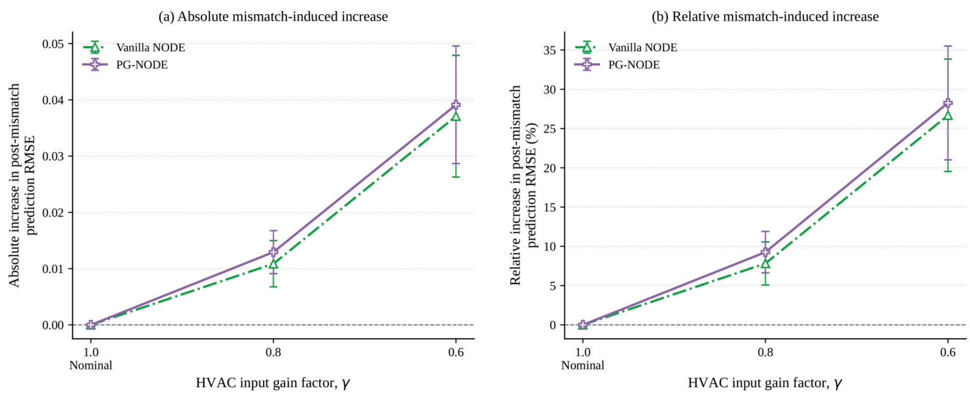

Figure 4. Sensitivity of the fixed Vanilla NODE and PG-NODE predictors to controlled HVAC-input mismatch: (**a**) absolute increase in post-mismatch prediction RMSE relative to the matched nominal condition and (**b**) relative increase in prediction RMSE. Changes were calculated within matched zone–seed–model runs and then aggregated across five thermal zones and five model-training seeds ($n = 25$ runs per model and mismatch setting). Markers represent mean values, and error bars indicate normal approximation 95% confidence intervals of the mean. Model parameters remained fixed throughout the experiment. The two predictors exhibited broadly comparable degradation as the imposed HVAC-input mismatch increased, with no clear robustness advantage for PG-NODE.

7.3. Gaussian Noise Robustness and Directional Regularization Sensitivity

Table 4 and Figure 5 summarize prediction performance under the Gaussian noise protocol defined in Section 6.2. Prediction RMSE increased for all models as the noise level increased. LSTM retained the lowest mean RMSE at both nonzero noise levels, achieving 0.3599 at $\sigma = 0.3$ and 0.6608 at $\sigma = 0.6$. The three NODE-based models exhibited broadly similar RMSE under the same perturbations; therefore, PG-NODE did not provide a consistent prediction accuracy advantage under noisy inputs. The principal difference was observed in RC-prior directional consistency. PG-NODE maintained mean RDI values of 3.65% and 6.25% at $\sigma = 0.3$ and 0.6, respectively, remaining substantially below those of the other learned nonlinear models. Thus, under the evaluated Gaussian perturbations, RC-prior directional regularization primarily preserved stronger agreement with the selected RC-inspired prior rather than reducing absolute prediction error.

Table 4. Prediction performance under nonzero Gaussian noise perturbations across five thermal zones and five model-training seeds.

σ	Model	RMSE	RDI (%)
0.3	RC	0.4328 ± 0.0146	0.00 *
	LSTM	0.3599 ± 0.0180	23.33 ± 3.39
	Vanilla NODE	0.4132 ± 0.0121	10.42 ± 6.41
	Regularized NODE	0.4124 ± 0.0137	12.36 ± 4.48
	PG-NODE	0.4141 ± 0.0142	3.65 ± 2.65
0.6	RC	0.8183 ± 0.0168	0.00 *
	LSTM	0.6608 ± 0.0136	24.85 ± 3.23
	Vanilla NODE	0.7522 ± 0.0154	17.17 ± 7.46
	Regularized NODE	0.7564 ± 0.0157	12.25 ± 3.67
	PG-NODE	0.7596 ± 0.0119	6.25 ± 3.24

Note: Values are reported as mean ± standard deviation across five thermal zones and five model-training seeds ($n = 25$ per model and noise level). RMSE denotes the normalized prediction RMSE defined in Equation (34). * indicates that the RC baseline has RDI = 0 by construction because it defines the reference direction. The noise-free baseline ($\sigma = 0$) is reported in Table 3 and is included in Figure 5 for visual comparison.

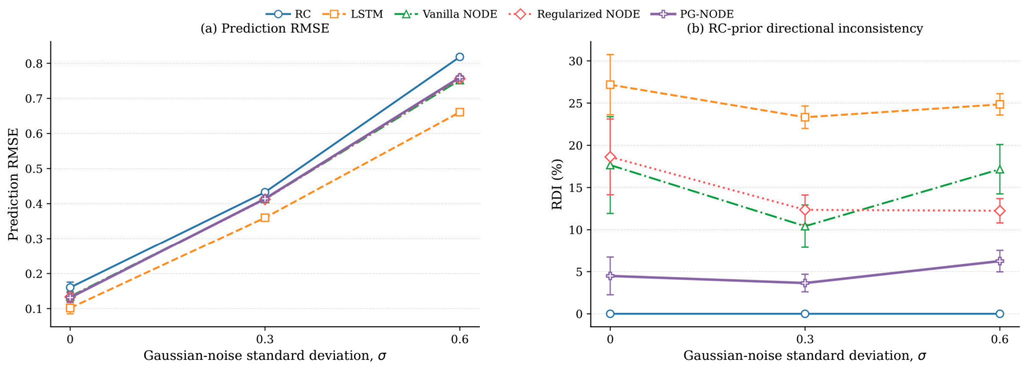

Figure 5. Prediction performance under increasing Gaussian noise levels: (**a**) Prediction RMSE and (**b**) RDI. Markers represent mean values across five thermal zones and five model-training seeds ($n = 25$ runs per model and noise level), and error bars indicate normal approximation 95% confidence intervals of the mean. The RC baseline has zero RDI by construction because it defines the reference direction.

Figure 6 summarizes the validation-only sensitivity of PG-NODE to the directional regularization weight λ_{dir}. The lowest validation RMSE of 0.1287 was obtained at $\lambda_{\mathrm{dir}} = 1$, whereas stronger regularization generally reduced RDI at the expense of prediction accuracy. Increasing λ_{dir} from 0 to 10 reduced the mean RDI from 14.96% to 4.02%, while increasing the mean RMSE from 0.1290 to 0.1361. At $\lambda_{\mathrm{dir}} = 100$, the mean RDI further decreased to 0.60%, whereas the mean RMSE increased to 0.1485. These results indicate a clear accuracy–prior-consistency trade-off rather than a monotonic improvement with increasing regularization strength, consistent with recent physics-guided modeling studies [18–20]. Accordingly, as described in Section 3.2, $\lambda_{\mathrm{dir}} = 10$ was retained as the nominal setting for subsequent experiments rather than as an RMSE-optimal or universally optimal value.

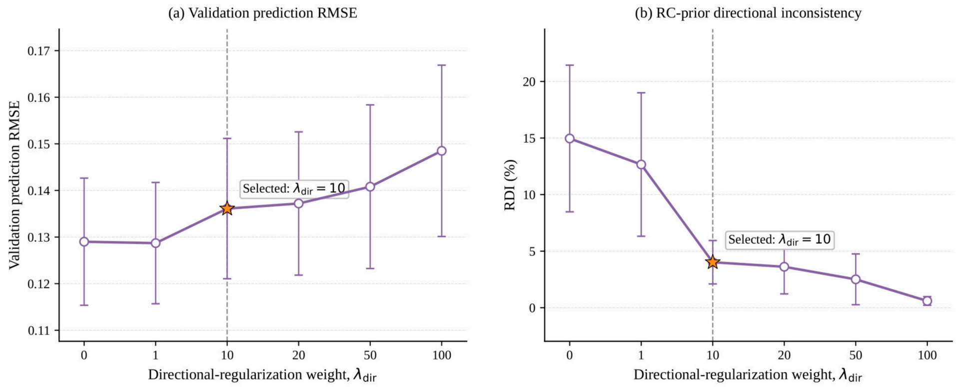

Figure 6. Validation-only sensitivity of PG-NODE to the directional regularization coefficient: (**a**) Validation prediction RMSE and (**b**) RDI. Markers represent mean values across five thermal zones and five model-training seeds ($n = 25$ runs per λ_{dir} setting), and error bars indicate normal approximation 95% confidence intervals of the mean.

7.4. Ablation Analysis of Directional Regularization and Restricted Adaptation

Following the directional regularization sensitivity analysis in Section 7.3, Table 5 isolates the effects of RC-prior directional regularization and restricted final-layer adaptation. Relative to M0, introducing directional regularization in M1 reduced the mean RMSE from 0.1346 to 0.1314 and the mean RDI from 17.66% to 4.49%, corresponding to reductions of approximately 2.4% and 74.6%, respectively. Thus, the directional regularization substantially improved RC-prior consistency while slightly improving prediction accuracy.

Table 5. Ablation results for RC-prior directional regularization and restricted final-layer adaptation.

Variant	RC-Prior Directional Regularization	Restricted Final-Layer Adaptation	RMSE	RDI (%)
M0	No	No	0.1346 ± 0.0349	17.66 ± 14.62
M1	Yes	No	0.1314 ± 0.0360	4.49 ± 5.71
M2	Yes	B = 1	0.1199 ± 0.0363	13.08 ± 3.43
M3	Yes	B = 5	0.1214 ± 0.0355	13.18 ± 3.48

Notes: Values are reported as mean ± standard deviation across five thermal zones and five model-training seeds, with $n = 25$ evaluations per variant. All variants were evaluated using the same noise-free held-out test protocol as the prediction baselines in Table 3. M0 and M1 correspond to the vanilla NODE and fixed PG-NODE baselines, respectively. Detailed M2/M3 adaptation settings are provided in Supplementary Table S2.

Restricted final-layer adaptation in M2 and M3 further reduced the mean RMSE to 0.1199 and 0.1214, respectively, but increased the corresponding mean RDI values to 13.08% and 13.18%. Increasing the transition-buffer length from $B = 1$ to $B = 5$ provided no clear additional benefit. These results indicate an accuracy–prior-consistency trade-off: restricted adaptation improved prediction accuracy relative to the fixed M1 model but weakened agreement with the selected RC prior. M2 and M3 are auxiliary ablation variants and are not part of the main PG-NODE-MPC controller; the corresponding results are visualized in Supplementary Figure S2.

7.5. Multi-Window Closed-Loop Controller Comparison

Table 6 summarizes the closed-loop outcomes for the evaluated BACK SPACE simulation case. After accounting for model-realization-seed and evaluation-window effects using the fixed-block model described in Section 6.4, the analysis supported a lower reference-tracking RMSE for PG-NODE-MPC relative to RBC, with a p-value of 0.019, but did not support a difference in total objective, with a p-value of 0.700. Relative to PG-NODE-MPC, the analysis supported lower reference-tracking RMSE and total objective for RC-EKF-MPC, with both comparisons yielding p-values below 0.001.

Table 6. Multi-window closed-loop performance of the three controllers.

Controller	RMSE	Tariff-Weighted Normalized Control Effort	Total Objective	Control Variation
PG-NODE-MPC	0.247 ± 0.183	32.18 ± 14.29	927.80 ± 1099.34	6.07 ± 4.12
RBC	0.269 ± 0.140	34.35 ± 12.57	940.51 ± 888.21	100.64 ± 92.44
RC-EKF-MPC	0.221 ± 0.194	33.89 ± 14.50	851.62 ± 1072.74	8.74 ± 10.89

Note: Values are descriptive means ± standard deviations across 25 matched seed–window cells, formed by crossing five model-realization seeds with five non-overlapping 48 h test windows. These cells were not treated as independent replicates in the inferential analysis. RMSE denotes the normalized reference-tracking RMSE described in Section 6.2. Tariff-weighted normalized control effort and the total objective are algorithmic comparison metrics under the fixed nominal weighting configuration; they do not represent measured physical energy consumption, monetary cost, or formal occupant thermal-comfort performance.

Figure 7 presents the control trajectories for one representative seed–window realization. The fixed-block analysis of cumulative control variation supported lower cumulative control variation for PG-NODE-MPC than for RBC, with $p < 0.001$. By contrast, the difference between PG-NODE-MPC and RC-EKF-MPC was not statistically significant, with $p = 0.106$.

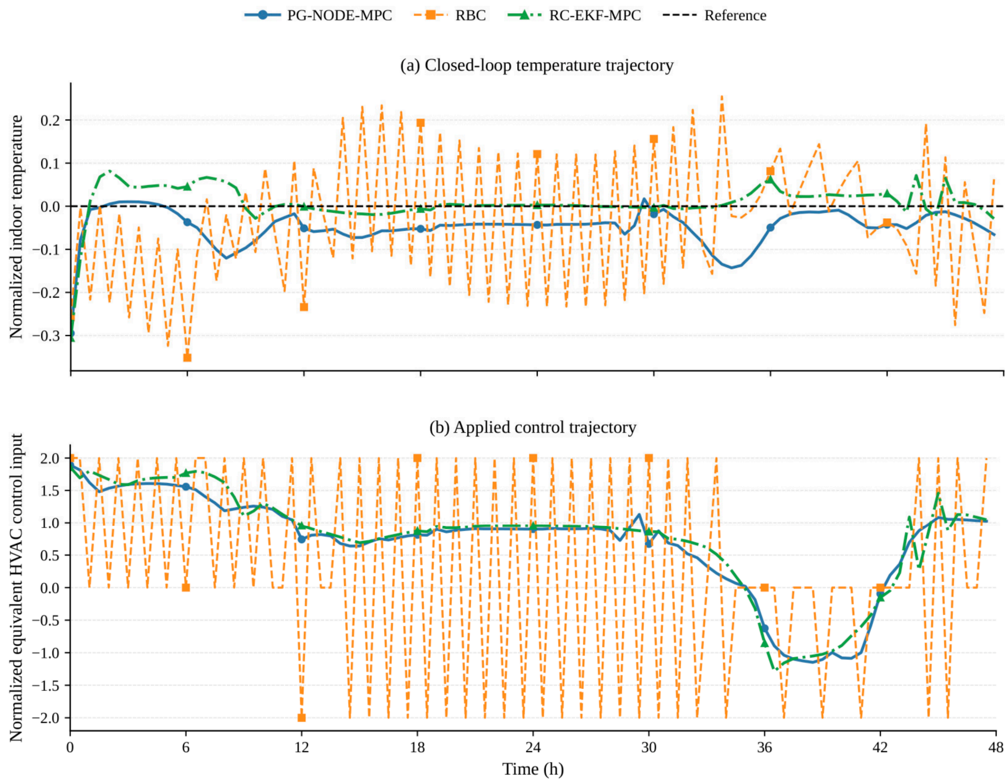

Figure 7. Illustrative closed-loop trajectories from the first predefined 48 h evaluation window for model-realization seed 0: (**a**) Normalized indoor temperature trajectories and (**b**) normalized equivalent HVAC control inputs. The trajectories are presented as a descriptive illustration of one representative seed–window realization.

For the tariff-weighted normalized control effort summarized in Figure 8, the fixed-block analysis supported a lower value for PG-NODE-MPC than for RC-EKF-MPC, with $p < 0.001$, whereas the comparison with RBC did not reach the 0.05 significance level, with $p = 0.070$. Supplementary Table S6 indicates that RC-EKF-MPC achieved the lowest mean absolute deviation from the reference, while the validation-only sensitivity analysis in Supplementary Table S5 shows that the qualitative comparison between the two MPC-based controllers remained stable across the evaluated weight range. Overall, the controllers exhibited distinct algorithmic performance profiles rather than the universal superiority of a single method. Complete fixed-block estimates and 95% confidence intervals for all prespecified comparisons are reported in Supplementary Table S9.

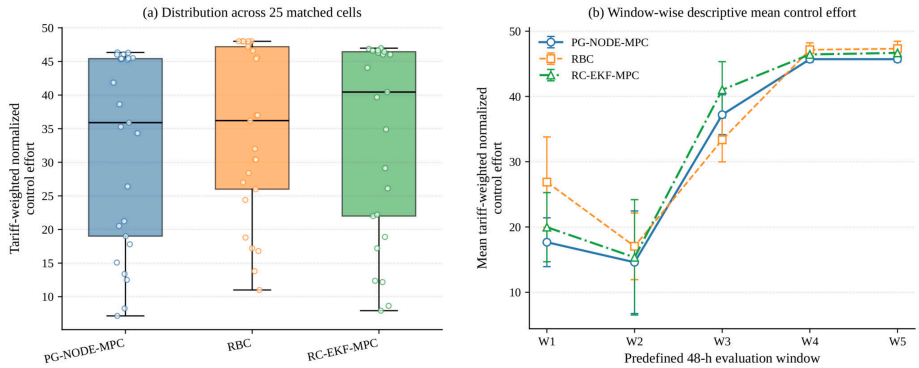

Figure 8. Tariff-weighted normalized control effort in the final multi-window closed-loop evaluation: (**a**) Distributions across the 25 matched seed–window cells and (**b**) descriptive window-wise means across five model-realization seeds. Error bars in panel (**b**) indicate ± 1 standard deviation across the five seeds and are provided for descriptive visualization. PG-NODE-MPC had the lowest descriptive mean tariff-weighted normalized control effort among the three controllers.

7.6. Computational Requirements and Timing Feasibility

As summarized in Table 7, the maximum observed MPC optimization time was 1.117 s, corresponding to approximately 0.062% of the 1800 s control interval, while the maximum process-level resident set size was approximately 307 MB. These results indicate substantial timing margin for the evaluated 30 min supervisory control configuration under the tested environment using only a central processing unit (CPU). Detailed information on the hardware and software environment, runtime measurement, memory accounting, and approximate computational complexity is provided in Supplementary Table S2.

Table 7. Computational requirements of the final PG-NODE-MPC implementation.

Computational Task	Mean ± SD	Maximum
48-step PG-NODE prediction	1.91 ± 0.25 ms	3.26 ms
MPC optimization	860 ± 168 ms	1117 ms
Peak process-level RSS	—	307 MB

Notes: SD denotes standard deviation, and RSS denotes resident set size. Runtime measurements correspond to the final PG-NODE-MPC configuration. Peak process-level RSS represents the maximum resident set size of the complete process running Python 3.13.9 and PyTorch 2.9.1+CPU rather than the standalone memory footprint of the PG-NODE model.

For practical integration, PG-NODE-MPC could operate as a supervisory software layer above existing equipment-level HVAC controllers rather than replacing local control loops. Model identification, retraining, and validation-based checkpoint selection would remain offline, whereas online execution would comprise state-feedback reinitialization, horizon prediction, and receding-horizon optimization. At each 30 min interval, the building management system (BMS) or a commercial HVAC supervisory control platform would provide the latest zone-temperature measurement, outdoor temperature and solar radiation forecasts, and the prescribed time-of-use weighting schedule to an edge computer or supervisory server running the fixed PG-NODE predictor and MPC optimizer. The first optimized action would then be translated through a site-specific interface into admissible setpoint adjustments or other supervisory commands, while existing safety interlocks, operating limits, operator overrides, and fallback logic would remain active. Deployment would additionally require mapping the normalized control coordinate to physical commands and validating communication latency, fault handling, and platform-specific interfaces.

7.7. Limitations and Future Work

Several limitations should be considered when interpreting the present results. First, although the prediction and robustness analyses cover five thermal zones, the main closed-loop comparison is restricted to the BACK SPACE zone of a single DOE prototype retail building under one climate setting. The five 48 h evaluation windows span only late October to late December and therefore do not constitute year-round validation or separate representative heating and cooling season assessments. The closed-loop findings should thus be interpreted as specific to the evaluated simulation case. Extension to other buildings, climates, or strongly coupled multi-zone systems would require renewed model identification or retraining, validation-based deployment-model selection, and, where necessary, an expanded state representation for inter-zone thermal interactions. Beyond these model-update requirements, scaling the framework to many thermal zones or multiple buildings would introduce additional computational and coordination challenges. In a centralized formulation, increasing numbers of thermal states, exogenous inputs, and supervisory control variables would enlarge both the thermal-state representation and the MPC decision problem, increasing training cost, memory demand, and online optimization time. The timing margins reported in Section 7.6 should therefore not be extrapolated directly to larger centralized formulations. Multi-zone operation would also require explicit coordination of inter-zone heat transfer, shared HVAC equipment, coupled operating constraints, and potentially competing zone-level reference-tracking objectives. Multi-building applications would further involve heterogeneous thermal dynamics, sensing availability, local constraints, and communication infrastructure. For sufficiently large systems, hierarchical or distributed MPC may therefore be preferable to a single centralized optimizer, although communication latency, decomposition consistency, convergence, and coordinated fallback operation would require dedicated evaluation.

The statistical generalizability of the closed-loop comparisons is also limited. Although the fixed-block analysis accounts for additive model-realization-seed and evaluation-window effects, its inference remains conditional on the five surrogate-model realizations and five predefined windows considered here. Broader statistical generalization would require additional independently trained surrogate realizations and a wider range of temporal evaluation windows.

Second, because all reported closed-loop controller comparisons relied on independently trained, fixed PG-NODE surrogate environments rather than direct EnergyPlus co-simulation or real-building operation, the resulting controller rankings are conditional on the fidelity and training configuration of those surrogate environments. They should therefore be interpreted as simulation-based algorithmic evidence rather than as closed-loop validation through direct EnergyPlus co-simulation, an experimental testbed, or a real building. The MPC further assumes perfect future outdoor temperature and solar radiation inputs, while occupancy-related and other unmeasured internal disturbances are represented only implicitly rather than forecast separately. In addition, the normalized HVAC input is an equivalent thermal-control coordinate without an explicit mapping to physical actuators, equipment capacity, thermal power, or electrical power. Accordingly, the reported control-effort and objective values are algorithmic comparison metrics rather than measures of physical energy consumption or monetary cost. The temperature metric likewise assesses tracking of the normalized reference state rather than serving as a formal assessment of occupant thermal comfort.

Third, the validation-based checkpoint screening verifies the sign of the local HVAC-response gain only over sampled validation transitions and does not guarantee reliable control gradients under out-of-distribution conditions. The robustness analyses are also limited to predefined Gaussian perturbations and controlled HVAC-input mismatch. Persistent sensor bias, forecast errors, actuator nonlinearities, occupancy variation, equipment faults, long-term model drift, and formal closed-loop stability remain outside the present evaluation scope.

Future validation should therefore proceed in stages. Suitable real-building datasets containing synchronized indoor temperature, weather, internal-load, and HVAC operating variables could first be used to evaluate predictive performance, robustness, and cross-building generalization under measured conditions [39,40]. Because static datasets alone cannot establish closed-loop controller performance without sufficiently resolved control and actuator information, subsequent validation should incorporate explicit actuator and equipment models and proceed through direct co-simulation, hardware-in-the-loop testing, staged BMS integration, field commissioning, and ultimately long-term field evaluation [41,42].

8. Conclusions

This study developed a PG-NODE framework for building thermal prediction and MPC by combining an RC-inspired thermal prior, a neural residual model, and RC-prior directional regularization. Across five thermal zones and five model-training seeds, LSTM achieved the lowest overall prediction RMSE, whereas PG-NODE exhibited substantially lower RDI than the other learned nonlinear models. The main observed benefit of the proposed regularization therefore lies in improving consistency with the selected RC-inspired prior rather than in universally improving prediction accuracy. The ablation study further showed that restricted final-layer adaptation reduced prediction error but increased RDI relative to the fixed PG-NODE, indicating an accuracy–prior-consistency trade-off under the evaluated settings.

Within the evaluated BACK SPACE multi-window closed-loop simulation case, PG-NODE-MPC achieved a lower reference-tracking RMSE than RBC after accounting for model-realization-seed and evaluation-window effects, while also producing substantially smoother control actions. PG-NODE-MPC had the lowest descriptive mean tariff-weighted normalized control effort, with the fixed-block analysis supporting a lower value relative to RC-EKF-MPC but not relative to RBC. RC-EKF-MPC achieved the lowest descriptive mean reference-tracking RMSE and total objective across the three controllers. These findings indicate distinct algorithmic performance profiles rather than universal superiority of any single controller within the evaluated surrogate-based simulation design.

From a practical perspective, the proposed framework offers a structured basis for intelligent-building supervisory decision support rather than a direct equipment-level control solution. By integrating thermal prediction, prior-consistency diagnostics, and receding-horizon optimization, it could support comparative assessment of supervisory actions in terms of reference tracking, control smoothness, and tariff-weighted normalized control effort. The observed timing margin suggests potential compatibility with low-frequency BMS supervisory operation, while existing local controllers, safety interlocks, operator overrides, and fallback logic would remain active. Practical deployment would nevertheless require site-specific command mapping, reliable forecasts and interfaces, and staged validation.

The present evidence remains limited to simulation-based evaluation of one prototype building, one climate, and five 48 h closed-loop windows in the BACK SPACE zone. As discussed in Section 7.7, future work should extend validation to measured real-building data and broader building and climate conditions, incorporate imperfect forecasts and explicit actuator and equipment models, and evaluate scalable multi-zone and multi-building formulations. Such extensions should address inter-zone thermal coupling, coordination of shared HVAC resources and operating constraints, and the potential need for hierarchical or distributed optimization, while deployment validation should progress through co-simulation, hardware-in-the-loop testing, staged BMS integration, and field deployment.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/buildings16153021/s1, Table S1: Validation-only sensitivity of PG-NODE to the RC-prior directional regularization weight λ_{dir}; Table S2: Complete reproducibility settings for data preprocessing, model development, surrogate-environment training, robustness protocols, controller implementation, closed-loop evaluation, fixed-block statistical analysis, and runtime benchmarking; Table S3: Validation-only sensitivity of the RC-EKF covariance settings; Table S4: Post hoc tariff-reweighting analysis under alternative peak/off-peak coefficient settings; Table S5: Validation-only sensitivity of the MPC controllers to the MPC objective weights; Table S6: Additional temperature-scale reference-tracking diagnostics for the final multi-window closed-loop evaluation; Table S7: Robustness of fixed predictors under controlled HVAC-input mismatch; Table S8: Zone-wise prediction performance under the noise-free held-out test condition; Table S9: Two-way fixed-block analyses of matched closed-loop controller differences across five model-realization seeds and five evaluation windows; Figure S1: Distribution of prediction performance across five thermal zones and five model-training seeds; Figure S2: Ablation accuracy–prior-consistency trade-off.

Author Contributions: Conceptualization, Y.W. and E.Z.; methodology, L.C. and E.Z.; software, L.C. and K.J.; validation, K.J. and X.Z.; formal analysis, L.C. and K.J.; investigation, Y.W.; resources, Y.W. and Y.L.; data curation, L.C. and X.Z.; writing—original draft preparation, L.C.; writing—review and editing, Y.W., E.Z., K.J., X.Z. and Y.L.; visualization, L.C.; supervision, E.Z.; project administration, E.Z. and Y.L.; funding acquisition, Y.W. and Y.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Gansu Provincial Key Laboratory of Solar Energy Utilization under its 2025 Open Project, grant number GSKLSEU2025-ZD01.

Data Availability Statement: The data and code supporting this study are available from the corresponding author upon reasonable request. Summary results and reproducibility settings are provided in the article and Supplementary Materials.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (GPT-5.6 Thinking, OpenAI; accessed on July 2026) to assist with the translation of selected text from Chinese into English and to improve the clarity, readability, and grammatical accuracy of the English text. The authors reviewed, revised, and verified all AI-assisted output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Nomenclature

Symbol	Definition
$h(t), \widehat{h}(t)$	Normalized indoor thermal state and its model-predicted counterpart
T_k^{obs}	Observed normalized indoor temperature at control step k
T_{out}	Outdoor temperature input
S	Solar radiation input
u_k	Normalized equivalent HVAC control input at control step k
$\mathbf{d}_k$	Exogenous-input vector containing outdoor temperature and solar radiation
$f_{\mathrm{RC}}(\cdot)$	RC-inspired reference component of the PG-NODE model
$f_{NN}(\cdot)$	Trainable neural residual component of the PG-NODE model
$\boldsymbol{\theta}_{\mathrm{RC}}$	Parameter vector of the RC-inspired reference component
$\boldsymbol{\theta}_{NN}$	Parameter vector of the neural residual component
$g_{\mathrm{RC},k}$	RC-prior local temperature change trend evaluated at step k
$\widehat{g}_i$	Local temperature change trend predicted by PG-NODE for sample i

λ_{dir}	Weight of the RC-prior directional regularization term
ϵ	Activity threshold used in RC-prior directional regularization
γ	HVAC-input gain factor used in the controlled model-mismatch protocol
m	Directional margin in the RC-prior regularization term
$\Delta\tau$	Numerical integration step; $\Delta\tau$= 1 corresponds to one 30-min sampling interval
N	MPC prediction horizon in control steps
λ_c	Reference-tracking penalty weight in the MPC objective
$\lambda_{\Delta u}$	Control-smoothness penalty weight in the MPC objective
c_t	Normalized time-of-use weighting coefficient
h_{ref}	Normalized indoor temperature reference state
TV_u	Cumulative variation of the applied control input
σ	Standard deviation of the Gaussian perturbation in normalized coordinates
η_t	Independent Gaussian perturbation applied at time step t
PG-NODE	Physics-guided neural ordinary differential equation
MPC	Model predictive control
RC	Resistance–capacitance
EKF	Extended Kalman filter
RBC	Rule-based controller
LSTM	Long short-term memory
RMSE	Root mean square error; used for prediction accuracy or closed-loop reference tracking as specified in the corresponding evaluation context
RDI	RC-prior directional inconsistency rate

References

1. Liang, Y.; Li, C.; Liu, Z.; Wang, X.; Zeng, F.; Yuan, X.; Pan, Y. Decarbonization Potentials of the Embodied Energy Use and Operational Process in Buildings: A Review from the Life-Cycle Perspective. *Heliyon* **2023**, *9*, e20190. [CrossRef] [PubMed]
2. Hu, S.; Zhou, X.; Yan, D.; Guo, F.; Hong, T.; Jiang, Y. A Systematic Review of Building Energy Sufficiency towards Energy and Climate Targets. *Renew. Sustain. Energy Rev.* **2023**, *181*, 113316. [CrossRef]
3. Drgoňa, J.; Arroyo, J.; Cupeiro Figueroa, I.; Blum, D.; Arendt, K.; Kim, D.; Ollé, E.P.; Oravec, J.; Wetter, M.; Vrabie, D.L.; et al. All You Need to Know about Model Predictive Control for Buildings. *Annu. Rev. Control* **2020**, *50*, 190–232. [CrossRef]
4. Stoffel, P.; Berktold, M.; Müller, D. Real-Life Data-Driven Model Predictive Control for Building Energy Systems Comparing Different Machine Learning Models. *Energy Build.* **2024**, *305*, 113895. [CrossRef]
5. Saloux, E.; Candanedo, J.A.; Vallianos, C.; Morovat, N.; Zhang, K. From Theory to Practice: A Critical Review of Model Predictive Control Field Implementations in the Built Environment. *Appl. Energy* **2025**, *393*, 126091. [CrossRef]
6. Khabbazi, A.J.; Pergantis, E.N.; Reyes Premer, L.D.; Papageorgiou, P.; Lee, A.H.; Braun, J.E.; Henze, G.P.; Kircher, K.J. Lessons Learned from Field Demonstrations of Model Predictive Control and Reinforcement Learning for Residential and Commercial HVAC: A Review. *Appl. Energy* **2025**, *399*, 126459. [CrossRef]
7. Bacher, P.; Madsen, H. Identifying Suitable Models for the Heat Dynamics of Buildings. *Energy Build.* **2011**, *43*, 1511–1522. [CrossRef]
8. Wang, Z.; Chen, Y. Data-Driven Modeling of Building Thermal Dynamics: Methodology and State of the Art. *Energy Build.* **2019**, *203*, 109405. [CrossRef]
9. Jiang, Z.; Wang, X.; Li, H.; Hong, T.; You, F.; Drgoňa, J.; Vrabie, D.L.; Dong, B. Physics-Informed Machine Learning for Building Performance Simulation—A Review of a Nascent Field. *Adv. Appl. Energy* **2025**, *18*, 100223. [CrossRef]
10. Ma, Z.; Jiang, G.; Hu, Y.; Chen, J. A Review of Physics-Informed Machine Learning for Building Energy Modeling. *Appl. Energy* **2025**, *381*, 125169. [CrossRef]
11. Drgoňa, J.; Tuor, A.R.; Chandan, V.; Vrabie, D.L. Physics-Constrained Deep Learning of Multi-Zone Building Thermal Dynamics. *Energy Build.* **2021**, *243*, 110992. [CrossRef]
12. Gokhale, G.; Claessens, B.; Develder, C. Physics Informed Neural Networks for Control Oriented Thermal Modeling of Buildings. *Appl. Energy* **2022**, *314*, 118852. [CrossRef]
13. Xiao, T.; You, F. Building Thermal Modeling and Model Predictive Control with Physically Consistent Deep Learning for Decarbonization and Energy Optimization. *Appl. Energy* **2023**, *342*, 121165. [CrossRef]

14. Chen, Y.; Yang, Q.; Chen, Z.; Yan, C.; Zeng, S.; Dai, M. Physics-Informed Neural Networks for Building Thermal Modeling and Demand Response Control. *Build. Environ.* **2023**, *234*, 110149. [CrossRef]
15. Yang, Z.; Gaidhane, A.D.; Drgoňa, J.; Chandan, V.; Halappanavar, M.M.; Liu, F.; Cao, Y. Physics-Constrained Graph Modeling for Building Thermal Dynamics. *Energy AI* **2024**, *16*, 100346. [CrossRef]
16. Ma, Z.; Jiang, G.; Chen, J. Physics-Informed Ensemble Learning with Residual Modeling for Enhanced Building Energy Prediction. *Energy Build.* **2024**, *323*, 114853. [CrossRef]
17. Tang, L.; Xie, H.; Wang, Y.; Xu, Z. Deeply Flexible Commercial Building HVAC System Control: A Physics-Aware Deep Learning-Embedded MPC Approach. *Appl. Energy* **2025**, *388*, 125631. [CrossRef]
18. Chen, Y.; Wang, H.; Chen, Z. Sensitivity Analysis of Physical Regularization in Physics-Informed Neural Networks (PINNs) of Building Thermal Modeling. *Build. Environ.* **2025**, *273*, 112693. [CrossRef]
19. Jiang, Z.; Wang, X.; Dong, B. Physics-Informed Modularized Neural Network for Advanced Building Control by Deep Reinforcement Learning. *Adv. Appl. Energy* **2025**, *19*, 100237. [CrossRef]
20. Wang, X.; Jiang, Z.; Zhan, S.; Liang, W.; Dong, B. Dual-Phase Evaluation Framework of Physics-Informed Machine Learning Model for Building Dynamics Modeling and Advanced Control. *Energy Build.* **2026**, *352*, 116791. [CrossRef]
21. Chen, R.T.Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D.K. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems 31*; Curran Associates, Inc.: Red Hook, NY, USA, 2018; pp. 6571–6583.
22. He, B.; Zhang, N.; Fang, C.; Su, Y.; Wang, Y. Flexible Building Energy Management with Neural ODEs-Based Model Predictive Control. *IEEE Trans. Smart Grid* **2024**, *15*, 4690–4704. [CrossRef]
23. Taboga, V.; Gehring, C.; Le Cam, M.; Dagdougui, H.; Bacon, P.-L. Neural Differential Equations for Temperature Control in Buildings under Demand Response Programs. *Appl. Energy* **2024**, *368*, 123433. [CrossRef]
24. Sabbagh, G.; Cimmino, M.; Delcroix, B. Physics-Informed Neural Ordinary Differential Equations for Multi-Zone Residential Thermal Modeling. *Energy Build.* **2026**, *350*, 116623. [CrossRef]
25. Wang, D.; Hu, W.; Wang, W. Physics-Informed Neural Networks with Continuous-Time Thermal Modelling for Model Predictive Control of Building Energy Systems in Demand Response. *Energy Build.* **2026**, *352*, 116864. [CrossRef]
26. Klanatsky, P.; Veynandt, F.; Heschl, C.; Stelzer, R.; Zogas, P.; Siokas, G.; Balomenos, A. Real Long-Term Performance Evaluation of an Improved Office Building Operation Involving a Data-Driven Model Predictive Control. *Energy Build.* **2025**, *338*, 115590. [CrossRef]
27. Zanetti, E.; Blum, D.; Fu, H.; Weyandt, C.; Pritoni, M.; Piette, M.A. Commercial Building HVAC Demand Flexibility with Model Predictive Control: Field Demonstration and Literature Insights. *Energy Build.* **2025**, *345*, 116097. [CrossRef]
28. Mork, M.; Redder, F.; Xhonneux, A.; Müller, D. Real-World Implementation and Evaluation of a Model Predictive Control Framework in an Office Space. *J. Build. Eng.* **2023**, *78*, 107619. [CrossRef]
29. Blum, D.; Wang, Z.; Weyandt, C.; Kim, D.; Wetter, M.; Hong, T.; Piette, M.A. Field Demonstration and Implementation Analysis of Model Predictive Control in an Office HVAC System. *Appl. Energy* **2022**, *318*, 119104. [CrossRef]
30. Hannula, E.; Häkkinen, A.; Uribe, F.; Solonen, A.; de Wiljes, J.; Roininen, L. Partially Stochastic Deep Learning with Uncertainty Quantification for Building Thermal Modeling. *Build. Environ.* **2026**, *289*, 114118. [CrossRef]
31. Stoffel, P.; Henkel, P.; Rätz, M.; Kümpel, A.; Müller, D. Safe Operation of Online Learning Data Driven Model Predictive Control of Building Energy Systems. *Energy AI* **2023**, *14*, 100296. [CrossRef]
32. Massano, M.; Patti, E.; Macii, E.; Acquaviva, A.; Bottaccioli, L. An Online Grey-Box Model Based on Unscented Kalman Filter to Predict Temperature Profiles in Smart Buildings. *Energies* **2020**, *13*, 2097. [CrossRef]
33. Zeng, T.; Barooah, P. An Adaptive Model Predictive Control Scheme for Energy-Efficient Control of Building HVAC Systems. *J. Eng. Sustain. Build. Cities* **2021**, *2*, 031001. [CrossRef]
34. Crawley, D.B.; Lawrie, L.K.; Winkelmann, F.C.; Buhl, W.F.; Huang, Y.J.; Pedersen, C.O.; Strand, R.K.; Liesen, R.J.; Fisher, D.E.; Witte, M.J.; et al. EnergyPlus: Creating a New-Generation Building Energy Simulation Program. *Energy Build.* **2001**, *33*, 319–331. [CrossRef]
35. U.S. Department of Energy. Prototype Building Models. Available online: https://www.energycodes.gov/prototype-building-models (accessed on 20 July 2026).
36. *ANSI/ASHRAE/IES Standard 90.1-2022*; Energy Standard for Sites and Buildings Except Low-Rise Residential Buildings. ASHRAE: Atlanta, GA, USA, 2022.
37. Wilcox, S.; Marion, W. *Users Manual for TMY3 Data Sets (Revised)*; NREL/TP-581-43156; National Renewable Energy Laboratory: Golden, CO, USA, 2008.
38. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]
39. Zhang, K.; Prakash, A.; Paul, L.; Blum, D.; Alstone, P.; Zoellick, J.; Brown, R.; Pritoni, M. Model Predictive Control for Demand Flexibility: Real-World Operation of a Commercial Building with Photovoltaic and Battery Systems. *Adv. Appl. Energy* **2022**, *7*, 100099. [CrossRef]

40. Wang, D.; Chen, Y.; Wang, W.; Gao, C.; Wang, Z. Field Test of Model Predictive Control in Residential Buildings for Utility Cost Savings. *Energy Build.* **2023**, *288*, 113026. [CrossRef]
41. Bird, M.; Daveau, C.; O'Dwyer, E.; Acha, S.; Shah, N. Real-World Implementation and Cost of a Cloud-Based MPC Retrofit for HVAC Control Systems in Commercial Buildings. *Energy Build.* **2022**, *270*, 112269. [CrossRef]
42. Morovat, N.; Athienitis, A.K.; Candanedo, J.A.; Nouanegue, H.F. Field Implementation of Model-Based Predictive Control in an All-Electric School Building: Impact of Occupancy on Energy Flexibility. *Energy* **2025**, *329*, 136852. [CrossRef]