

Article

Robust Closed-Loop Control of Industrial Systems Based on Cloud–Edge Device Collaboration

Wenjing Zhang ¹, Wenchao Zhang ^{2,*}, Xiao Ma ^{1,*}, Weijia Han ¹, Liang Wang ³ and Minghang Chen ¹

¹ School of Physics and Information Technology, Shaanxi Normal University, Xi'an 710119, China; wenjing-zhang@snnu.edu.cn (W.Z.); wjhan@snnu.edu.cn (W.H.); minghangchen@snnu.edu.cn (M.C.)

² Department of Community Management and Service, Shaanxi Normal University, Xi'an 710119, China

³ School of Artificial Intelligence and Computer Science, Shaanxi Normal University, Xi'an 710119, China; wangliang@snnu.edu.cn

* Correspondence: wenchao@snnu.edu.cn (W.Z.); xma@snnu.edu.cn (X.M.)

Abstract

Under China's dual-carbon strategic goals, large-scale public centralized heating plays a critical role in energy conservation. However, traditional manual open-loop control suffers from high response latency. Furthermore, existing unidirectional predictive methods lack dynamic feedback correction mechanisms. To address these issues, this study proposes an intelligent dual-system collaborative control architecture specifically designed for public heating systems. This architecture utilizes a cloud–edge device framework. It establishes a two-way linkage between heating equipment and the cloud decision platform. Consequently, it constructs an integrated regulation framework encompassing forward decision generation, reverse state verification, and dynamic feedback correction. Specifically, the forward module utilizes a Mamba-structured state-space model to generate data-driven boiler operation strategies. Meanwhile, the reverse module employs a Temporal Convolutional Network with Monte Carlo Dropout (TCN-MC Dropout). This probabilistic network enables state inversion evaluation with reliable uncertainty prediction intervals. These two modules are deeply coupled through an adaptive feedback correction mechanism. Together, they significantly improve system stability and operational robustness under complex thermal disturbances. Specifically, the proposed architecture achieves a room temperature compliance rate exceeding 96% and restricts temperature fluctuations to within ± 0.75 °C. Simultaneously, it reduces boiler energy consumption by over 16.4%. This solution has been successfully deployed in the heating network at Shaanxi Normal University as a representative real-world case study. Ultimately, it provides a practical technical reference for the intelligent upgrading and low-carbon transformation of public centralized heating systems.

Keywords: industrial control; cloud–edge device; smart heating; robustness; closed-loop control; Mamba; uncertainty quantification; TCN-MC Dropout



Academic Editor: Stefano Scanzio

Received: 7 July 2026

Revised: 6 September 2026

Accepted: 9 September 2026

Published: 11 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Under the dual-carbon goals, intelligently upgrading public heating systems is crucial, serving as a key breakthrough for reducing energy consumption and driving the green transformation of heating infrastructures. Centralized public heating represents a typical large-scale thermal regulation system. It exhibits inherent challenges such as strong thermal inertia, long time delays, and complex external disturbances [1,2]. The system's regulation efficiency and dynamic tuning accuracy directly determine energy utilization and user

comfort, reflecting the overall intelligence and refined management level of modern public heating services.

However, in typical public heating scenarios, traditional control systems still rely heavily on manual experience and open-loop modes [3–5]. Consequently, these operational practices suffer from delayed multi-region coordination and lack data-driven closed-loop optimization mechanisms, which usually introduces the following critical problems:

- 1 Insufficient overall coordination. Manual control struggles to achieve real-time load balancing across multiple heating zones. This causes supply–demand imbalances, leading to uneven indoor temperatures.
- 2 Significant control delays. Thermal systems possess strong thermal inertia. Manual control cannot predict sudden external weather disturbances, leading to large room temperature fluctuations and poor thermal comfort.
- 3 Passive and inefficient maintenance. Relying on manual inspections results in long response times for heating equipment faults. The lack of digital analysis prevents early anomaly prediction, increasing operational costs and safety risks.

Therefore, there is an urgent need to build a new intelligent regulation architecture featuring data-driven decision-making, autonomous perception, and adaptive optimization for public heating [6,7]. To address these issues, this paper proposes a dual-system collaborative closed-loop control scheme tailored for public heating systems. This framework establishes an intelligent thermal regulation mechanism adaptable to complex heating environments. Built upon a cloud–edge device collaborative architecture, the scheme integrates forward decision-making, reverse verification, and feedback correction. Specifically, the forward module uses a Mamba-structured state-space model to rapidly generate heating control strategies. Meanwhile, the reverse module adopts a Temporal Convolutional Network with Monte Carlo Dropout (TCN-MC Dropout) architecture for strategy pre-verification and risk assessment. Through probabilistic prediction, it enables dynamic monitoring and real-time correction of the regulated thermal parameters.

Actual operational data validate the excellent performance of this solution. It increases the room temperature compliance rate to over 96% while keeping temperature fluctuations within ± 0.75 °C. Concurrently, boiler energy consumption is reduced by more than 16.4%. Furthermore, this solution has been successfully deployed at Shaanxi Normal University. This practical deployment provides a practical and scalable model for the intelligent and green transformation of public heating systems.

The remainder of this paper is organized as follows: Section 2 reviews related work on public heating systems and intelligent regulation architectures, identifying the gaps that this study addresses. Section 3 introduces the system model and problem formulation, defining the core challenges of centralized thermal control. Section 4 details the proposed dual-system collaborative closed-loop control scheme, including forward Mamba decision-making and reverse TCN-MC Dropout verification. Section 5 presents the experimental setup, comparative results against baseline methods, and comprehensive evaluations of inference speed and energy efficiency. Finally, Section 6 concludes with the key findings, limitations, and directions for future work.

2. Related Work

Recently, remarkable advancements in time-series forecasting and cloud–edge device architectures have laid a solid technological foundation for intelligent closed-loop control in public heating systems. Driven by these developments, data-driven predictive modeling has attracted growing research interest in the field of thermal regulation.

Early studies primarily adopted Recurrent Neural Networks (RNNs) for multi-step thermal load forecasting. For instance, Powell et al. [8] demonstrated the effectiveness of

RNNs in predicting 24 h-ahead heating dynamics. However, large-scale thermal systems are characterized by long time delays and long sequence dynamics. In these scenarios, RNNs frequently suffer from gradient vanishing or exploding phenomena alongside low serial computation efficiency. To mitigate these structural defects, Hochreiter et al. [9] introduced Long Short-Term Memory (LSTM) networks, which excel at capturing long-term dependencies. Lee et al. [10] further verified the superiority of RNN and LSTM variants over conventional Artificial Neural Networks (ANNs) and Support Vector Regression (SVR).

Meanwhile, recent attention has shifted toward Transformer-based architectures due to their extraordinary capability in modeling long-range dependencies [11]. Nevertheless, the quadratic computational complexity ($\mathcal{O}(N^2)$) of the Transformer mechanism imposes prohibitive computational overheads for resource-constrained edge devices. This results in catastrophic inference latency that fails to meet the real-time demands of public heating control. In order to overcome these computational barriers, state-space models (SSMs) have recently emerged as a breakthrough technology. They provide an efficient alternative solution for Transformers. Specifically, the Mamba architecture introduces a selective state-space mechanism. It has achieved excellent long sequence modeling capabilities. Importantly, it maintains linear computational complexity. This breakthrough makes SSMs very suitable for real-time control on edge devices.

Additionally, cloud–edge architectures enable efficient data transmission, distributed processing, and remote centralized control [12,13]. This hierarchical approach combines edge perception with cloud analysis. It fits the demands of large centralized heating systems well. However, existing research shows clear limitations. Most solutions rely on open-loop control paradigms without feedback. A typical example is the traditional machine learning heat load forecasting system designed by Idowu et al. [14]. Recent studies have explored closed-loop control in other fields. For example, Ali et al. designed adaptive control for voltage regulation [15], Khan et al. applied vision control in industrial tasks [16], and Jiang et al. developed dual-feedback control for air conditioning [17]. Nevertheless, these methods lack real-time strategy verification and dynamic correction for heating dynamics. Consequently, they cannot provide a complete closed loop for heating regulation.

These cloud–edge schemes effectively resolve data transmission and response latency issues. However, they struggle to maintain operational stability when facing sudden temperature drops or severe building load fluctuations. This limitation aligns with the robustness defects of traditional schemes verified by Hansen et al. [18]. Ultimately, this severely compromises overall system robustness. While some studies attempt to introduce feedback mechanisms, they mostly rely on simple Proportional–Integral (PI) control. These conventional approaches cannot quantify prediction uncertainties or evaluate control risks. Therefore, they struggle to adapt to complex public heating scenarios. Such scenarios are typically characterized by strong disturbances, multi-regional differences, and frequent condition switching. In order to provide clear classification, Table 1 provides a multidimensional summary and comparison of the above prediction schemes.

To address the aforementioned limitations, this paper proposes a novel dual-system collaborative closed-loop control architecture. A core methodological innovation of this work is the selection of the Mamba model. Mamba is a cutting-edge SSM. It serves as the backbone for the forward decision-making module. Compared to traditional RNNs and Transformers, Mamba exhibits unique structural advantages for energy system optimization. It utilizes a selective state-space mechanism and hardware-aware algorithms. These features allow Mamba to capture long-range dependencies effectively while maintaining linear computational complexity ($\mathcal{O}(N)$). This capability perfectly overcomes the quadratic computational bottleneck of Transformers. Consequently, the model efficiently processes

the long time delays and strong thermal inertia inherent in large-scale heating systems. It achieves this without incurring a prohibitive computational overhead. Thus, it guarantees real-time and efficient sequence modeling on resource-constrained edge nodes.

Furthermore, we integrate TCN-MC Dropout into the reverse module. This integration establishes a complete and robust closed loop. TCNs offer distinct advantages for time-series verification. They utilize dilated causal convolutions. This structure exponentially expands the receptive field to capture complex temporal patterns. Unlike traditional RNNs, TCNs support highly parallel computations. This feature significantly accelerates inference speeds on edge devices. Additionally, we couple the TCN architecture with MC Dropout. This combination enables reliable uncertainty quantification. Conventional PI or open-loop schemes often struggle with complex external thermal disturbances. Our proposed design overcomes these limitations by dynamically assessing potential control risks. Ultimately, this mechanism ensures robust strategy pre-verification and safe closed-loop correction.

Table 1. Summary of the schemes mentioned previously.

Schemes	Core Methodologies	Computational Complexities	Control Paradigms	Edge Deployment	Uncertainty Quantification	Robustness
Powell et al. [8]	RNN	$O(N)$	Open loop	Low	×	Poor
Hochreiter et al. [9]	LSTM	$O(N)$	Open loop	Medium	×	Medium
Lee et al. [10]	RNN/LSTM Variants	$O(N)$	Open loop	Medium	×	Medium
Tan et al. [11]	Transformer	$O(N^2)$	Open loop	Poor	×	Medium
Idowu et al. [14]	Traditional ML heat load forecasting	$O(N)$	Open loop	Medium	×	Medium
Ali et al. [15]	DRL-PID adaptive control (TD3)	$O(N)$	Closed loop	High	×	Strong
Khan et al. [16]	Real-time closed-loop control	$O(N)$	Closed loop	High	×	Strong
Jiang et al. [17]	Cloud-edge cooperative dual-feedback closed-loop control	$O(N)$	Closed loop	High	×	Strong
This article	Mamba + TCN-MC Dropout	$O(N)$	Closed loop	High	✓	Strong

Abbreviations: RNN: Recurrent Neural Network; LSTM: Long Short-Term Memory; ML: Machine Learning; DRL: Deep Reinforcement Learning; PID: Proportional-Integral-Derivative; TD3: Twin Delayed Deep Deterministic Policy Gradient; TCN-MC Dropout: Temporal Convolutional Network with Monte Carlo Dropout.

3. System Model and Problem Formulation

Large-scale public centralized heating is characterized by multivariable coupling, strong inertia, large time delays, and significant external disturbances [19]. The operational process spans multiple interconnected stages. These include heat generation, transmission, storage, and dissipation. Concurrently, the system must satisfy the heating demands of diverse regions. As shown in Figure 1, the centralized heating system at Shaanxi Normal University precisely exemplifies this scenario. Although it achieves basic intelligent scheduling, it still faces critical challenges. These include differentiated regional loads and insufficient fine-grained control. Therefore, it provides an ideal environment for researching and verifying the intelligent upgrading of public heating systems.




Figure 1. Schematic diagram of centralized heating system in large public buildings.

Based on a cloud–edge device architecture, this paper addresses the regulation challenges of public centralized heating systems. The selected campus network features complex functional zoning and significant thermal load differences, which represent the complexities of large-scale public heating networks. This deployment provides real-world operating conditions and solid engineering verification for the proposed scheme. From a regulation perspective, intelligent public heating presents a complex dynamic modeling problem that requires bidirectional correlation between forward decision-making and reverse verification, mutually constraining and collaborating to achieve precise closed-loop regulation and enhance system robustness [20].

Forward decision modeling establishes a causal mapping from environmental factors and demands to boiler operating parameters. Essentially, it predicts the predicted control output given external environmental conditions and target room temperatures. This approach provides the theoretical foundation for generating the initial control strategy. As the system’s core heat source, the boiler’s output power directly determines the heating intensity of each region. Based on the principle of heat balance, this thermal output power is expressed as follows:

$$Q_b(t) = c_w \cdot m_w \cdot (T_s(t) - T_r(t)) \cdot \eta_b, \quad (1)$$

where $Q_b(t)$ is the thermal output power of the boiler at time t , c_w is the specific heat capacity of water, m_w is the circulating water flow rate, $T_s(t)$ is the boiler supply water temperature at time t , $T_r(t)$ is the boiler return water temperature at time t , and η_b is the thermal efficiency of the boiler.

The proposed heat output model is based on general physical principles. It can be adapted to different boilers after parameter calibration. Specific engineering parameters include the thermal efficiency, design heat load, and temperature–load relationship. These parameters, along with actual operational data from Shaanxi Normal University, will be detailed in the experimental chapter.

Based on this model, forward decision-making integrates pipeline heat loss and building thermal dynamics. It establishes a dynamic mapping from historical data, weather forecasts, and target regional temperatures to boiler operating parameters. This process enables the rapid generation of data-driven control strategies.

Meanwhile, reverse room temperature verification complements forward decision modeling. Its core objective is to construct a transfer model from boiler operating param-

ters to regional temperatures. Essentially, this process represents a dynamic response that predicts system outputs based on known control inputs. This verification provides critical theoretical support for initial control strategies while enhancing overall system robustness.

Heating systems exhibit strong thermal inertia and nonlinear characteristics. Consequently, predictions from a single forward model are highly susceptible to disturbances and prone to deviations. This vulnerability necessitates verification and correction via an inverse model. The core logic of this inversion is to predict the expected regional room temperature $T_{\text{pred},i}(t)$ after strategy implementation. The prediction integrates the forward model outputs, specifically, the boiler supply water temperature $T_s(t)$ and the number of operating boilers $N_b(t)$, with external factors like the outdoor temperature $T_{\text{out}}(t)$ and building thermal inertia. The central constraint relationship is expressed as follows:

$$\left| T_{\text{pred},i}(t) - T_{\text{set},i} \right| \leq \Delta T_{\text{max}}, \quad (2)$$

where ΔT_{max} is the allowable room temperature fluctuation range, which is set to ± 1 °C according to the actual heating demand. This constraint relationship directly determines the qualification of the forward decision-making strategy.

Therefore, reverse room temperature verification must both accurately predict the expected room temperature and quantify prediction uncertainty. By calculating the room temperature prediction variance $\sigma^2(t)$, the system evaluates the reliability of the current strategy. This evaluation prevents regulation deviations caused by external disturbances. Furthermore, it provides a quantitative basis for correcting forward decisions, thereby enhancing system robustness.

By integrating forward decision-making and reverse verification, this paper constructs a dual-system collaborative closed-loop framework. A cloud–edge device architecture serves as the core infrastructure for this system. The primary objective is to minimize energy consumption while satisfying regional thermal comfort constraints. Concurrently, the framework aims to enhance system robustness under strong disturbance conditions. The overall optimization objective is formulated as follows:

$$\min J = \lambda_1 \cdot G(t) + \lambda_2 \cdot \sum_{i=1}^3 |T_{\text{in},i}(t) - T_{\text{set},i}| + \lambda_3 \cdot \sigma^2(t), \quad (3)$$

where J represents the comprehensive optimization objective function. A smaller value of J indicates superior system performance across three dimensions: energy consumption, room temperature comfort, and regulation robustness. Additionally, $G(t)$ denotes the hourly natural gas consumption of the boiler. The term $T_{\text{in},i}(t)$ is the actual indoor temperature of the i -th region at time step t . $T_{\text{set},i}$ is its corresponding target setpoint. Finally, $\sigma^2(t)$ represents the prediction variance from the reverse verification. This variance quantifies uncertainty and characterizes system robustness.

To eliminate optimization bias caused by varying dimensional scales, all sub-items in the objective function undergo Min–Max normalization. This process strictly maps all variables to the $[0, 1]$ interval. Furthermore, the weight coefficients λ_1 , λ_2 , and λ_3 are not assigned arbitrarily. Their selection is grounded in historical operational data and practical heating regulation priorities. Specifically, indoor thermal comfort (λ_2) takes the highest priority. Energy efficiency (λ_1) serves as the secondary focus. System robustness (λ_3) acts as an auxiliary adjustment. Based on these principles, the initial baseline weights are set to $\lambda_1 = 0.3$, $\lambda_2 = 0.5$, and $\lambda_3 = 0.2$. Operators can dynamically adjust these hyperparameters. This flexibility allows the system to shift its optimization focus according to specific real-world scenario requirements.

It should be noted that the Min–Max normalization described above is exclusively applied to balance the dimensional scales of the sub-items during the offline simulation and optimization phase. However, for the actual field deployment evaluation, the comprehensive index J is directly computed using the raw, unnormalized physical quantities multiplied by the same weight coefficients.

4. The Proposed Scheme

4.1. Construction of the Cloud–Edge Device Architecture

The cloud–edge device collaborative architecture supports hierarchical control and closed-loop linkage in intelligent heating systems. Large-scale heating systems involve numerous devices, wide spatial distribution, and complex couplings [21]. To address these challenges, we construct a hierarchical framework that enables bidirectional interaction. It integrates multi-source data acquisition at the device end, on-site preprocessing at the edge, and global optimization in the cloud. This architecture ensures real-time local control while supporting global strategy optimization, serving as a stable technical carrier for the dual-system collaborative decision-making process. Multi-source data acquisition and hierarchical processing build the perception foundation, converting physical states into standardized digital inputs for intelligent models [22].

As shown in Figure 2, our framework deploys a comprehensive sensing system across three dimensions: the environment, equipment, and users. The environmental unit collects real-time weather data and forecasts. Sensors across the boiler cluster and pipeline network monitor operating parameters, including supply/return temperatures, pressures, gas flows, and valve openings. Indoor sensors record average room temperatures at T -minute intervals to reflect terminal heating demands. At the edge, raw data undergo standardized preprocessing, including 3σ outlier removal, linear interpolation for missing values, and timestamp alignment. The processed data are then encrypted and transmitted to a cloud database, achieving real-time digital representation of the physical system.




Figure 2. Cloud–edge device collaborative architecture diagram.

To ensure feasible and practical optimization strategies, the architecture strictly adheres to operational constraints [23]. We embed physical boundaries and management rules into the virtual model. Hard constraints enforce mandatory limits for active boiler counts, supply temperatures, and indoor temperatures. Soft rules implement graded heating strategies for different functional zones. Simulation constraints also account for external disturbances like weather and occupant behavior. This setup enables accurate simulation under diverse conditions and supports robustness verification of control strategies.

Bidirectional data interaction between the cloud and the edge drives the closed-loop control [24]. The edge uploads preprocessed real-time data to synchronize cloud model status. The cloud then sends control instructions, such as boiler counts and supply temperature setpoints, back to the edge for execution. After execution, actual room temperatures and equipment responses are fed back to the cloud. This feedback corrects cloud model parameters to minimize simulation deviations and allows adaptive adjustment of the forward strategy based on target deviations. This complete cycle of perception, decision, execution, and feedback synchronizes the physical site with cloud intelligence, enabling dynamic adaptation to evolving system states.

4.2. Overall Architecture of Dual-System Coordinated Control

This paper constructs a collaborative closed-loop control architecture that combines a forward decision-making system and a reverse verification system. The forward system generates an initial data-driven control strategy based on multi-source data. The reverse system evaluates execution effects and quantifies uncertainty. Both systems rely on a cloud–edge device platform for bidirectional linkage and iterative optimization, forming a closed-loop control system that is adaptable to diverse operating conditions [25].

As shown in Figure 3, the two subsystems operate on the same 5×15 time-series window. The forward model predicts boiler status and supply temperatures, revealing how heating must be performed. The reverse model predicts the future 5 h room temperature trajectory ($\hat{Y}_{t+1:t+5}$), revealing what happens if heating is performed in that way. Together, they form a dynamic bidirectional closed loop through four steps: First, the Mamba-based forward model processes the shared window to generate an initial boiler strategy. Second, the system performs a causal injection, overwriting the boiler parameters in the last time step of the reverse model's window with the proposed strategy. The TCN-MC Dropout model then evaluates this strategy and outputs a predicted temperature trajectory with prediction intervals. Third, a strategy corrector compares the prediction with the target setpoint and adjusts the supply temperature accordingly. If severe undercooling is predicted, the system forces one boiler open as a fallback. Fourth, the system injects the corrected strategy back into the reverse model for second-round prediction, ensuring the final temperature trajectory matches the finalized boiler plan. Through these rounds of evaluation and correction, the architecture establishes a true dynamic closed loop. The finalized strategy then proceeds to physical execution.

After execution, the system uploads actual operating data and room temperature feedback to the collaborative data layer. This feedback serves as input for the next regulation round and also corrects model parameters, synchronously updating both the forward Mamba decision model and the reverse TCN-MC Dropout verification model. This continuous predict–verify–feedback–correct cycle distinguishes the framework from traditional non-interactive cascaded structures, covering data input, strategy generation, verification, execution, feedback, and model iteration.




Figure 3. Flowchart of dual-system coordination and control.

4.3. Mamba-Based Forward Boiler Decision Model

Mamba offers exceptional long sequence modeling capabilities and linear computational efficiency [26,27]. Leveraging these advantages, the proposed Mamba-based forward boiler decision model drives initial strategy generation within the dual-system collaborative architecture. Together with the reverse room temperature verification model, it forms a complete decision-making and verification closed loop. This forward model processes multi-source time-series state data from the collaborative data layer. It rapidly generates data-driven boiler control strategies while satisfying operational constraints. Ultimately, this output serves as the initial decision-making basis for subsequent closed-loop control.

To achieve precise control, this study defines the operational decision-making of large-scale industrial heating as a multivariate time-series hybrid prediction task. The core objective is to generate directly executable boiler control commands. These commands are derived from historical states, real-time environmental data, and future weather forecasts aggregated by the data layer. This process ensures standard indoor temperatures across all zones. Furthermore, it provides the initial basis for subsequent dual-system closed-loop regulation. Centralized heating systems exhibit significant thermal inertia. To capture these dynamics, the problem is formulated mathematically. Given a system state sequence $\{S_{t-L+1}, S_{t-L+2}, \dots, S_t\}$ over L consecutive historical time steps, alongside weather forecasts and room temperature targets at time $t + k$, the model solves for the boiler control vector $\mathbf{u}_{t+k}$. Here, Δ_t represents the time step size, L denotes the input sequence length, and k is the prediction step size. These parameters can be flexibly adjusted to match the operational traits and control needs of various heating systems.

The single-time-step input feature vector $\mathbf{X}_t$ consists of four types of features: time encoding, environment and room temperatures, boiler operating status, and future constraints. Its rigorous mathematical definition is as follows:

$$\mathbf{X}_t = [\mathbf{X}_{\text{time}}, \mathbf{X}_{\text{state}}, \mathbf{X}_{\text{boiler}}, \mathbf{X}_{\text{future}}]^T. \quad (4)$$

Here, $\mathbf{X}_{\text{time}} = \left[\sin\left(2\pi \cdot \frac{h_t}{24}\right), \cos\left(2\pi \cdot \frac{h_t}{24}\right) \right]$ represents the time encoding features. The sub-vector $\mathbf{X}_{\text{state}} = [T_{\text{outdoor}}(t), T_{\text{panel}}(t), T_{\text{indoor}}(t)]$ denotes the environmental and indoor states at time t . Specifically, $T_{\text{outdoor}}(t)$, $T_{\text{panel}}(t)$, and $T_{\text{indoor}}(t)$ represent the outdoor temperature, panel temperature, and average indoor temperature, respectively. The term $\mathbf{X}_{\text{boiler}} = [T_{1,\text{sup}}(t), T_{1,\text{ret}}(t), \dots, T_{4,\text{sup}}(t), T_{4,\text{ret}}(t)]$ comprises the operating parameters of the boiler cluster. In this sub-vector, $T_{i,\text{sup}}(t)$ and $T_{i,\text{ret}}(t)$ represent the supply and return water temperatures of the i -th boiler. Finally, $\mathbf{X}_{\text{future}} = [T_{\text{outdoor}}(t+k), T_{\text{indoor}}(t+k)]$ defines the future boundary constraints at $t+k$ time steps. Within this block, $T_{\text{outdoor}}(t+k)$ denotes the forecast outdoor temperature, and $T_{\text{indoor}}(t+k)$ represents the target indoor temperature setpoint.

The model input is a two-dimensional tensor $\mathcal{X}_t = [X_{t-L+1}, X_{t-L+2}, \dots, X_t]^T$. This tensor is formed by concatenating feature vectors of L consecutive time steps in chronological order. The sequence length L can be adjusted flexibly based on the thermal inertia traits of different heating systems. Meanwhile, the model output is an eight-dimensional hybrid control vector $\mathbf{u}_{t+k}$ for the future time $t+k$. This vector contains both the boiler on/off statuses and the outlet water temperatures. Its rigorous definition is formulated as follows:

$$\mathbf{u}_{t+k} = [s_1(t+k), T_1(t+k), s_2(t+k), T_2(t+k), s_3(t+k), T_3(t+k), s_4(t+k), T_4(t+k)]^T, \quad (5)$$

where $s_i(t+k)$ represents the on/off status of the i -th boiler at time $t+k$, and $T_i(t+k)$ denotes the outlet water temperature of the i -th boiler at time $t+k$. Specifically, the forward decision model predicts the operating statuses and outlet temperatures of the four boilers. These parameters can be adjusted flexibly based on the equipment configuration of actual engineering scenarios.

Figure 4 illustrates the overall architecture of the Mamba-based forward boiler decision model. The model takes multi-source time-series state data from the collaborative data layer as input. Through hierarchical feature extraction and time-series modeling, it rapidly generates the data-driven boiler control strategy while satisfying all operational constraints.




Figure 4. Architecture diagram of the Mamba-based positive decision-making model.

The model first inputs the real-time system state S_t into the feature extraction and encoding layer. This layer sequentially performs data normalization, time-series encoding, and constraint embedding. Specifically, data normalization eliminates dimensional differences, while time-series encoding injects positional information. Constraint embedding incorporates hard constraints, such as boiler operating boundaries and room temperature standards. This step ensures the executability of the generated strategy from the outset.

Next, the preprocessed feature vectors enter a three-layer stacked Mamba temporal prediction network. Each Mamba block consists of layer normalization, causal convolution, selective state-space computation, and a feedforward network. Residual connections are also integrated into the architecture. Within each block, the causal convolution layer strictly guarantees temporal causality. Concurrently, the core selective state-space layer efficiently captures long-term dependencies. This mechanism accurately models the long time delays caused by building thermal inertia.

The high-level features output from the Mamba network then flow into the strategy output and loss calculation layer. Here, the control strategy output layer generates an action a_t , which defines the number of active boilers and the water supply temperature. This initial strategy is sent to the physical system for execution and to the reverse verification system for pre-verification. Meanwhile, the loss calculation module uses historical operational data as labels. It employs a hybrid loss function combining the mean squared error and constraint loss to calculate prediction errors. Through backpropagation, the module updates network parameters to continuously optimize strategy generation accuracy.

This model supports both real-time inference and offline training modes. In the real-time inference mode, the system performs only forward propagation. Its millisecond-level strategy generation successfully meets the real-time requirements of industrial control. Conversely, the offline training mode utilizes historical batch data for periodic iterations. This process continuously improves the model's generalization ability across diverse operating conditions.

Data-driven deep learning control models typically have black-box characteristics, which reduce decision trustworthiness in engineering applications. To address this issue, this study also adopts univariate perturbation sensitivity analysis as the interpretability verification method for the forward decision model.

This method follows the control variable principle. It fixes all input features except one, and observes the response pattern of control outputs to the single variable change. It can verify whether model decisions conform to the physical mechanism of heating systems, and effectively reflects the decision traceability and physical rationality of the forward model. The corresponding experimental verification and result analysis are presented in Section 5.

4.4. Inverse Room Temperature Validation Model Based on TCN-MC Dropout

As discussed earlier, TCNs excel at rapid time-series modeling on edge devices [28,29]. To provide reliable uncertainty quantification, the reverse validation module integrates a TCN architecture with MC Dropout. This powerful combination enables Bayesian inference within the network. Specifically, by keeping Dropout active during the inference stage, the model executes multiple forward propagations. This mechanism successfully approximates the output as a Gaussian distribution, thereby yielding both the predictive mean and its corresponding prediction intervals [30].

Inverse room temperature verification serves as a vital evaluation module within large-scale industrial centralized heating closed-loop control systems. This study models this verification as a Multi-Input Single-Output (MISO) time-series probabilistic prediction problem. The core objective is to accurately predict future regional indoor temperatures.

This prediction relies on several inputs aggregated by the data layer, including boiler operation strategies from the forward module, historical environmental data, past room temperatures, and future meteorological constraints.

The inverse room temperature verification is a probabilistic prediction task for multiple steps. It achieves precise forecasting and rigorous uncertainty quantification. Let $\{S_{t-L+1}, \dots, S_t\}$ denote the historical state window of length L . Let U_t represent the data-driven boiler control strategies from the forward module. The inverse model predicts the indoor temperature $\hat{y}_{t+k}$ at each future step $t+k$ for $k = 1, \dots, K$. It also provides a 95% prediction interval $[LB_{95\%}, UB_{95\%}]$. Here, L is the input sequence length and K is the prediction horizon. In the proposed system, L and K equal 5. The time step is one hour. These values can change for other industrial heating scenarios.

Take the input feature vector $\mathbf{X}_t$ for a single time step. It aligns with the 15-dimension layout of the forward module. It consists of four feature types: time encoding, environment and room temperatures, boiler operating status, and future constraints. Its rigorous mathematical definition is as follows:

$$\mathbf{X}_t = [\mathbf{X}_{\text{time}}, \mathbf{X}_{\text{state}}, \mathbf{X}_{\text{boiler}}, \mathbf{X}_{\text{future}}]^T. \quad (6)$$

Here, $\mathbf{X}_{\text{time}} = \left[\sin\left(2\pi \cdot \frac{h_t}{24}\right), \cos\left(2\pi \cdot \frac{h_t}{24}\right) \right]$ represents the time encoding features. The variable h_t is the current hour. The sub-vector $\mathbf{X}_{\text{state}} = [T_{\text{outdoor}}(t), T_{\text{panel}}(t), T_{\text{indoor}}(t)]$ denotes the environmental and indoor states at time t . Specifically, $T_{\text{outdoor}}(t)$, $T_{\text{panel}}(t)$, and $T_{\text{indoor}}(t)$ represent the outdoor temperature, panel temperature, and average building indoor temperature. The term $\mathbf{X}_{\text{boiler}} = [T_{1,\text{sup}}(t), T_{1,\text{ret}}(t), \dots, T_{4,\text{sup}}(t), T_{4,\text{ret}}(t)]$ contains the operating parameters of the boiler cluster. In this vector, $T_{i,\text{sup}}(t)$ and $T_{i,\text{ret}}(t)$ represent the supply and return water temperatures of the i -th boiler. In closed-loop operation, the forward module control output U_t replaces the boiler terms at the final time step of the historical window. All other features remain identical to the forward input. Finally, $\mathbf{X}_{\text{future}} = [T_{\text{outdoor}}(t+5), T_{\text{target}}(t+5)]$ defines the future boundary constraints at the target step $t+5$. Within this block, $T_{\text{outdoor}}(t+5)$ denotes the forecast outdoor temperature. $T_{\text{target}}(t+5)$ represents the target indoor temperature setpoint.

Consequently, the complete input to the model is the sequence tensor:

$$\mathcal{X}_t = [\mathbf{X}_{t-L+1}, \mathbf{X}_{t-L+2}, \dots, \mathbf{X}_t]^T. \quad (7)$$

Thus, the inverse module integrates recent historical states and the forward boiler control plan within a single unified window. The inverse module generates probabilistic outputs via the Monte Carlo Dropout mechanism. This quantifies the prediction risk. For each lead time $t+k$ ($k = 1, \dots, 5$), the output vector is as follows:

$$\mathcal{Y}_{t+k} = [\hat{y}_{t+k}, LB_{95\%}, UB_{95\%}]. \quad (8)$$

Here, $\hat{y}_{t+k}$ represents the mean predicted indoor temperature. $LB_{95\%}$ and $UB_{95\%}$ denote the lower and upper bounds of the 95% prediction interval from the Monte Carlo sampling. This provides a strict and quantifiable risk assessment for the closed-loop system.

Figure 5 illustrates the overall architecture of the TCN-MC Dropout-based reverse verification model. The model receives system status data collected from the cloud-edge data layer and control strategies output by the forward Mamba model. It processes these inputs through three core stages: feature extraction, temporal modeling, and uncertainty quantification. Ultimately, the model outputs multi-step room temperature predictions accompanied by prediction intervals. This probabilistic output provides critical support for pre-verifying and dynamically correcting the forward strategy.




Figure 5. Architecture diagram of the reverse verification model based on TCN-MC Dropout.

The model takes two primary inputs: the system state data S_t collected from the cloud and edge, and the initial regulation strategy a_t output by the forward Mamba model. First, these data enter the feature extraction and encoding layer. This layer sequentially performs data normalization, time-series encoding, and policy encoding. Specifically, data normalization eliminates dimensional differences, while time-series encoding injects positional information. Meanwhile, policy encoding fuses discrete and continuous features, establishing a solid foundation for subsequent time-series modeling.

Next, the fused sequence tensor, with an input length of $L = 5$ and a feature dimension of $D = 15$, enters the stacked causal TCN for temporal modeling. The tensor first passes through a 1×1 convolution layer to map the 15-dimensional features into 64-dimensional hidden channels. It is then processed by three stacked TCN blocks. Each block integrates causal convolutions, batch normalization, ReLU activation, residual connections, and MC Dropout (with a probability of $p = 0.5$ applied to the blocks and output heads). Within the network, the causal convolutional layers strictly guarantee temporal causality to prevent data leakage, while the MC Dropout layer remains active during both training and inference to enable uncertainty quantification.

After passing through LayerNorm, the feature at the final time step is extracted and fed into a cascading fully connected structure comprising 64, 32, and 5 neurons, respectively, to output the point predictions of indoor temperature for the next five steps.

During the inference stage, MC Dropout is employed with a Monte Carlo sampling frequency of $N_{MC} = 20$ forward passes. This value meets the real-time computational limits on edge devices. Note that $N_{MC} = 50$ is used in offline experiments for statistical stability. For each prediction step ($t + 1 \sim t + 5$), the point prediction is determined by the mean of these 20 forward passes. Furthermore, to evaluate thermal comfort and risk under uncertainty, the model constructs a 95% prediction interval using the empirical quantiles of 2.5% and 97.5% from the predicted samples rather than relying on a traditional prediction interval for conditional expectations. This prediction interval effectively captures the dispersion of future room temperatures and epistemic model uncertainty. Regarding computational overhead, MC inference requires approximately 20 times the computation time of a single forward pass while remaining well within the limits of hourly online verification requirements. Finally, these prediction results and uncertainty bounds flow into the forward model correction module.

4.5. Adaptive Feedback Correction Mechanism

Building on the probabilistic forecasts and uncertainty bounds output by the reverse verification module, this section designs a targeted adaptive feedback correction mechanism to close the entire control loop. The mechanism converts the pre-verification conclusions of the reverse model into executable adjustment signals for the forward strategy. It realizes iterative optimization of control schemes before actual physical execution.

The mechanism comprises a TCN module and a rule-based adjuster. The TCN module predicts room temperature trajectories and quantifies uncertainty to guide adjustments under normal conditions. The rule-based adjuster generates discrete boiler commands and enforces safety limits, bridging predictions to executable actions. The two modules cycle through evaluation, action, and re-verification until the comfort requirement is met or the maximum iterations are reached.

The adaptive feedback correction mechanism bridges virtual strategy evaluation and physical operation [31,32]. It establishes two complementary feedback paths: a fast in-cycle strategy iteration for real-time refinement (Section 4.5.1), and slow cross-cycle model maintenance for thermal drift adaptation (Section 4.5.2). Section 4.5.3 details their hierarchical execution across the integrated architecture.

4.5.1. In-Cycle Strategy Iteration

This fast closed loop operates within each control cycle. Rather than retraining neural networks online, it uses the inverse model as a virtual plant to iteratively refine boiler control actions until comfort requirements are met.

At decision time t , the forward model f_{θ_f} generates an initial plan $\mathbf{u}_t$ based on the historical window $\mathbf{X}_{t-4:t}$. The inverse model g_{θ_i} then evaluates this plan by injecting $\mathbf{u}_t$ into its inputs. It outputs a multi-step temperature trajectory $\{\hat{Y}_{t+k}\}_{k=1}^5$.

During the iterative correction process, the TCN prediction module and the rule-based adjuster form a cyclic collaborative workflow. The prediction module first quantifies the deviation from the comfort target. The adjuster then uses this deviation and preset rules to generate control adjustments. The adjusted strategy is fed back to the prediction module for re-verification, and this cycle repeats until the temperature error converges or the maximum iteration limit is reached.

A rule-based adjuster refines $\mathbf{u}_t$ via coarse-to-fine iterations. At iteration m , the short-horizon comfort proxy $\bar{y}^{(m)}$ and indoor comfort error $e^{(m)}$ are defined as follows:

$$\bar{y}^{(m)} = \frac{1}{3} \sum_{k=1}^3 \hat{Y}_{t+k}^{(m)}, \quad (9)$$

$$e^{(m)} = T_{\text{set}} - \bar{y}^{(m)}. \quad (10)$$

The dynamic adjustment step size Δ_m is formulated as follows:

$$\Delta_m = \max\left(\Delta_{\min}, \min\left(\Delta_0 \gamma^m, |e^{(m)}|\right)\right), \quad (11)$$

where τ is the tolerance, Δ_0 is the initial step, $\gamma \in (0, 1)$ is the decay factor, and $\Delta_{\min}$ is the minimum step.

The adjuster increases supply temperatures by Δ_m if $e^{(m)} > \tau$, and decreases them if $e^{(m)} < -\tau$. All commands are clipped to $[T_{\min}, T_{\max}]$. A cold-start safeguard forces boiler ignition if all units are off but the forecast is critically low. To monitor the system state, a comprehensive cost $J_t = \lambda_1 G_t + \lambda_2 E_t + \lambda_3 \sigma_t^2$ is computed using the MC Dropout variance σ_t^2 . Algorithm 1 and Table 2 detail this edge-configurable procedure.

Algorithm 1 In-cycle coarse-to-fine iterative strategy correction at t

Input: Historical window $\mathbf{X}_{t-4:t}$, target set-point T_{set} , models $f_{\theta_f}, g_{\theta_i}$

Parameters: Corrector hyperparameters $(\tau, \Delta_0, \gamma, \Delta_{\min}, M, T_{\min}, T_{\max}, \delta_{\text{cs}})$

- 1: $\mathbf{u}_t \leftarrow f_{\theta_f}(\mathbf{X}_{t-4:t})$ ▷ Initial strategy via forward model
- 2: $\mathbf{X}' \leftarrow \text{InjectBoiler}(\mathbf{X}_{t-4:t}, \mathbf{u}_t)$
- 3: $\{\hat{Y}_{t+k}\}_{k=1}^5 \leftarrow g_{\theta_i}(\mathbf{X}')$ ▷ Predict future 5-step temperature
- 4: $\bar{y} \leftarrow \frac{1}{3} \sum_{k=1}^3 \hat{Y}_{t+k}$ ▷ Average of first 3-step prediction
- 5: $e \leftarrow T_{\text{set}} - \bar{y}$ ▷ Temperature deviation
- 6: $m \leftarrow 0$ ▷ Iteration counter
- 7: **while** $|e| > \tau$ **and** $m < M$ **do**
- 8: $\Delta \leftarrow \max(\Delta_{\min}, \min(\Delta_0 \gamma^m, |e|))$ ▷ Adaptive adjustment step size
- 9: **if** $e > \tau$ **then**
- 10: $\mathbf{u}_t \leftarrow \text{IncreaseSupply}(\mathbf{u}_t, \Delta; T_{\min}, T_{\max})$
- 11: **else if** $e < -\tau$ **then**
- 12: $\mathbf{u}_t \leftarrow \text{DecreaseSupply}(\mathbf{u}_t, \Delta; T_{\min}, T_{\max})$
- 13: **end if**
- 14: **if** $m = 0$ **and** $\text{AllOff}(\mathbf{u}_t)$ **and** $e > \tau + \delta_{\text{cs}}$ **then**
- 15: $\mathbf{u}_t \leftarrow \text{ColdStart}(\mathbf{u}_t)$ ▷ Cold start compensation
- 16: **end if**
- 17: $\mathbf{X}' \leftarrow \text{InjectBoiler}(\mathbf{X}_{t-4:t}, \mathbf{u}_t)$
- 18: $\{\hat{Y}_{t+k}\}_{k=1}^5 \leftarrow g_{\theta_i}(\mathbf{X}')$ ▷ Re-verify correction effect
- 19: $\bar{y} \leftarrow \frac{1}{3} \sum_{k=1}^3 \hat{Y}_{t+k}$
- 20: $e \leftarrow T_{\text{set}} - \bar{y}$
- 21: $m \leftarrow m + 1$
- 22: **end while**
- 23: $J_t \leftarrow \lambda_1 G_t + \lambda_2 E_t + \lambda_3 \sigma_t^2$ ▷ Comprehensive cost
- 24: **Output:** $\mathbf{u}_t, \{\hat{Y}_{t+k}\}_{k=1}^5, J_t$, convergence flag $\mathbb{I}(|e| \leq \tau)$

Note: The typical step sequence is $2.0 \rightarrow 1.0 \rightarrow 0.5^\circ\text{C}$ until $|e| \leq \tau$.

Table 2. Default hyperparameters of the in-cycle iterative adjuster and offline update policy.

Symbol	Meaning	Typical Value
τ	Comfort band; stop if $ e \leq \tau$	0.3°C
Δ_0	Initial supply temperature step	2.0°C
γ	Step decay per iteration	0.5
$\Delta_{\min}$	Minimum single-step adjustment	0.5°C
M	Maximum adjustment iterations	5
$\bar{y}$	Feedback signal	Mean of $\hat{Y}_{t+1}, \hat{Y}_{t+2}, \hat{Y}_{t+3}$
Termination	Convergence or cap	$ e \leq \tau$ or $m = M$
θ_f, θ_i	Online NN update	Frozen
Offline update	Retrain + PICP calibration	Periodic batch

4.5.2. Cross-Cycle Model Maintenance

This slow-loop path addresses gradual drifts in building thermal dynamics and equipment performance [33,34]. It systematically decouples real-time input updates from offline parameter learning through two distinct operational phases.

During every control cycle, the edge node ingests measured boiler states, zonal temperatures, and weather data after each dispatch. This continuous execution feedback refreshes the rolling history $\mathbf{X}_{t-4:t}$ to support subsequent forward decisions.

Conversely, the network weights θ_f and θ_i remain strictly frozen during online inference. Once sufficient operational logs accumulate, the cloud initiates periodic offline model maintenance. It retrains f_{θ_f} and updates g_{θ_i} using a closed-loop setting where forward predictions are injected into the inverse inputs [35]. Additionally, the prediction intervals are recalibrated on validation data to maintain the target coverage [36].

The two feedback paths form a two-layer closed-loop system with clear boundaries and complementary time scales. The fast loop runs within each hourly cycle. It adjusts the current strategy without updating network weights to ensure real-time accuracy. The slow loop runs offline across multiple cycles. It updates model parameters and prediction intervals using accumulated data to correct long-term drift. The fast loop provides real-time samples for the slow loop to retrain the model. After retraining, the slow loop sends the updated model back to the edge for the fast loop. The fast loop ensures real-time response, while the slow loop guarantees long-term effectiveness. Together, they form a complete adaptive closed-loop system.

4.5.3. Cloud–Edge Device Integration

This hierarchical architecture distributes computational loads in a data-driven manner. Edge nodes execute real-time virtual simulation and strategy correction (Algorithm 1). Simultaneously, the cloud platform manages heavy periodic retraining, uncertainty calibration, and model versioning. This effectively couples fast-loop decision-making with slow-loop model adaptation.

5. Experimental Results and Analysis

This study selects the centralized heating system of Shaanxi Normal University as the verification scenario. Although the actual deployment and verification of this study were conducted in the heating system of Shaanxi Normal University, the core design logic of this architecture is based on general thermal dynamics. Its direct transferability to other systems, however, requires further validation due to differences in thermophysical properties.

The total heating area is approximately 1.5 million square meters. The standard heating period is 121 days. However, due to sudden cold spells, the system actually operated for 129 days (5 days of early heating and 3 days of extended heating). The average indoor temperature was maintained at approximately 22.4 °C, while the average outdoor temperature was lower than the previous season.

The dataset was collected by the Logistics Management Department of Shaanxi Normal University. It covers three heating seasons from 2023 to 2026 (November to March), with data recorded every 20 min, totaling 7279 valid samples. The raw dataset is sampled every 20 min. For model training, we aggregate three consecutive samples by averaging them to construct a one-hour time step. This matches the control frequency of the heating system. The dataset is chronologically split into training, validation, and test sets at a 7:2:1 ratio. In preprocessing, non-heating season data are removed, shutdown indicators “/” are mapped to 0, and samples with missing values are discarded. Sine and cosine functions encode hourly timestamps to preserve diurnal rhythms. All temperature variables are linearly scaled to stabilize gradient convergence. A sliding window of size 5 reconstructs temporal samples, enabling the model to use historical dynamics and future forecasts to output control commands at the target step. To avoid ambiguity, we consistently use the term ‘data sampling’ for the 20 min cycle, ‘model step’ for the one-hour aggregation, and ‘prediction horizon’ for the 5 h forward window throughout this paper.

Building upon this foundation, the boiler operates at an actual thermal efficiency of approximately 92%. The total designed heat load is 55 MW, which corresponds to an hourly heat output of 198 GJ. Measured data indicate that the system heat load decreases by approximately 9.25 GJ for every 1 °C increase in outdoor temperature. This variation corresponds to a reduction of roughly 286 m³ in the boiler’s hourly natural gas consumption. The subsequent validation utilizes these engineering parameters. The experiments comprehensively evaluate the proposed dual-system collaborative control architecture through both simulation and actual deployment.

The backend was fully developed in Python 3.9, covering the complete workflow of deep learning model construction, training, and data processing. The frontend adopts the Vue3 framework to build responsive visualization and interactive interfaces. The proposed dual-system collaborative control architecture was systematically evaluated through both simulation experiments and field deployment.

To minimize training variability caused by random weight initialization and Dropout sampling, all deep learning models were trained and evaluated in five independent runs with distinct random seeds. All reported metrics, line plots, and bar charts are the average results of these five runs.

5.1. Simulation Experiments

5.1.1. Positive Decision Experiment

The forward decision model is a multivariate time-series forecasting module. It determines the on/off statuses and supply water temperatures of the four boilers, generating initial strategies for subsequent inverse verification and closed-loop control.

Figure 6 shows the training loss decreasing steadily. A sharp drop within the first 20 epochs indicates rapid capture of thermal inertia and load patterns. The loss then converges to approximately 0.05 after 40 epochs without rebounding, demonstrating robust generalization and an absence of overfitting.




Figure 6. Loss curve of the forward Mamba model training.

The discrete boiler on/off state critically impacts control reliability and heating safety [37]. Figure 7 evaluates classification performance on a test set covering high, medium, and low loads. The ON state accuracy is 98.43%, with a false-negative rate of only 1.57%. The OFF state accuracy is 83.44%, mainly due to misjudgments during load transitions. Overall classification accuracy reaches 92.09% with a macro F1-score of 91.27%. These metrics confirm reliable initial decision support for boiler regulation.




Figure 7. Boiler status confusion matrix.

5.1.2. Reverse Verification Experiment

The reverse verification module performs multi-step room temperature forecasting and uncertainty quantification. Before physical execution, it evaluates whether the forward strategy achieves the target thermal comfort and minimizes prediction risks.

Table 3 summarizes point prediction errors across one-to-four-step horizons on the 2024 test set and the 2025 new dataset. The model achieves its highest accuracy at the single-step prediction ($t + 1$), with a Mean Absolute Error (MAE) of 0.1786 °C. Although the error increases with the prediction horizon, the fourth-step MAE remains at 0.2354 °C on the test set, staying well within 0.3 °C without exponential divergence. This confirms that the model captures thermal inertia and suppresses error accumulation during multi-step forecasting.

Table 3. Point prediction errors of the model (dropout disabled) on the 2024 and 2025 datasets.

Prediction Step	Winter 2024 Test Set		Winter 2025 New Dataset	
	MAE (°C)	Max. Error (°C)	MAE (°C)	Max. Error (°C)
$y(t + 1)$	0.1786	0.4102	0.1943	0.8851
$y(t + 2)$	0.2076	0.7293	0.2079	0.7558
$y(t + 3)$	0.2331	1.0690	0.2747	0.6209
$y(t + 4)$	0.2354	1.2777	0.2792	1.0890

Table 4 reports the Prediction Interval Coverage Probability (PICP) at the 90%, 95%, and 99% target levels. Across all steps, empirical coverage rates closely approximate theoretical targets, with maximum deviations within 3 percentage points. The 95% coverage stabilizes between 92% and 96%, confirming that MC Dropout generates reliable bounds across different temporal distributions.

Table 4. PICP on the 2024 and 2025 datasets.

Prediction Step	Winter 2024 Test Set			Winter 2025 New Dataset		
	90%	95%	99%	90%	95%	99%
$y(t + 1)$	92.71%	95.76%	99.29%	89.96%	94.31%	98.83%
$y(t + 2)$	92.94%	95.53%	97.41%	89.55%	93.26%	99.26%
$y(t + 3)$	92.71%	94.82%	96.94%	90.81%	92.52%	97.77%
$y(t + 4)$	89.59%	92.00%	97.88%	89.30%	93.06%	97.28%

We further evaluated probability distributions using Negative Log-Likelihood (NLL) and the Continuous Ranked Probability Score (CRPS). As shown in Table 5, both scores remain consistently low. The Gaussian CRPS reaches 0.0657 on the 2024 test set and 0.0500 on the 2025 dataset. The strongly negative NLL values indicate a high likelihood that true observations align with predicted distributions.

Table 5. Probabilistic distribution scores on the 2024 and 2025 datasets.

Prediction Step	Winter 2024 Test Set		Winter 2025 New Dataset	
	NLL	CRPS	NLL	CRPS
Overall	−0.7337	0.0657	−0.9510	0.0500
$y(t + 1)$	−1.0718	0.0449	−1.4762	0.0269
$y(t + 2)$	−0.8131	0.0595	−0.9865	0.0452
$y(t + 3)$	−0.6055	0.0732	−0.7443	0.0586
$y(t + 4)$	−0.4444	0.0854	−0.5970	0.0695

Regarding interval sharpness, the 95% prediction intervals maintain constrained widths. The Normalized Mean Prediction Interval Width (NMPIW) remains below 0.25 on the test set and 0.36 on the new dataset. Winkler scores, which penalize uncaptured true values, range from 0.4202 ($t + 1$) to 0.7651 ($t + 4$) on the test set. These metrics demonstrate that the model generates sharp, informative intervals rather than overly conservative estimates. This capability allows the dual-system framework to filter out highly uncertain strategies before physical execution, enhancing robustness without sacrificing efficiency. Note that the results above are obtained in the open-loop setting. The closed-loop comparison and the corresponding coverage reduction are detailed in the comparative experiment section.

Temporal weight analysis shows that the TCN effectively captures thermal inertia. The autoregressive room temperature dominates short-term predictions but gradually decays over longer horizons as the model shifts reliance toward environmental and heating features. This physically consistent behavior, along with the lightweight architecture, prevents overfitting and supports practical deployment.

5.1.3. Univariate Perturbation Experiment

To verify the decision traceability and physical consistency of the forward Mamba model, we conducted a univariate perturbation sensitivity experiment on 300 test windows. In each case, we perturbed a single input variable while holding all other features constant, quantifying the resulting changes in the average number of active boilers and the mean supply water temperature. The baseline prediction averaged 1.57 active boilers with a mean supply temperature of 49.8 °C.

We perturbed the outdoor temperature by $\{-5, -2, 0, +2, +5\}$ °C. Meanwhile, we perturbed both the indoor building temperature and target setpoint by $\{-2, -1, 0, +1, +2\}$ °C.

As illustrated in Figure 8 and summarized in Figure 9, the forward model exhibits clear and traceable sensitivity to input perturbations. As shown in Figure 8a, the outdoor temperature yields the most significant response. The system sharply increases both boiler usage and supply temperature as the outdoor temperature rises, confirming high model sensitivity to external weather conditions.




Figure 8. Univariate perturbation experiment on test windows.

More importantly, the historical building temperature depicted in Figure 8b produces a physically consistent, inverse trend. A decrease in indoor temperature steadily raises the mean supply temperature and activates more boilers, whereas an increase lowers them. This behavior perfectly matches the expected heat supply logic for maintaining indoor comfort. In contrast, the target setpoint perturbations shown in Figure 8c produce a mild and somewhat nonlinear response, exerting a much smaller influence on the overall control strategy compared to environmental factors.

Overall, this analysis confirms that the forward Mamba model generates traceable, physically consistent control decisions rather than opaque outputs. While outdoor temperature dominates the response magnitude, the building temperature response strictly aligns with practical heating control principles. Combined with explicit in-cycle feedback correction, these results strongly support the decision traceability and practical trustworthiness of the proposed framework. Note that these marginal sensitivities should be interpreted within the full multivariate control context rather than as standalone proof of complete physical causality.




Figure 9. Grouped bar chart summary of the univariate perturbation experiment on the test set.

5.1.4. Comparative Experiment of Forward Model

To verify the necessity of Mamba as the forward backbone under thermal inertia, we compared Mamba, LSTM, and a Transformer under the same data split, features, and training budget. A sequence-length ablation was conducted with input windows of 5, 12, 24, and 48 h.

Figure 10 reports mean $\pm$ SD (Standard Deviation) over $N = 5$ independent training runs. In Figure 10a,b, Mamba achieves the lowest supply temperature MAE and Root Mean Square Error (RMSE) at all four horizons. At 5 h, its MAE is 0.169°C , versus 0.238°C for LSTM and 0.204°C for the Transformer. As the window lengthens, LSTM and Transformer errors rise more clearly, while Mamba's stays near 0.16 – 0.18°C . At 24 h, the MAE gap to the Transformer is significant ($p = 0.022$).

Figure 10c shows that boiler state F1 remains comparable across models. LSTM and the Transformer score slightly higher at short horizons, while Mamba improves with longer windows and reaches 0.77 at 48 h. Mamba is therefore selected mainly for continuous temperature regression rather than discrete state classification.




Figure 10. Sequence-length ablation of forward models on the test set.

Figure 11 quantifies the temperature advantage of Mamba over the best baseline at each horizon. Both Δ MAE and Δ RMSE remain positive at 5–48 h. The advantage is most evident when the LSTM error increases at longer windows while Mamba stays stable.




Figure 11. Gap of Mamba relative to the best baseline at each input horizon.

In summary, Mamba provides the most accurate and stable supply temperature estimates under an extended thermal history. This supports its use as the forward backbone in the proposed closed-loop controller.

5.1.5. Comparative Experiment on the Reverse Model

To evaluate the causal convolutional network in the inverse module, we compared the TCN-MC model with an LSTM-MC baseline. Both models use the same closed-loop protocol.

Both models share the same setup. This includes the closed-loop input construction, data split, output horizon ($H = 5$), hidden size (64), MC Dropout inference ($N = 50$, chosen for offline evaluation to reduce sampling variance and ensure calibration), and validation set interval calibration. The forward Mamba model is frozen during inverse training. The last-step boiler channels in each input window are replaced by the forward prediction. This reflects online closed-loop conditions. Point prediction is evaluated by RMSE and MAE. Interval quality is measured by calibrated PICP and MPIW. Each model is trained five times with independent random seeds. The error bars in Figures 12 and 13 show mean $\pm$ SD.


The interval calibration procedure is as follows: For each prediction horizon, MC Dropout first produces a raw prediction interval on the validation set. A per-horizon scale factor is then fitted so that the empirical PICP approaches the nominal 95% target. If the raw MC interval is too narrow, a residual-quantile fallback is used instead. The fitted calibrator is frozen and applied unchanged to the test set. No test labels are used during calibration.

Table 6 summarizes the benchmark configurations. Test set results are shown in Figures 12 and 13.

Table 6. Comparison setup for the inverse models.

Item	TCN-MC	LSTM-MC
Input	$(B, 5, 15)$, closed-loop boiler channels	Same
Output	5-step indoor temperature + MC interval	Same
Hidden size	64	64
Backbone depth	3 causal TCN blocks	2-layer LSTM
Training loss	Huber loss on point prediction	Same
Inference	MC Dropout ($N = 50$) + validation calibration	Same

Figure 12 shows that the overall point prediction accuracy is similar for both models. The TCN-MC model achieves a mean RMSE of 1.78 ± 0.09 °C. The LSTM-MC model achieves 1.74 ± 0.11 °C. Their mean MAE values are 1.53 ± 0.10 °C and 1.47 ± 0.12 °C, respectively. Paired t -tests show no significant difference ($p > 0.05$).

**Figure 12.** Overall performance comparison of TCN-MC and LSTM-MC models on the closed-loop test set (mean RMSE, mean MAE, and mean MPIW).

The step-wise results in Figure 13 reveal different multi-step behaviors. Figure 13a,b show that the TCN-MC model has a lower RMSE and MAE at $t + 1$ and $t + 2$. The LSTM-MC model performs better at $t + 3$, $t + 4$, and $t + 5$. Online correction mainly uses near-term horizons ($t + 1$ to $t + 3$). The TCN-MC model's advantage at $t + 1$ and $t + 2$ is therefore more relevant for closed-loop control.

Regarding uncertainty estimation, the two models show a clear trade-off. Figure 13c reports the horizon-specific PICP. For the TCN-MC model, the calibrated coverage is 0.92, 0.72, 0.80, 0.94, and 0.94 at $t + 1$ to $t + 5$, with a mean PICP of 0.86 ± 0.10 . Thus, the TCN-MC model remains below the nominal 95% level at $t + 2$ and $t + 3$. The LSTM-MC model achieves higher coverage at all horizons, with a mean PICP of 0.98 ± 0.02 . Figure 13d shows that the LSTM-MC model uses wider intervals. Its mean MPIW is 11.22 ± 1.38 °C. The TCN-MC model maintains a tighter mean MPIW of 8.30 ± 1.85 °C ($p = 0.011$). The widest LSTM intervals appear at $t + 3$ and $t + 5$.

This undercoverage is an important limitation for risk-aware control. Because prediction intervals support risk assessment and iterative correction, a PICP below 0.95 means that true indoor temperatures may fall outside the reported band more often than intended. The risk is largest at $t + 2$ and $t + 3$, which are still used in near-term correction. In practice, this may lead to over-confident decisions and delayed correction when the reverse model underestimates uncertainty. Therefore, interval-based stopping and adjustment rules should be interpreted with caution at these horizons. Conservative safeguards, such as a larger temperature band or additional iteration checks, remain necessary.

The undercoverage is mainly attributed to closed-loop distribution shift. The boiler channel at the last time step is replaced by the forward Mamba prediction. This makes the inference input differ from the training distribution. Forward errors also accumulate with the prediction horizon. As a result, validation-calibrated intervals can underestimate uncertainty on the closed-loop test set, especially at mid-term steps.




Figure 13. Step-wise evaluation of TCN-MC and LSTM-MC models on the closed-loop test set across four key metrics.

Overall, the LSTM-MC model is a competitive baseline with higher interval coverage. The TCN-MC model remains more suitable for the inverse verification task for three reasons: First, it achieves comparable point prediction accuracy. Second, it produces significantly tighter prediction intervals. Third, it performs better at $t + 1$ and $t + 2$, the horizons most used in closed-loop correction. At the same time, its tighter intervals come with undercoverage relative to the nominal 95% target. Combined with efficient parallel causal convolutions on short input sequences, the TCN-MC model is selected as the core inverse verification module. Improving closed-loop calibration and horizon-wise coverage is left as future work.

5.1.6. Ablation Experiment

To quantify the performance contribution of each module, we conducted a framework-level ablation study under unified offline simulation conditions. The test used $N = 200$ historical windows randomly sampled from the test set. Results are reported as mean $\pm$ SD over five independent evaluation runs, and three metrics are adopted. Four configurations are evaluated as follows: First, the Full framework retains forward Mamba prediction, reverse TCN-MC with MC Dropout, and iterative closed-loop correction. Second, *w/o* Reverse Verification keeps the reverse model with MC Dropout but disables in-cycle iterative correction. Third, *w/o* Uncertainty retains reverse prediction and iterative correction but disables MC Dropout so that only point predictions are used. Fourth, Open-Loop uses the forward model alone, and the reverse model serves only as an indoor temperature evaluator without participating in strategy correction.

The compliance threshold in this ablation study was set to $\pm 2.5\text{ }^{\circ}\text{C}$. This threshold accounts for the simulation error of the offline virtual environment and is adopted only for the ablation study to assess the relative contribution of each module. Figure 14a reports the room temperature compliance rate ($|T_{\text{in}} - T_{\text{set}}| \leq 2.5\text{ }^{\circ}\text{C}$), with Full reaching 91.0%, *w/o* Verification reaching 85.4%, and *w/o* Uncertainty reaching 78.2%. Figure 14b shows the cumulative heat proxy E_Q , with Full achieving the minimum of 14,987 GJ and others higher. Figure 14c presents the temperature fluctuation σ_T , with Full achieving the lowest, $1.546\text{ }^{\circ}\text{C}$, *w/o* Verification achieving a higher value, and *w/o* Uncertainty and Open-Loop performing equally worse. The closed-loop architecture outperforms Open-Loop across all metrics. Uncertainty contributes most, while reverse verification adds further gains. This validates the necessity of combining forward decision, reverse verification, and MC-based uncertainty in the heating control framework.




Figure 14. Ablation experiment with dual-system closed-loop control.

Crucially, the *w/o* Uncertainty configuration serves as a fully automated yet open-loop baseline. Its compliance rate (78.2%) is virtually identical to the pure open-loop case (78.0%), while the Full framework achieves 91.0%. This confirms that the performance gain is not merely due to replacing manual control with automation; instead, it is the closed-loop verification and uncertainty quantification mechanisms that deliver the substantial improvement.

5.1.7. Robustness Under Disturbed Operating Conditions

It should be noted that these robustness tests are simulated disturbances injected into the offline simulation environment. They are not physical fault validations on the actual plant. The purpose of these simulations is to evaluate the control logic's responsiveness under extreme operating conditions.

We further examined closed-loop robustness under four typical disturbances. Each episode is a 24 h rolling simulation with a fixed setpoint $T_{\text{set}} = 21\text{ }^{\circ}\text{C}$. The disturbance is injected from $t = 6\text{ h}$ to $t = 18\text{ h}$. The unperturbed run uses the same windows without the fault. Four scenarios were considered: extreme cold, indoor temperature sensor drift, communication interruption, and single-boiler outage. The tolerance $\delta = 1.5\text{ }^{\circ}\text{C}$ represents the acceptable deviation under fault conditions. This threshold is larger than the strict $\pm 1\text{ }^{\circ}\text{C}$ physical constraint in Section 4.4 because it accounts for the degraded performance during disturbances. It also differs from the $\pm 2.5\text{ }^{\circ}\text{C}$ statistical threshold used in the ablation study (Section 5.1.6), which evaluated overall seasonal performance. Compliance is defined as $|T_{\text{in}} - T_{\text{set}}| \leq \delta$ with $\delta = 1.5\text{ }^{\circ}\text{C}$. The second metric is the maximum absolute deviation during the disturbance window.

Figure 15 shows representative indoor temperature trajectories. The dashed line is the unperturbed response. It stays slightly below T_{set} and remains nearly flat. The gray band marks the disturbance interval. Figure 15 presents the indoor temperature responses under

four disturbances: Extreme cold first lowers the temperature and then causes an overshoot near $t = 18$ h. Sensor drift produces only a mild offset. Communication interruption yields the largest drop because the control strategy is frozen while the outdoor load changes. Boiler outage reduces heating capacity, causing a temperature drop that does not fully recover within 24 h.




Figure 15. Indoor temperature responses in four disturbance scenarios.

Figure 16 reports mean $\pm$ SD over $N = 5$ independent evaluation runs, where each run averages 12 episodes and error bars denote the run-to-run SD. The unperturbed compliance achieves $95.9 \pm 0.9\%$. Extreme cold reduces it to $90.6 \pm 1.4\%$. Sensor drift remains comparable at $95.9 \pm 0.9\%$. Communication interruption drops sharply to $68.5 \pm 2.1\%$. Boiler outage falls to $75.9 \pm 1.9\%$. Figure 16b shows the corresponding maximum deviation. The unperturbed peak error is $0.96 \pm 0.02^\circ\text{C}$. The four disturbances raise it to $1.51 \pm 0.05^\circ\text{C}$, $1.04 \pm 0.02^\circ\text{C}$, $1.95 \pm 0.08^\circ\text{C}$, and $1.95 \pm 0.10^\circ\text{C}$, respectively.




Figure 16. Robustness metrics in four disturbance scenarios.

Collectively, the closed-loop controller maintains bounded degradation under all four faults. Communication interruption is the most severe case. Boiler outage and extreme

cold cause moderate losses. Sensor drift has a limited impact within $\delta = 1.5\text{ }^{\circ}\text{C}$. These results support the robustness of the dual-system framework under nominal disturbed heating conditions.

However, significant performance degradation is observed under communication interruption and boiler outage. These cases reveal substantial limitations of the current framework. Specifically, during communication interruption, the edge node loses real-time data updates from the cloud, causing the control strategy to freeze. During a boiler outage, the physical heat source is lost, and the controller lacks a redundant backup mechanism to maintain thermal comfort. Therefore, practical deployment must incorporate local data caching, fault-tolerant sensing, and redundant heat source scheduling to mitigate these risks.

5.1.8. Comparison with Conventional Closed-Loop Baselines

To further address the need for a conventional closed-loop baseline, we added an offline simulation comparison under the same protocol as the framework ablation. Real plant deployment of a Proportional–Integral–Derivative Controller (PID) or industrial Model Predictive Control (MPC) is subject to practical constraints in our edge setting, particularly regarding operational safety and real-time latency requirements. Therefore, we conducted a simulation-based comparison to provide a rigorous evaluation under controlled and reproducible conditions.

Three controllers are evaluated on $N = 200$ historical test windows. Results are reported as mean $\pm$ SD over five independent runs. The full dual system uses forward Mamba prediction, reverse TCN-MC with uncertainty, and iterative correction. The PI/PID baseline starts from the last-step historical boiler setpoints and adjusts supply temperature via room temperature error feedback. The MP baseline is a simplified finite-horizon controller with discrete supply temperature increments and receding-horizon execution. Both PI/PID and MPC use the reverse point predictor as the indoor temperature plant model. They do not use the forward Mamba policy.

Figure 17 illustrates the results. The full dual system achieves a compliance rate of $90.5 \pm 0.9\%$. PI/PID and the simplified MPC both reach $78.4 \pm 0.7\%$. The paired t -test gives $p < 0.001$ for full versus either baseline. The full system also uses less heat energy ($14,997 \pm 422$ GJ versus about 15,644 GJ) and yields a smaller temperature fluctuation ($1.549 \pm 0.017\text{ }^{\circ}\text{C}$ versus $1.866 \pm 0.023\text{ }^{\circ}\text{C}$). Under this discrete-action setting, PI/PID and the simplified MPC perform almost identically. This indicates that the main gain comes from the dual-system design rather than from classical feedback alone.




Figure 17. Offline comparison of the full dual-system against PI/PID and a simplified MPC baseline.

5.2. Actual Deployment

For practical deployment, this study aligns with the intelligent heating upgrade requirements at Shaanxi Normal University. The system covers three heating zones: residential, teaching, and student areas.

To upgrade the existing platform, we adopt a layered distributed architecture and a microservices framework. The deployment uses a device–server–application structure for district heating control. The device deployment is briefly described below from several perspectives.

5.2.1. Device Layer

The upgrade adds 50 T4 socket-type Internet of Things (IoT) temperature panels across the three zones. Each panel supports NB-IoT, LoRa, and 4G, with a temperature accuracy of ± 0.3 °C. The collection and reporting cycles are configurable from 1 min to 30 days. In the pilot project, minute-level reporting was used. Table 7 lists the main specifications.

Table 7. Specifications of the T4 IoT temperature acquisition panel.

Item	Specification
Product	T4 socket-type temperature and humidity collector
Form factor	86-type wall socket, 88 × 88 × 38 mm
Display	1.2-inch LCD
Network	NB-IoT/LoRa/4G
Temperature range	−20 °C to +80 °C
Temperature accuracy	± 0.3 °C (supports ± 0.1 °C)
Collection cycle	Configurable, 1 min to 30 days
Reporting cycle	Configurable, 1 min to 30 days
Power supply	AC 100–240 V
Standby power	<0.1 W
Local storage	>10,000 records
LoRa coverage	Up to about 1 km diameter
Connectivity scale	Up to 100,000 devices (vendor specification)

The original boiler data acquisition and control equipment is fully reused. Existing PLC/SCADA interfaces remain unchanged.

5.2.2. Server Layer

A dedicated on-campus server hosts forward and inverse model inference, database services, and platform integration. It is equipped with an Intel Core Ultra 9 285K CPU, 64 GB DDR5 memory, two 1 TB SSDs, and one NVIDIA RTX 4090 GPU. The in-cycle iterative correction in Algorithm 1 runs on this server. Table 8 summarizes the hardware configuration.

Table 8. Configuration of the on-campus inference and data server.

Component	Configuration
Motherboard	ASUS PRIME Z890-P (ASUS, Taipei, Taiwan)
CPU	Intel Core Ultra 9 285K (Intel, Santa Clara, CA, USA)
Memory	Kingston DDR5 5600 MHz, 32 GB × 2 (64 GB total) (Kingston, Fountain Valley, CA, USA)
Storage	ZHITAI Ti600 1 TB SSD × 2 (YMTC, Wuhan, China)
GPU	NVIDIA RTX 4090 (48 GB, blower edition) (NVIDIA, Santa Clara, CA, USA)
Cooling	ASUS PRIME Master 360 ARGB (ASUS, Taipei, Taiwan)
Power supply	MAG A850GL, 850 W, ATX 3.1 (MSI, New Taipei City, Taiwan)

5.2.3. Application and Monitoring Layer

All operational data streams are synchronized via unified timestamps. In a 20 min cycle, the system ingests refreshed room temperatures, boiler statuses, and meteorological data. However, the forward and inverse models execute their control decisions using an aggregated one-hour time step, consistent with the training phase. The closed-loop inference engine then generates optimized control commands and logs them into the central database. Meanwhile, model retraining and prediction interval calibration are executed offline as batch tasks, isolating the online control loop from heavy computational overhead. As illustrated in Figure 18, the interactive monitoring interface presents core performance indicators and supports detailed data drill-down.




Figure 18. Real-time smart heating AI brain system interface.

5.2.4. Information Security

To ensure enterprise-grade security for campus operations, all communications and control traffic run strictly within an isolated internal Local Area Network (LAN) under the campus IT security framework. System access is governed by Role-Based Access Control (RBAC) and strict account authentication, preventing unauthorized or malicious manual command overrides. Together, these measures safeguard the physical infrastructure against external cyber threats and ensure reliable daily thermal management.

5.2.5. Software Platform

The platform runs on Windows with a MySQL database. Data transmission uses HTTPS. Open APIs connect the system to the existing heating platform. A dedicated assessment data table unifies boiler operation records and room temperatures for management and reuse.

Furthermore, this study conducted a comparative experiment between the proposed dual-system coordinated control method and the school's original manual experience-based approach. The comparison uses the complete 2024–2025 heating season (manual control) as the baseline and the complete 2025–2026 heating season (dual-system control) as the pilot. Both seasons share the same standard calendar period (November to March). Since only one physical heating system was available, these two seasons represent the sequential deployment of manual operation and the proposed dual-system operation on the same system. The comprehensive optimization objective function J serves as the core evaluation criterion.

The reported comprehensive index J values in this section are calculated directly from the raw operational data as described in Section 3. Figure 19 presents the performance improvements across multiple dimensions. First, as shown in Figure 19a, the mean value of the comprehensive index J decreases from 2.12 to 1.51. This represents a 28.6% reduction. Second, Figure 19b shows that the room temperature compliance rate across the three

zones increases from 83.1% to 96.7%. Concurrently, the temperature fluctuation range narrows from ± 1.3 °C to ± 0.75 °C (Figure 19c). This corresponds to a 42.3% reduction in temperature fluctuation. As shown in Figure 19d, the 2024–2025 season consumes approximately 7,680,291.5 m³. The 2025–2026 season consumes approximately 6,420,723.7 m³. This corresponds to an absolute reduction of 1,259,567.8 m³, or 16.4%. Importantly, despite the 2025–2026 season experiencing harsher weather conditions (lower outdoor temperatures) and operating for eight extra days due to sudden cold spells, the absolute gas consumption still decreases. Since rigorous weather-normalized analysis was not available in the current field experiment, we present this result as an observed reduction between the two seasons, which indicates a substantial improvement in operational efficiency under more demanding heating load conditions. Finally, in terms of the model's internal performance, the inverse model prediction variance falls from 0.217 to 0.091. This represents a 57.8% decrease in prediction uncertainty.

In terms of validation performance, the operational data from the industrial centralized heating system at Shaanxi Normal University indicate certain improvements achieved by the proposed collaborative architecture. Compared to traditional manual control, the proposed scheme contributes to reducing natural gas energy consumption, mitigating room temperature fluctuations, and increasing the room temperature compliance rate. These results reflect the performance characteristics of the system regarding precision, energy conservation, and robustness, thereby indicating its engineering practicality and certain application value.




Figure 19. Comparison of dual-system regulation performance indicators.

6. Conclusions

Public centralized heating systems are rapidly evolving toward energy conservation, requiring data-driven heat source efficiency alongside safe operation and thermal comfort. Traditional manual control modes, hampered by severe propagation lags and high energy

consumption, struggle to meet these complex demands. To address this, we propose an intelligent dual-system collaborative control architecture integrating a Mamba predictive model for forward decision-making with TCN-MC Dropout reverse verification. In this tightly coupled mode, the forward system generates an initial boiler control strategy, while the reverse system continuously analyzes actual operational data to verify validity, correct deviations, and assess risks. Together, these subsystems dynamically adapt to varying heating phases, optimizing energy efficiency without compromising indoor comfort.

A field deployment across three heating zones at Shaanxi Normal University validates the architecture's real-world effectiveness. Compared to traditional manual methods, the proposed framework significantly improves comprehensive optimization metrics. It effectively reduces natural gas consumption and minimizes temperature fluctuations while maintaining a room temperature compliance rate well above national standards, successfully achieving a strict balance between thermal comfort and resource efficiency.

Despite the promising results in this specific setting, it should be noted that the proposed framework has currently only been validated in a single university heating installation. Transferring it to other buildings or heating systems requires further investigation. Differences in building thermal inertia, radiator types, and historical control strategies may cause distribution shifts and require model retraining. Furthermore, its performance heavily relies on the completeness and accuracy of edge sensor data. Extreme scenarios, such as large-scale sensor failures or severe communication disruptions, may compromise the reverse verification module. This sensitivity to hardware and network anomalies remains a critical operational bottleneck that requires further architectural optimization. Future research will explore multi-modal data fusion and fault-tolerant control mechanisms to enhance overall system robustness.

Author Contributions: Conceptualization, W.Z. (Wenjing Zhang) and X.M.; methodology, W.Z. (Wenjing Zhang) and W.Z. (Wenchao Zhang); software and visualization, W.Z. (Wenjing Zhang); validation, W.H. and M.C.; formal analysis and investigation, W.Z. (Wenjing Zhang) and L.W.; resources and data curation, W.Z. (Wenchao Zhang) and M.C.; writing—original draft preparation, W.Z. (Wenjing Zhang); writing—review and editing, W.Z. (Wenchao Zhang) and X.M.; supervision and project administration, W.Z. (Wenchao Zhang) and X.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Natural Science Basis Research Plan in Shaanxi Province of China, grant number 2025JC-YBMS-674; the Technological Innovation Guidance Program (Fund)-Three Reforms-Performance-Based Evaluation, grant number 2025ZC-YYDP-51; the Scientist + Engineer Team Building in Qinchuangyuan, grant number 2025QCY-KXJ-169; the Key Research and Development Program of Shaanxi, grant number 2025CY-YBXM-059; the Shaanxi Key Project—Key Core Technology Research and Development—Future Emerging Industries Field, grant number 2025CG-GJHX-16; and the National Natural Science Foundation of China, grant number 62571306.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data are contained within the article.

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

References

1. Wang, Z.; Liu, H.; Qian, Z.; Wu, C.; Fang, Z.; Ma, X.; Chen, Z. Optimization and application of control strategies for hybrid heat pump heating systems in public buildings. *Appl. Therm. Eng.* **2026**, *302*, 131965. [[CrossRef](#)]
2. Mazhar, A.R.; Liu, S.; Shukla, A. A state of art review on the district heating systems. *Renew. Sustain. Energy Rev.* **2018**, *96*, 420–439. [[CrossRef](#)]

3. Mahmoud, M.; Ramadan, M.; Naher, S.; Pullen, K.; Olabi, A.G. The impacts of different heating systems on the environment: A review. *Sci. Total Environ.* **2021**, *766*, 142625. [[CrossRef](#)] [[PubMed](#)]
4. Gjoka, K.; Rismanchi, B.; Crawford, R.H. Fifth-generation district heating and cooling systems: A review of recent advancements and implementation barriers. *Renew. Sustain. Energy Rev.* **2023**, *171*, 112997. [[CrossRef](#)]
5. Østergaard, D.S.; Smith, K.M.; Tunzi, M.; Svendsen, S. Low-temperature operation of heating systems to enable 4th generation district heating: A review. *Energy* **2022**, *248*, 123529. [[CrossRef](#)]
6. Liu, L.; Zhang, H.; Liu, Y. A smart and transparent district heating mode based on industrial Internet of things. *Int. J. Energy Res.* **2021**, *45*, 824–840. [[CrossRef](#)]
7. Lee, D.; Lee, S.T. Artificial intelligence enabled energy-efficient heating, ventilation and air conditioning system: Design, analysis and necessary hardware upgrades. *Appl. Therm. Eng.* **2023**, *235*, 121253. [[CrossRef](#)]
8. Powell, K.M.; Sriprasad, A.; Cole, W.J.; Edgar, T.F. Heating, cooling, and electrical load forecasting for a large-scale district energy system. *Energy* **2014**, *74*, 877–885. [[CrossRef](#)]
9. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)] [[PubMed](#)]
10. Lee, J.; Cho, Y. National-scale electricity peak load forecasting: Traditional, machine learning, or hybrid model? *Energy* **2022**, *239*, 122366. [[CrossRef](#)]
11. Tan, Q.; Xue, G.; Xie, W. Cakformer: Transformer model for long-term heat load forecasting based on Cauto-correlation and KAN. *Energy* **2025**, *324*, 135460. [[CrossRef](#)]
12. Yang, C.; Wang, Y.; Lan, S.; Wang, L.; Shen, W.; Huang, G.Q. Cloud-edge-device collaboration mechanisms of deep learning models for smart robots in mass personalization. *Robot. Comput. Integr. Manuf.* **2022**, *77*, 102351. [[CrossRef](#)]
13. He, S.; Ren, J.; Wang, J.; Huang, Y.; Zhang, Y.; Zhuang, W.; Shen, S. Cloud-edge coordinated processing: Low-latency multicasting transmission. *IEEE J. Sel. Areas Commun.* **2019**, *37*, 1144–1158. [[CrossRef](#)]
14. Idowu, S.; Saguna, S.; Åhlund, C.; Schelén, O. Applied machine learning: Forecasting heat load in district heating system. *Energy Build.* **2016**, *133*, 478–488. [[CrossRef](#)]
15. Ali, A.K.; Al-Obaidi, M.A.; Ismail, A.H.; Mohammed, M.N.; Sadeq, A.M. Deep Reinforcement Learning Integrated PID for Hybrid Adaptive Control Approach in Automatic Voltage Regulation Systems. *Energies* **2026**, *19*, 2693. [[CrossRef](#)]
16. Khan, A.; Hussain, S.; Jawad, M.; Ahmad, W.; Jahanzaib, M. A real-time closed-loop control system for adaptive robotic welding: A step toward intelligent welding. *Int. J. Adv. Manuf. Technol.* **2025**, *141*, 5223–5245. [[CrossRef](#)]
17. Jiang, A.; Wei, H.; Deng, J.; Qin, H. Cloud-edge cooperative model and closed-loop control strategy for the price response of large-scale air conditioners considering data packet dropouts. *IEEE Trans. Smart Grid* **2020**, *11*, 4201–4211. [[CrossRef](#)]
18. Hansen, L.P.; Sargent, T.J. Robust control and model uncertainty. *Am. Econ. Rev.* **2001**, *91*, 60–66. [[CrossRef](#)]
19. Bhamare, D.; Zolanvari, M.; Erbad, A.; Jain, R.; Khan, K.; Meskin, N. Cybersecurity for industrial control systems: A survey. *Comput. Secur.* **2020**, *89*, 101677. [[CrossRef](#)]
20. Dong, Y.; Zhang, H.; Li, C.; Guo, S.; Leung, V.C.; Hu, X. Fine-tuning and deploying large language models over edges: Issues and approaches. *IEEE Commun. Mag.* **2026**, *64*, 94–101. [[CrossRef](#)]
21. Galloway, B.; Hancke, G.P. Introduction to industrial control networks. *IEEE Commun. Surv. Tutor.* **2012**, *15*, 860–880. [[CrossRef](#)]
22. Duan, Q.; Wang, S.; Ansari, N. Convergence of networking and cloud/edge computing: Status, challenges, and opportunities. *IEEE Netw.* **2020**, *34*, 148–155. [[CrossRef](#)]
23. Gu, H.; Zhao, L.; Han, Z.; Zheng, G.; Song, S. AI-enhanced cloud-edge-terminal collaborative network: Survey, applications, and future directions. *IEEE Commun. Surv. Tutor.* **2023**, *26*, 1322–1385. [[CrossRef](#)]
24. Wu, Y. Cloud-edge orchestration for the Internet of Things: Architecture and AI-powered data processing. *IEEE Internet Things J.* **2020**, *8*, 12792–12805. [[CrossRef](#)]
25. Benonysson, A.; Bøhm, B.; Ravn, H.F. Operational optimization in a district heating system. *Energy Convers. Manag.* **1995**, *36*, 297–314. [[CrossRef](#)]
26. Bazdyrev, A. Linear-time sequence modeling with selective state spaces for the Ukrainian language. In *System Analysis and Data Mining*; Zgurovsky, M., Pankratova, N., Eds.; Springer: Cham, Switzerland, 2026; Volume 609, pp. 250–265.
27. Zhang, H.; Zhu, Y.; Wang, D.; Zhang, L.; Chen, T.; Wang, Z.; Ye, Z. A survey on visual mamba. *Appl. Sci.* **2024**, *14*, 5683. [[CrossRef](#)]
28. Farha, Y.A.; Gall, J. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 3575–3584.
29. Wang, H.; Zhang, Z. TATCN: Time series prediction model based on time attention mechanism and TCN. In Proceedings of the 2022 IEEE 2nd International Conference on Computer Communication and Artificial Intelligence (CCAI), Beijing, China, 6–8 May 2022; pp. 26–31.
30. Hewage, P.; Behera, A.; Trovati, M.; Pereira, E.; Ghahremani, M.; Palmieri, F.; Liu, Y. Temporal convolutional neural (TCN) network for an effective weather forecasting using time-series data from the local weather station. *Soft Comput.* **2020**, *24*, 16453–16482. [[CrossRef](#)]

31. Daryanavard, S.; Porr, B. Closed-loop deep learning: Generating forward models with backpropagation. *Neural Comput.* **2020**, *32*, 2122–2144. [[CrossRef](#)] [[PubMed](#)]
32. Daryanavard, S. Real-Time Predictive Artificial Intelligence: Deep Reinforcement Learning for Closed-Loop Control Systems and Open-Loop Signal Processing. Ph.D. Thesis, University of Glasgow, Glasgow, UK, 2024.
33. Bimba, A.T.; Idris, N.; Al-Hunaiyyan, A.; Mahmud, R.B.; Shuib, N.L.B.M. Adaptive feedback in computer-based learning environments: A review. *Adapt. Behav.* **2017**, *25*, 217–234. [[CrossRef](#)]
34. Astrom, K.J. Adaptive feedback control. *Proc. IEEE* **1987**, *75*, 185–217. [[CrossRef](#)]
35. Ljung, L. Prediction error estimation methods. *Circuits Syst. Signal Process.* **2002**, *21*, 11–21. [[CrossRef](#)]
36. Niv, Y.; Schoenbaum, G. Dialogues on prediction errors. *Trends Cogn. Sci.* **2008**, *12*, 265–272. [[CrossRef](#)] [[PubMed](#)]
37. Che, Z.; Sun, J.; Na, H.; Yuan, Y.; Qiu, Z.; Du, T. A novel method for intelligent heating: On-demand optimized regulation of hydraulic balance for secondary networks. *Energy* **2023**, *282*, 128900. [[CrossRef](#)]

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