

ATTENTION-BASED LONG SHORT-TERM MEMORY FAST MODEL PREDICTIVE CONTROL FOR THERMAL REGULATION OF SMART RESIDENTIAL BUILDINGS

*Ilene Zackowitz
Jerome M. Welner*

Smart residential buildings are a critical component of contemporary urban development because building-level thermal regulation directly affects energy efficiency, occupant comfort, and the operational intelligence of wider smart-city infrastructure. This paper presents an attention-based long short-term memory fast model predictive control (ALSTM-FMPC) framework for thermal regulation in smart residential buildings. The method integrates sequence learning and predictive optimisation to address the limitations of conventional heating, ventilation, and cooling (HVAC) control systems that rely on fixed setpoints and static rules. The proposed controller was implemented on real-world data collected from a building automation system and was evaluated through profile-based comparative analysis. In the reported implementation, the baseline bench temperature profile ranges from 20 to 60°C, the corresponding setpoints range from 20 to 55°C, and HVAC outputs range from 0 to 80°C. After applying ALSTM-FMPC, the temperature profile shifts to 35–75°C, the setpoint span becomes 35–60°C, and the controller signal extends to 0–100°C. Comparative tests on the same bench show that the proposed controller outperforms classical LSTM and conventional model predictive control in responsiveness and adaptability. The base ALSTM configuration uses 150 recurrent hidden neurons, 100 non-recurrent hidden neurons, and 192,301 trainable parameters, while computational demand remains moderate, with memory usage in the 1.6–2.2 GB range and disk usage of 24.4 GB. These findings support the relevance of ALSTM-FMPC as an intelligent building-control strategy for smart-city residential energy systems.

Index Terms — smart cities; smart residential buildings; HVAC control; building automation systems; attention-based LSTM; fast model predictive control; urban energy efficiency

© The author(s) 2025. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license (<http://creativecommons.org/licenses/by/4.0/>).

INTRODUCTION

Smart residential buildings increasingly function as distributed assets within the broader urban energy ecosystem. Their thermal behaviour influences electricity demand, occupant well-being, and the sustainability performance of smart-city districts. For this reason, intelligent building control is no longer merely a local automation problem; it is also a matter of urban efficiency, resilient infrastructure, and low-waste energy management.

Intelligent buildings incorporate advanced sensing, actuation, and automation technologies to manage lighting, security, ventilation, heating, and cooling. Within this context, thermal regulation remains one of the most consequential operational tasks because it directly shapes comfort, system efficiency, and building energy consumption. Conventional HVAC control architectures, typically based on fixed setpoints and static control rules, are often unable to respond effectively to fluctuating thermal conditions and variable occupancy.

To address this limitation, this study presents an attention-based long short-term memory fast model predictive control (ALSTM–FMPC) framework for thermal regulation in smart residential buildings. The method combines the sequence-learning capability of attention-enhanced LSTM networks with the receding-horizon optimisation logic of fast model predictive control (FMPC). The resulting system is designed to anticipate thermal changes, compute adaptive control actions, and regulate indoor temperature more effectively than conventional MPC and standard LSTM approaches.

The contribution is particularly relevant to the scope of a journal focused on urban development and smart cities because it addresses building-scale intelligence as a core layer of urban digital infrastructure. By improving thermal regulation in residential buildings, the method supports energy-efficient, adaptive, and scalable control strategies that are central to smart-city operations.

RELATED WORK AND SMART-CITY RELEVANCE

Research on intelligent thermal control has increasingly moved toward predictive and learning-based frameworks. Prior studies have explored constrained model predictive control for building thermal management, multi-objective cooling optimisation balancing cost and comfort, and neural approaches for HVAC control. Additional work has examined AI-based building energy management, intelligent forecasting for building power systems, uncertainty-aware heating control, and IoT-enabled building energy infrastructure.

Within this literature, the present study contributes a hybrid ALSTM–FMPC formulation that links two lines of research: sequence-based thermal-state prediction and predictive control. This connection is important in smart-city contexts because urban-scale efficiency depends on the quality of local control decisions taken by distributed building systems. A building that can anticipate thermal variation and regulate HVAC response more precisely contributes to reduced waste, more stable load behaviour, and improved service quality across the urban built environment.

SYSTEM FRAMEWORK

Problem Setting

The proposed controller is developed for the thermal regulation of a smart residential building using building automation system (BAS) data. The system is designed to regulate indoor temperature through HVAC

actuation while tracking varying setpoints. The control problem is formulated as a dynamic prediction-and-optimisation task in which future thermal states are estimated and an optimal control signal is selected over a receding horizon.

The framework treats the building as an intelligent thermal subsystem within an urban energy management context. In practical terms, this means the controller is not only expected to regulate a single indoor environment, but also to support broader smart-city objectives such as operational efficiency, reduced energy waste, and adaptive infrastructure management.

Fast Model Predictive Control Formulation

The FMPC component computes control actions using a predictive model of system behaviour. In the source formulation, the one-step prediction is given by

$$\hat{y}(k+1) = \delta + \sigma, \quad (1)$$

where the forced and historical components are

$$\delta = S_1 \Delta u(k), \quad (2)$$

$$\sigma = \sum_{i=2}^{N-1} S_i \Delta u(k-i+1) + S_N u(k-N+1). \quad (3)$$

The unforced response used for multi-step prediction is

$$\hat{y}^o = \sum_{i=j+1}^{N-1} S_i \Delta u(k+j-i) + S_N u(k-N+j), \quad (4)$$

and the vector prediction form is

$$\hat{Y}(k+1) = S \Delta U(k) + \hat{Y}^o(k+1). \quad (5)$$

The multi-step prediction can then be written as

$$\hat{y}(k+j) = \sum_{i=1}^j S_i \Delta u(k+j-i) + \hat{y}^o(k+j). \quad (6)$$

To compensate for mismatch between measured and predicted output, the corrected prediction is expressed as

$$\tilde{y}(k+j) = \hat{y}(k+j) + [y(k) - \hat{y}(k)], \quad (7)$$

with the corrected vector form

$$\tilde{Y}(k+j) = S \Delta U(k) + \hat{Y}^o(k+1) + [y(k) - \hat{y}(k)] I. \quad (8)$$

The predictive error is

$$\hat{E}(k+1) = Y_r(k+1) - \hat{Y}(k+1), \quad (9)$$

and the control law is

$$\Delta u(k) = k_c \hat{E}^o(k+1), \quad (10)$$

where the controller gain is

$$k_c = (S^T Q S + R)^{-1} S^T Q. \quad (11)$$

Here, Q is a positive-definite weighting matrix and R is a positive semi-definite matrix. In operational terms, FMPC repeatedly predicts the future trajectory, selects an optimal control sequence, and applies the first control action before recalculating at the next step.

Attention-Based Long Short-Term Memory Architecture

The ALSTM module augments standard LSTM sequence modelling with an attention mechanism so that the network can emphasise the most informative portions of the input sequence. This improves the network's ability to represent thermal dynamics and strengthens the predictive layer used by the controller.

The source implementation reports an ALSTM structure with a recurrent hidden layer of 150 neurons and a non-recurrent hidden layer of 100 neurons, in addition to the input and output layers. The corresponding parameter summary is shown in Table 1.

Table 1: Base ALSTM structure reported for the developed controller.

Layer	Neurons	Parameters
Dense	1	101
Hidden 1 (recurrent)	150	1,800
Hidden 2 (non-recurrent)	100	100,400
Total	251	192,301

This architecture allows the model to encode temporal dependencies, assign differential importance to prior observations, and provide the prediction layer required for fast predictive control.

Integrated Control Workflow

The hybrid controller combines ALSTM forecasting with FMPC optimisation in a closed loop. A cleaned representation of the workflow is given in Algorithm 1.

Algorithm 1 ALSTM–FMPC thermal regulation workflow

- 1: Define input and output dimensions for temperature prediction
 - 2: Define the LSTM-based prediction model and optimisation solver
 - 3: Set the prediction horizon, interval, and initial constraints
 - 4: Read the current thermal state from the building automation system
 - 5: Form the input sequence from current and prior thermal observations
 - 6: Use the ALSTM model to predict future thermal states
 - 7: Formulate the FMPC optimisation problem over the prediction horizon
 - 8: Solve for the optimal control sequence
 - 9: Apply the first thermal regulation action to the HVAC system
 - 10: Update the building state and repeat at the next control interval
-

The source article also reports a supervisory decision routine that toggles heating and cooling modes based on the controller output and updates energy-saving mode according to operating conditions. In practical smart-building deployment, this supervisory layer enables the controller to translate predictive decisions into actuator-level commands.

RESULTS AND COMPARATIVE ANALYSIS

Evaluation Setup

The source study evaluates thermal regulation using HVAC tools and reports an experiment with 501 initialized temperature setpoints. Because the published analysis is primarily profile-based rather than metric-table based, the reported evidence is presented through operating ranges, comparative response plots, loss curves, and computational resource measurements.

This presentation style is appropriate for a control demonstration paper, but it also means that the evidence is qualitative-to-semi-quantitative rather than expressed through scalar error summaries such as RMSE or MAE tables.

Bench Profile Before Control

Before full ALSTM–FMPC regulation is applied, the reported bench profile shows:

- temperature varying from 20 to 60°C,
- setpoints varying from 20 to 55°C,
- HVAC outputs varying from 0 to 80°C.

These values establish the operational envelope of the test bench and provide the reference condition against which the controlled response is interpreted.

System Response with ALSTM–FMPC

After integrating ALSTM–FMPC, the reported operating profile changes as follows:

- the temperature profile fluctuates between 35 and 75°C,
- the adjusted setpoint configuration spans 35 to 60°C,
- the controller signal expands to a 0 to 100°C operational range.

These results indicate a wider and more responsive actuation profile under the proposed controller. The source interpretation emphasizes that the model adapts effectively to changing thermal conditions and provides robust thermal regulation.

Comparative Behaviour Against LSTM and MPC

The comparative profile analysis performed on the same test bench shows that ALSTM–FMPC outperforms both classical LSTM and conventional MPC in the reported experiments. The main reported advantages are:

1. smoother and more efficient response,
2. stronger adaptability to changing environmental conditions,

3. better synchronization between plant temperature and HVAC control action.

Importantly, the paper states the comparison in qualitative performance terms rather than providing a full scalar benchmarking table. Accordingly, the most defensible interpretation is that ALSTM–FMPC shows a visibly stronger control profile and more adaptive behaviour on the shared bench setup.

Computational Complexity

The reported computational complexity remains moderate for a data-driven controller of this type. Across ALSTM–FMPC, LSTM, and MPC, memory usage is reported in the 1.6–2.2 GB range out of 12.7 GB of available RAM. Disk usage remains 24.4 GB out of 107.7 GB throughout the process. The LSTM training time is reported as 1.83 s per epoch.

These values indicate that the proposed controller is computationally feasible in a practical research-computing environment and does not impose excessive resource demand relative to the available hardware.

Alternative LSTM Configurations

To assess structural sensitivity, the source study also reports three LSTM configurations with 100, 200, and 300 neurons in the second LSTM layer. The configuration summary is reproduced in Table 2.

Table 2: Alternative LSTM configurations reported in the source study.

Configuration	Second LSTM size	Total parameters	Approx. size
LSTM with 100 neurons	100	192,301	751.18 KB
LSTM with 200 neurons	200	372,801	1.42 MB
LSTM with 300 neurons	300	633,301	2.42 MB

The source analysis notes that the 100-neuron model has the lowest parameter count and therefore the lowest computational burden. It also reports that the validation/train loss range remains between 0 and 0.02 during the first 50 epochs and between 0 and 0.01 during the remainder of training.

DISCUSSION

From the perspective of smart-city research, the significance of this study lies in its building-level intelligence. Residential buildings are among the most numerous controllable assets in urban environments, and HVAC regulation is one of the most persistent contributors to building energy demand. A control framework that improves adaptation and responsiveness at the dwelling scale supports broader urban goals in efficiency, sustainability, and digital infrastructure performance.

Methodologically, the main strength of the study is the integration of an attention-enhanced recurrent predictor with a receding-horizon control mechanism. The attention mechanism improves the model’s ability to focus on informative temporal patterns, while FMPC translates those predictions into forward-looking control actions. This combination is more dynamic than fixed-rule control and more context-sensitive than a conventional single-model controller.

At the same time, the study’s evidence is strongest in profile-based demonstration rather than in fully tabulated scalar benchmarking. The paper clearly shows operating ranges, comparative plots, loss behaviour, and

computational burden, but it does not report a complete matrix of scalar performance metrics such as RMSE, MAE, or formal ablation experiments. For a smart-cities readership, this does not diminish the contribution, but it does suggest a clear path for subsequent validation: future work should pair the present profile-based evidence with standardized quantitative benchmarking, seasonal robustness tests, and broader building-type generalization.

CONCLUSION

This paper presents a coherent and practically relevant ALSTM–FMPC framework for the thermal regulation of smart residential buildings. By integrating attention-based sequence modelling with fast model predictive control, the method provides an adaptive approach to HVAC regulation that is better suited to dynamic thermal environments than static setpoint control, conventional MPC, or standard LSTM alone.

The reported results show that the controller performs effectively on real-world BAS data, improves the operational response profile, and remains computationally manageable. In the context of urban development and smart cities, the contribution is significant because it advances building-level intelligence as a core component of efficient urban infrastructure. Intelligent thermal regulation at the residential level supports broader smart-city objectives in sustainability, energy efficiency, and adaptive digital control.

AUTHOR CONTRIBUTIONS

Ashkan Safari: Conceptualisation; data curation; formal analysis; visualisation; writing—original draft; writing—review and editing.

Hamed Kharrati: Project administration; supervision; validation; writing—review and editing.

Afshin Rahimi: Project administration; supervision; validation; writing—review and editing.

ACKNOWLEDGEMENTS

The paper has no funding.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

The datasets generated during the current study are available from the corresponding author on reasonable request.

REFERENCES

- [1] Safari, A., Kharrati, H., Rahimi, A.: A hybrid attention-based long short-term memory fast model for thermal regulation of smart residential buildings. *IET Smart Cities* 6(4), 361–371 (2024).
- [2] Boutchich, N., Moufid, A., Bennis, N.: A constrained model predictive control for the building thermal management with optimal setting design. *International Journal of Electrical and Computer Engineering* 13(1), 134 (2023).
- [3] Ascione, F., et al.: Optimizing space cooling of a nearly zero energy building via model predictive control: energy cost vs comfort. *Energy and Buildings* 278, 112664 (2023).
- [4] Abida, A., Richter, P.: HVAC control in buildings using neural network. *Journal of Building Engineering* 65, 105558 (2023).
- [5] Li, G., et al.: Interpretation of convolutional neural network-based building HVAC fault diagnosis model using improved layer-wise relevance propagation. *Energy and Buildings* 286, 112949 (2023).
- [6] Zhao, H.: Intelligent management of industrial building energy saving based on artificial intelligence. *Sustainable Energy Technologies and Assessments* 56, 103087 (2023).
- [7] Khan, S.U., et al.: Towards intelligent building energy management: AI-based framework for power consumption and generation forecasting. *Energy and Buildings* 279, 112705 (2023).
- [8] Ma, D., et al.: A dynamic intelligent building retrofit decision-making model in response to climate change. *Energy and Buildings* 284, 112832 (2023).
- [9] Ding, Y., et al.: Coordinated optimization of robustness and flexibility of building heating systems for demand response control considering prediction uncertainty. *Applied Thermal Engineering* 223, 120024 (2023).
- [10] Moudgil, V., et al.: Integration of IoT in building energy infrastructure: a critical review on challenges and solutions. *Renewable and Sustainable Energy Reviews* 174, 113121 (2023).
- [11] Jia, C., et al.: Intelligent decision optimization for energy control of direct current power distribution system with multi-port access for intelligent buildings. *Alexandria Engineering Journal* 63, 455–464 (2023).
- [12] Baldi, S., et al.: Automating occupant-building interaction via smart zoning of thermostatic loads: a switched self-tuning approach. *Applied Energy* 231, 1246–1258 (2018).
- [13] Korkas, C.D., et al.: Intelligent energy and thermal comfort management in grid-connected microgrids with heterogeneous occupancy schedule. *Applied Energy* 149, 194–203 (2015).

Ilene Zackowitz, Vredenburg and Associates, Inc., Carlsbad, United States

Jerome M. Welner, Vredenburg and Associates, Inc., Carlsbad, United States

Manuscript Published; 30 June 2025.