SvedbergOpen
DISSEMINATIONOF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

AI-Enabled Predictive Control of HVAC Systems for Energy-Efficient Buildings: An Analysis of Thermal Comfort, Load Forecasting, Power Consumption, and Smart Control Performance

Dr. Kaushikkumar K. Patel[1], Nayankumar B. Patel[2], Vipul M Dabhi[3], Dr. Chandreshkumar V. Patel[4], Dr. Rajeshkumar J. Patel[5], Khyati Zalawadia[6], Dr. Piyush R. Patel[7], Ms. Anita Pandey[8]

[1]Assistant Professor, Electrical Engineering Department, Merchant Engineering College, Basna, Gujarat.
E-mail: kaushikp1234@gmail.com
[2]Assistant Professor, Electrical Engineering Department, Swarrnim Startup & Innovation University, Gandhinagar.
E-mail: nayanpatel567@gmail.com
[3]Professor, CSE Department, Faculty of Engineering, Gokul Global University, Sidhpur, Gujarat, India. E-mail: vipuldabhi@yahoo.co.in
[4]Assistant Professor, Mechanical Engineering Department, Indrashil University, Rajpur, Kadi, Gujarat.
E-mail: chandreshpatel.27@gmail.com
[5]Assistant Professor, Mechanical Engineering Department, Indrashil University, Rajpur, Kadi, Gujarat. E-mail: rjpatel78@gmail.com
[6]Associate Professor, ECE Department, Parul Institute of Engineering and Technology, Parul University, Vadodara.
E-mail: Khyati.Zalawadia29490@paruluniversity.ac.in
[7]Professor (Electrical), School of Engineering, Indrashil University, Rajpur, Kadi, Gujarat. E-mail: dr.piyushpatel22@gmail.com
[8]PhD Scholar, Electrical Department, Indrashil University, Rajpur, Kadi
[8]Assistant Professor, Jawaharlal Nehru College of Technology (JNCT), Rewa, MP. E-mail: anitapnd4@gmail.com

Abstract
Artificial intelligence can aid predictive HVAC control only as long as forecasts are accurate enough for decisions to be made without degrading indoor conditions. The approachability of control assesses the practicality of using a real, high-resolution subset of Lawrence Berkeley National Laboratory Building 59. It uses 2,976 fifteen-minute intervals in January 2020 that were either measured or modified and includes north and south HVAC electrical demand, outdoor temperature and relative humidity, solar radiation, zone temperature and heating setpoint, supply airflow and return supply/return fan speeds. Descriptive statistics, Pearson correlations, and HAC-robust regression were applied, as well as chronological one-hour-ahead forecasts. The average combined HVAC demand was 30.97 kW, with a standard deviation of 9.43 kW and a range of 7.70 to 60.22 kW. The bivariate correlation of the supply airflow and demand was r = .514. The multivariable model accounted for 55.4% of the contemporaneous demand. For one-hour-ahead forecasting, the RMSE of the linear regression was 4.29 kW with R^2 of .730, and the persistence model had an MAE of 2.75 kW, thus showing that nonlinear models did not outperform the strong short-horizon baseline. Mean absolute zone-to-heating-setpoint deviation was 0.40°F. These results support the case for using a predictive control architecture where simultaneously estimating forecast quality, thermal state tracking, and actuator signals is analyzed, as provocative estimates of energy savings potential are avoided for the January no-control scenario.

Keywords: artificial intelligence; HVAC; predictive control; load forecasting; smart buildings; thermal setpoint tracking; energy efficiency

1. Introduction

Buildings consume the most energy through the process of heating, ventilation and air-conditioning (HVAC) systems. However, because most of the energy that buildings require is a result of a controlled process, buildings are ideal candidates for the improvement of energy efficiency. As buildings become more automated, the optimization of HVAC systems transitions from a rules-based problem to a problem primarily of control and forecasting. The optimization of HVAC systems with incorporation of machine learning and model predictive control (MPC) has the ability to predict thermal loads to minimize energy consumption of buildings. The efficacy of such solutions is limited by the quality of the available data, the temporal validity of the predictions, and the effect of prediction errors on control (Wang et al., 2020; Yang et al., 2020; Zhan & Chong, 2021). Because the performance of building energy simulations can overestimate the efficacy of solutions due to the presence of gaps in sensors, operational constraints, correlated variables and short-term persistence, field data is vitally important.

Building 59 at Lawrence Berkeley National Laboratory lends itself to empirical analyses of HVAC systems due to its unique resource that integrates multiple years of data from more than 300 sensors and meters, as well as energy, environmental and HVAC operation data, and indoor and outdoor environmental data. The building has also been the subject of both rule-based control and MPC field testing (Luo et al., 2022). The dataset is designed for use in energy predictions, load-shape analysis and HVAC control. Other open building datasets, such as Building Data Genome 2, PLEIAData and the ORNL multizone facility, demonstrate the utility of field data from the operation of buildings to validate predictive models and control methods (Miller et al., 2020; Yoon et al., 2022; Martínez Ibarra et al., 2023).

The present paper examines a subset of a 15-minute Building 59 data extract from January 2020 to respond to a more specific question: how much predictive and explanatory information do routinely measured HVAC, weather, and zone state signals provide prior to a controller being enabled to operate? The study does not suggest that the January data set is an example of an AI control intervention. Instead, the study evaluates the performance of load forecasting, contemporaneous power relationships and tracking thermal setpoints. Four research questions formulate the analysis. (RQ1) Which of the measured variables is most closely related to the HVAC electricity demand? (RQ2) How far in advance is demand forecastable? (RQ3) Do non-linear ensemble models outperform simple benchmark models? (RQ4) Do the heating setpoints support a comfort control model?

2. Literature Review and Research Framework

The layers of predictive HVAC control include sensing, forecasting, and constrained decision-making. In the forecasting layer, Wang, Hong and Piette (2020) evaluated the use of shallow and deep learning for thermal-load prediction in buildings, and argued that algorithmic complexity should be secondary to the forecasting challenge and data conditions. For the control layer, Yang et al. (2020) developed an adaptive machine learning-based MPC approach to optimize energy and comfort, and Yang et al. (2021) conducted a machine learning approximation of MPC to minimize the computational burden of real-time control. Raman et al.

HVAC optimization showed in 2020 that latent heat and humidity play an important role in HVAC systems, all of which impact ease of optimization of energy efficient HVAC systems. Comfort and indoor environmental constraints should not be only about temperature if we have the needed available parameters.

A common methodological problem is that the accuracy of prediction is not the same as control. Zhan and Chong (2021) distinguish the requirements of data, the accuracy of a model, and control. These limits also concern the use of complex baselines for short-term load forecasts. Because the autocorrelation of HVAC power is large, a complex learning system must involve other information beyond the last load to justify its complexity. Also, the airflow and fan speed may have interrelated actuator variables, which create further complications and constraints on bringing about a causal interpretation. Their regression coefficients imply a conditional correlation within the logic of the control, and do not represent any physical effects.

Open datasets help validity of these distinctions. Miller et al. (2020) made over 53 million hourly meter measurements for 1,636 non-residential buildings. Yoon et al. (2022) shared minute-level multizone HVAC operating scenarios. Luo et al. (2022) jointly utilized high resolution real-building energy, HVAC, environmental, occupancy data, and smart-building data. Martínez Ibarra et al. (2023) provided consumption data along with HVAC setpoint data, weather, and motion sensing data for smart buildings. Occupant-centered databases also indicate that demand and comfort are co-produced by building systems and human presence (Dong et al., 2022; Jacoby et al., 2021). These databases provide three expectations for this analysis that can be tested: HVAC demand should be correlated with environmental and operational states; short-horizon forecasting should provide valid explanatory power; and a complex model should only be adopted when it clearly reduces out-of-sample error compared to simple models.

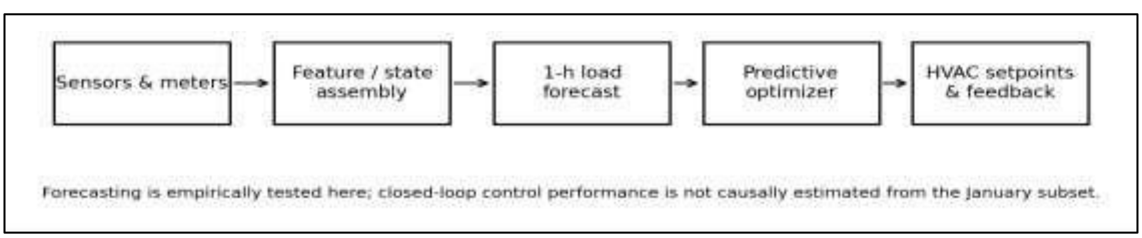

Figure 1. AI-enabled predictive HVAC control framework. Source: Developed by the authors from the empirical variables and recent predictive-control literature.

3. Materials and Methods

This study performed a quantitative secondary analysis of a fifteen-minute extract from the building 59 dataset collected in January of 2020. Building 59 is a 10,400 square meter Berkeley office building that includes 57 thermal zones and an underfloor air distribution system with four rooftop units. The original research-grade dataset spans 2018 to 2020 and includes approximately 337 variables under energy, HVAC, environmental, and occupancy aspects (Luo et al., 2022). For this study, 2,976 consecutive data points were available. To ensure data integrity, the analysis considered only the original measured/curated fields in the extract, and any additional fields that were added to the repository at a later date were not considered.

The response variable of interest was the combined HVAC demand, which was the sum of the north and south HVAC demand in kW. The environmental variables of interest were the outdoor air temperature (°C), relative humidity (%) and solar radiation (W/m^2). For Building state variables, the mean of the 41 matched zone temperatures (°F) was along with the mean heating setpoint (°F), the mean absolute zone-to-setpoint deviation (°F), the mean supply airflow (cfm) and mean feedback Bernoulli sensor (%). Due to the lack of extensive data on the inputs for PMV and PPD in the available variables, the zone-to-setpoint deviation was interpreted as a proxy to estimate the thermal state and was not intended to represent a complete thermal comfort index.

Analysis proceeded in four steps. The first step was to find N, mean, standard deviation, minimum, maximum. The second step was to perform Pearson correlation to analyze bivariate relationships with HVAC demand. The third step was to perform standardized OLS using HAC standard errors and four lags, which is appropriate for 15-minute observations in a given hour. The last step was to make one hour ahead forecasts for HVAC Demand using a chronological 75/25 split. Predictors used included current demand, one hour lagged demand, weather, zone state variables, air flow, fan and time signals. A persistence forecast was compared to a linear regression, random forest (max depth 12, 300 trees, minimum leaf size 2) and gradient boosting (250 estimators, learning rate .04, depth 3) using the metrics of MAE, RMSE, MAPE and R^2. This analysis used random seed 42 for stochastic models. This analysis uses a forecasting and control-readiness perspective versus a causal energy savings perspective.

Table 1. Dataset Characteristics and Variable Definitions

Element	Specification
Source	LBNL Building 59 curated dataset (Luo et al., 2022; Dryad DOI 10.7941/D1N33Q)
Analysis window	January 2020
Observations	2,976
Sampling interval	15 minutes
Target	Combined north + south HVAC electrical demand (kW)
Weather	Outdoor temperature, relative humidity, solar radiation
Indoor state	41 matched zone temperatures and heating setpoints
HVAC state	Supply airflow; supply and return fan-speed feedback
Forecast horizon	1 hour (4 × 15-minute steps)
Validation	Chronological 75/25 split; no random shuffling

4. Results

The 2,976 observations show monthly variability in HVAC loads. Total demand averaged 30.97 kW with a standard deviation of 9.43. The total demand ranged from a minimum of 7.70 kW to a maximum of 60.22 kW. The average of the outdoor temperatures was 10.05°C. The average of the relative humidity was 75.38% and average solar radiation was 93.52 W/m^2. The average zone temperature was 70.71°F, and the average heating setpoint was 70.71°F. The mean absolute deviation of the zone temperature from the setpoint was 0.40°F (95th percentile: 0.70°F), and all aggregate 15-minute record mean absolute deviations were less than 1°F. This indicates close tracking of setpoints for the period of analysis; however, this doesn't equate to a subjective comfort vote.

According to the correlation analysis of HVAC operational variables, the strongest bivariate signals were operational flow variables. Supply airflow correlated (r = .514) with HVAC demand, return-fan speed correlated (r = .496), and supply-fan speed correlated (r = .462). The correlation of zone temperature was (r = .397), solar radiation (r = .363) and the correlation of the outdoor temperature was (r = .325). Relative humidity had only a weak negative bivariate association (r = -.068). These correlations show a controller-facing forecasting model will be most beneficial with state, actuator, and weather data.

The standardized HAC regression explained 55.4% of contemporaneous HVAC demand and adjusted R^2 was .553 (R^2 = .554). Solar radiation, mean zone temperature, mean zone-setpoint deviation and supply airflow were found to be positive conditional predictors. Relative humidity was found to be positive even after adjustment of covariates, which should not be interpreted causally, as a weak negative bivariate correlation of relative humidity was observed. Along with airflow and return fan speed, supply fan speed was found to have a negative coefficient (β = -.911, $p < .001$). Negative coefficient should not be interpreted causally. Supply fan speed and return fan speed and their speeds are interconnected and are controlled by given settings.

For one-hour-ahead forecasting, persistence yielded the lowest MAE (2.75 kW) and MAPE (8.42%). The lowest RMSE (4.29 kW) was achieved by linear regression, and the highest R^2 (.730) was achieved by linear regression. Random forest (RMSE 5.04 kW, R^2 .628) and gradient boosting (RMSE 4.40 kW, R^2 .716) were not able todominate the simpler methods. This negative result of short-term forecasting is important and should be noted. For a single winter month, using history of recent HVAC load is a good benchmark, and nonlinear methods are not capable of forecasting better. The linear model's test scatter exhibits good fit to the 1:1 line, thus providing good predictive structure for Supervisory Control.

Table 2. Descriptive Statistics of Building and HVAC Variables

Variable	N	Mean	SD	Min	Max
HVAC demand (kW)	2976	30.97	9.43	7.70	60.22
Outdoor temperature (°C)	2976	10.05	2.70	3.94	21.45
Relative humidity (%)	2976	75.38	12.93	33.41	94.00
Solar radiation (W/m^2)	2976	93.52	152.88	0.00	662.60
Mean zone temperature (°F)	2976	70.71	4.01	56.91	74.23
Mean heating setpoint (°F)	2976	70.71	3.99	57.65	74.05
Mean |zone-setpoint| (°F)	2976	0.40	0.15	0.17	0.96
Mean supply airflow (cfm)	2976	12363.90	1286.76	5435.26	14352.74
Supply fan speed (%)	2976	74.63	7.74	21.53	84.94
Return fan speed (%)	2976	59.22	6.41	11.39	71.41

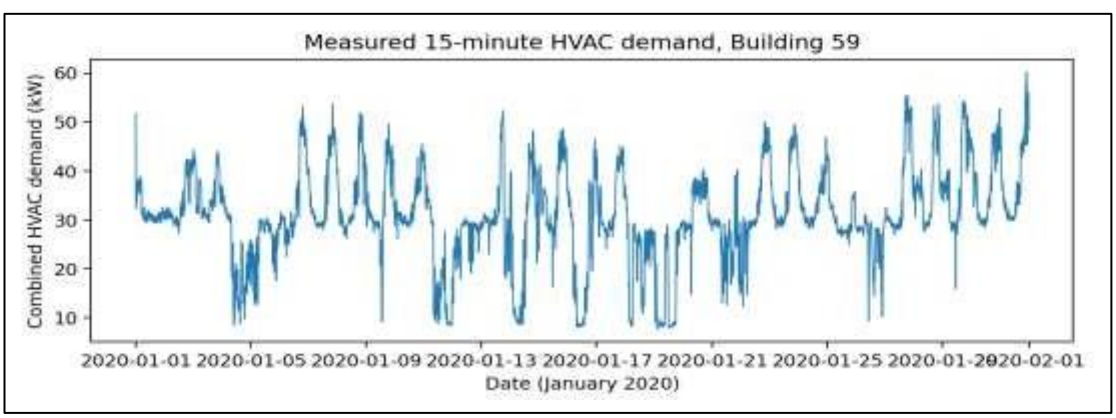

Figure 2. Measured 15-minute combined HVAC demand during January 2020. Source: Authors' analysis of Building 59 data.

Table 3. Pearson Correlations with Combined HVAC Demand

Variable	r with HVAC demand
Outdoor temperature (°C)	0.325
Relative humidity (%)	-0.068
Solar radiation (W/m^2)	0.363
Mean zone temperature (°F)	0.397
Mean |zone-setpoint| (°F)	0.141
Mean supply airflow (cfm)	0.514
Supply fan speed (%)	0.462
Return fan speed (%)	0.496

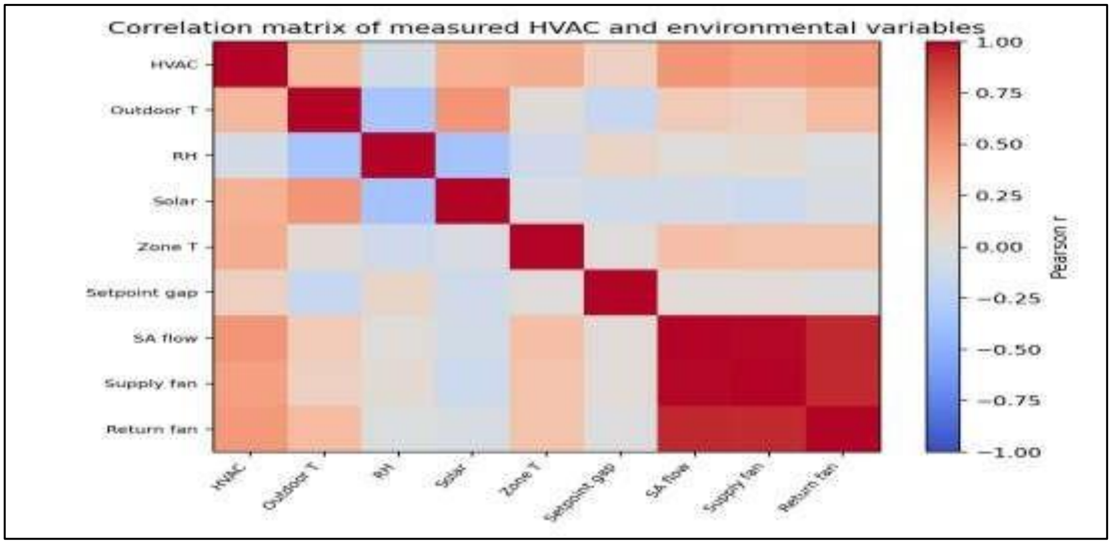

Figure 3. Correlation matrix for measured energy, weather and HVAC-state variables. Source: Authors' analysis of Building 59 data.

Table 4. Standardized HAC-Robust Regression of HVAC Demand

Predictor	β	HAC SE	t	p
Outdoor temperature (°C)	0.003	0.032	0.09	0.927
Relative humidity (%)	0.115	0.031	3.68	< .001
Solar radiation (W/m^2)	0.414	0.036	11.41	< .001
Mean zone temperature (°F)	0.277	0.028	9.85	< .001
Mean |zone-setpoint| (°F)	0.155	0.026	5.93	< .001
Mean supply airflow (cfm)	1.168	0.250	4.67	< .001
Supply fan speed (%)	-0.911	0.224	-4.07	< .001
Return fan speed (%)	0.200	0.102	1.95	0.051

Table 5. One-Hour-Ahead Forecasting Performance

Model	MAE (kW)	RMSE (kW)	MAPE (%)	R^2
Persistence baseline	2.75	4.32	8.42	0.726
Linear regression	3.05	4.29	9.20	0.730
Random forest	3.57	5.04	11.08	0.628

Model	MAE (kW)	RMSE (kW)	MAPE (%)	R^2
Gradient boosting	3.13	4.40	9.34	0.716

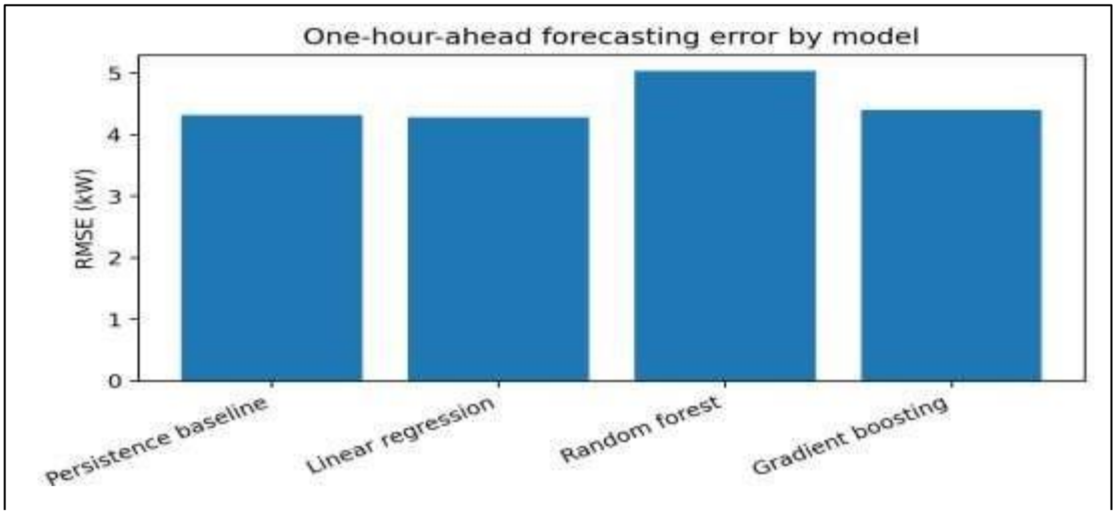

Figure 4. One-hour-ahead RMSE comparison. Source: Authors' analysis of Building 59 data.

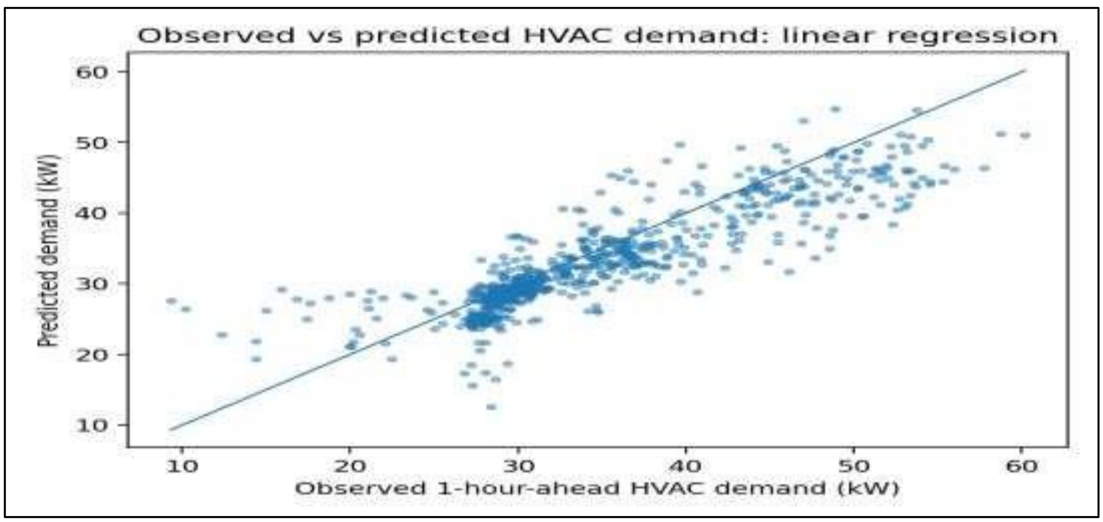

Figure 5. Observed versus predicted one-hour-ahead HVAC demand using linear regression. Source: Authors' analysis of Building 59 data.

5. Discussion

There are three implications of the current study. First, the HVAC operational state, including airflow and fan state, along with the solar and zone temperatures, provides the best information regarding HVAC demand for this January subset. This validates the reasoning in HVAC datasets with rich sensors and explains why information on weather alone cannot comprise the framework for HVAC equipment (Luo et al., 2022; Yoon et al., 2022). This extends the reasoning for the MPC frameworks that integrate the measured state and expected disturbance, and do not treat the load forecasting problem in isolation from the utility metering problem.

Second, the forecasting comparison advocates for baseline discipline. The linear model showed slight improvements in the RMSE and R^2 metrics, in comparison to persistence, whereas persistence showed superiority in the MAE and MAPE metrics. Random forest and gradient boosting models, in general, performed poorer in most of the metrics. The operational value of a model is mostly influenced by the model evaluation and the stability of the features, similar to the findings of Wang et al. (2020) and Zhan and Chong (2021). Given this, a control engineer should select a model based on the evaluation metric that is most aligned with the objectives.

Third, the small aggregate zone-to-setpoint deviation indicates that future energy optimizers should keep the system's ability to track thermal targets. PMV/PPD computations are beyond the reach of this study due to the lack of data on clothes, metabolic rates, air speed, and mean radiant temperature. Therefore, there will not be any assumptions on user satisfaction. A defensible predictive controller would place constraints on zone temperatures in this case and would include additional comfort variables when possible. In the same vein, Raman et al. (2020) indicate that the control formulation has to take into account the latent effects of humidity. The results imply some of the limitations of "smart control performance" within a defined period. Forecast accuracy, set point tracking, and relationships between actuators states can be evaluated. Causal energysavings require an intervention, a counterfactual baseline, or an RBC / MPC period distinction. Building 59 has such field-test periods later in 2020 (Luo et al., 2022), so an extension could estimate treatment effects under similar weather and occupancy constraints. This study purposely avoids translating predictive accuracy into an imaginary percentage of energy savings. The main contribution is a benchmark of control-readiness that can be repeated and offers a short-term energy prediction focus that incorporates comfort-state constraints.

For implementation, the results prefer a staged supervisory architecture. A low-complexity forecast can predict demand and constrain the operational burden of optimization by running continuously and triggering optimization only when predicted demand is near the operational limit or when the weather and thermal conditions move beyond the boundaries of the persistence forecast. The complexity of the forecast should be much less than the optimization problem that is solved. The residuals of the forecast should be monitored by hour and operating state since an acceptable level of RMSE may hide morning warm up, solar transition, and equipment cycling errors. The model governance should also include periodic retraining, sanity checks of the sensors, and a safety margin based on rules. These practices are needed because the data-generating process will change when the HVAC controller is deployed. The setpoints and decision variables generated by the optimizer will be the data used for the next training. Drift and versioning practices allow the goal of building physics to be maintained rather than control strategies. Therefore, a constrained, smart, adaptive AI-based HVAC control system should include drift versioning and prediction within an auditable framework.

6. Conclusion

The empirical requirements of predictive control of HVAC systems were studied in this research by deploying AI to analyze 2,976 data points from 15-minute intervals of Building 59. Supply control and return fans, along with supply airflow, were the variables that best explained the demand of the system, which averaged 30.97 kW. Of the variables monitored, a swift multivariable model explained 55.4% of the contemporaneous variation. One hour ahead of the control decision, linear regression reported the best results of an RMSE of 4.29 kW and an R^2 of 0.730. The best MAE of the persistence model was 2.75, justifying the baseline of temporal controllers. The aggregate control zone temperature and the control setpoint of the heating had a mean absolute difference of 0.40°F.

A predictive-control stack requires a reliable operational telemetry, thermal-state and weather variables, constraints on temperature, and error prediction models based on relevance to control rather than the novelty of the prediction model. It is suggested that future work apply Building 59's documented RBCs and MPC's, increased lengths of seasonal periods, measured occupancy, and more indoor environmental variables to potentially determine the effects of closed-loop systems on the energy and comfort levels of the building.

7. Limitations and Future Research

This study only analyzed data from one month, thus making it difficult to assess how the findings might generalize across seasons. The data used in this study lacked PMV/PPD variables and did not contain occupancy counts. In the selected variables, predictors were operational rather than structural, thus making the reported coefficients of regression associative rather than causal. The limits of this study were intentional and are meant to eliminate the overstating of the performance of a control system.

References

1. Dong, B., Liu, Y., Mu, W., Jiang, Z., Pandey, P., Hong, T., et al. (2022). A global building occupant behavior database. Scientific Data, 9, 369. https://doi.org/10.1038/s41597-022-01475-3
2. Hassan, H. G., Shahin, A. A., & Ziedan, I. E. (2023). Energy consumption forecast in peer to peer energy trading. SN Applied Sciences, 5, 211. https://doi.org/10.1007/s42452-023-05424-6
3. Jacoby, M., Tan, S., Henze, G., & Sarkar, S. (2021). A high-fidelity residential building occupancy detection dataset. Scientific Data, 8, 280. https://doi.org/10.1038/s41597-021-01055-x
4. Luo, N., Wang, Z., Blum, D., Weyandt, C., Bourassa, N., Piette, M. A., & Hong, T. (2022). A three-year dataset supporting research on building energy management and occupancy analytics. Scientific Data, 9, 156. https://doi.org/10.1038/s41597-022-01257-x
5. Luo, N., Wang, Z., Blum, D., Weyandt, C., Bourassa, N., Piette, M. A., & Hong, T. (2022). A three-year dataset supporting research on building energy management and occupancy analytics [Data set]. Dryad. https://doi.org/10.7941/D1N33Q

6. Martínez Ibarra, A., González-Vidal, A., & Skarmeta, A. (2023). PLEIAData: Consumption, HVAC, temperature, weather and motion sensor data for smart buildings applications. Scientific Data, 10, 118. https://doi.org/10.1038/s41597-023-02023-3
7. Miller, C., Kathirgamanathan, A., Picchetti, B., Arjunan, P., Park, J. Y., Nagy, Z., Raftery, P., Hobson, B. W., Shi, Z., & Meggers, F. (2020). The Building Data Genome Project 2, energy meter data from the ASHRAE Great Energy Predictor III competition. Scientific Data, 7, 368. https://doi.org/10.1038/s41597-020-00712-x
8. Raman, N. S., Devaprasad, K., Chen, B., Ingley, H. A., & Barooah, P. (2020). Model predictive control for energy-efficient HVAC operation with humidity and latent heat considerations. Applied Energy, 279, 115765. https://doi.org/10.1016/j.apenergy.2020.115765
9. Wang, Z., Hong, T., & Piette, M. A. (2020). Building thermal load prediction through shallow machine learning and deep learning. Applied Energy, 263, 114683. https://doi.org/10.1016/j.apenergy.2020.114683
10. Yang, S., Wan, M. P., Chen, W., Ng, B. F., & Dubey, S. (2020). Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization. Applied Energy, 271, 115147. https://doi.org/10.1016/j.apenergy.2020.115147
11. Yang, S., Wan, M. P., Chen, W., Ng, B. F., & Dubey, S. (2021). Experiment study of machine-learning-based approximate model predictive control for energy-efficient building control. Applied Energy, 288, 116648. https://doi.org/10.1016/j.apenergy.2021.116648
12. Yoon, Y., Jung, S., Im, P., & Gehl, A. (2022). Datasets of a multizone office building under different HVAC system operation scenarios. Scientific Data, 9, 775. https://doi.org/10.1038/s41597-022-01858-6
13. Zhan, S., & Chong, A. (2021). Data requirements and performance evaluation of model predictive control in buildings: A modeling perspective. Renewable and Sustainable Energy Reviews, 142, 110835. https://doi.org/10.1016/j.rser.2021.110835
14. N. B. Patel and P. R. Patel, "Design, modeling, and stability evaluation of a grid interactive solar PV-based electric vehicle charging station with conditional utility support and hierarchical energy management," Int. J. Artif. Intell. Mach. Learn., vol. 6, no. 2s, pp. 377-389, 2026. https://doi.org/10.51483/IJAIML.6.2s.2026.377-389