Data Requirements and Performance evaluation for Control-Oriented Models: a Case Study on Internal Heat Gain

Sicheng Zhan[1], Yue Lei[1], Yuan Jin[2], Da Yan[2], Adrian Chong[1]
[1]Department of Building, School of Design and Environment, University of Singapore, Singapore
[2]Department of Building Science, School of Architecture, Tsinghua University, China

Abstract

Model predictive control has shown its great potential in improving building performance but is bottlenecked by the difficulty in constructing control-oriented models. This paper attempts to bridge the knowledge gaps between data requirements, model capability, and control performance by conducting simulation experiments. Considering its importance in building operations, internal heat gain was selected as the subject. Actual data was fed into high-fidelity models to emulate model identification and control application of real buildings. It was shown that higher input granularity resulted in higher prediction accuracy, especially for longer prediction horizons. However, the superiority was diminished regarding the control performance. Lastly, critical discussions on control-oriented modeling and further research directions were provided to promote the MPC application.

Key Innovations

- Connected input data granularity, model capability, and control performance in MPC
- Holistically evaluated different inputs for internal heat gain in control-oriented models
- Examined control-oriented models for prediction, extrapolation, and control optimization

Practical Implications

Lower prediction error of control-oriented models does not necessarily mean better control performance. It is preferable to have an input for internal heat gain estimation in control-oriented models, especially for longer prediction horizons.

Introduction

Model predictive control (MPC) has been tested in buildings since the 1990s (Henze et al., 1997). Able to incorporate different system dynamics and disturbances, it shows great potential for improving the building operation performance (Drgoňa et al., 2020). However, there is still a limited number of actual implementations (Benndorf et al., 2018). One major barrier is the cost of obtaining a sufficient control-oriented model, which is the foundation of the entire framework. 10% of model discrepancy could lead to 5% more energy cost and 100% more comfort violation (Bengea et al., 2011).

Past studies have proposed various modeling methods for MPC, such as physics-based models (Sturzenegger et al., 2015), data-driven models (Ferreira et al., 2012), and hybrid models (Dong and Lam, 2014). Resistor-capacitor (RC) model is a typical hybrid model. It has been popular because it shares the advantages of the other two model types, and it is suitable for optimization (Zhan and Chong, 2021). Therefore, it is adopted in this study.

In addition to the difficulty of modeling building dynamics, it is also hard to determine how accurate a model should be and to assess the modeling effort in advance (Killian and Kozek, 2016). Due to the diversity across buildings, studies catering to specific buildings are hardly generalizable. Yet, since the models are ultimately used for control, it is critical to connect model characteristics to control performance. Only few researches focused on this and explored the influence of different model configurations. Picard et al. (2017) varied the number of states in the model and spotted the minimum amount that guarantees the control performance. Blum et al. (2019) tested several practical factors for model identification and concluded with a couple of modeling suggestions. Arroyo et al. (2020) found that a centralized multi-zone model and a simplified single-zone model achieved similar prediction accuracy, except the former got more robust control performance.

With many other unexplored factors and the complicated intrinsic mechanism, MPC application in buildings still has knowledge gaps between data requirements, model capability, and control performance. Particularly, how to quantify the model capability so that it can inform the ultimate control performance remains obscure. Thus, the model capability refers to how well a model can serve the control optimization, which requires not just high prediction accuracy against normal operation data, but also solid extrapolation capability. Besides, it is unclear how data inputs, together with model structure, affect model identification and the resulting model capability. Accordingly, the type of input data is varied for RC models to investigate the interrelationships in this study. Meanwhile, the models are extensively tested for extrapolation capability and control performance.

Occupant behavior is a major source of uncertainty in building operations (Tian et al., 2018). Correspondingly, internal heat gain is a significant component in building loads prediction, especially when the model is used for control optimization (Wang et al., 2019). Measurement granularity refers to how accurately the measurement represents the object. Various levels of granularity for internal heat gain inputs have been used in control-oriented models. Some models used no input for this and expected the model to incorporate the uncertainty (Ferreira et al., 2012). Under the shortage of real-time measurement, Váňa et al. (2014) approximated the internal heat gain with a ratio-based design schedule. At a higher level, the profiles were estimated based on the electricity consumption trend (De Coninck and Helsen, 2016). Among the measurements that are not commonly available in building operations, CO_2 concentration is the most used for control-oriented modeling (Maasoumy et al., 2014). Considering its importance and the absence of a consensus, internal heat gain is selected as the case study object in this paper.

Objectives

This paper attempts to shed some light on the three research questions when designing control-oriented RC models for MPC in buildings:

1. How to quantify model capability and its relationship with the control performance?
2. What is the optimal combination of input data granularity and model complexity to achieve better model capability?
3. Which level of internal heat gain measurements is required to build a reliable model?

A simulation-based approach is adopted to tackle these questions. In the rest of this paper, we first introduce the simulation framework and the experiment setup. Then, the results of model identification and control experiments are presented respectively. Subsequently, the research questions are discussed based on observations from the experiment results. Upon concluding this study, we also point out directions for further investigation.

Methodology

Simulation framework

As shown in figure 1, the simulation framework consists of emulation models, control-oriented RC models, and corresponding MPC controllers. In each experiment, the emulator served as a real building to generate synthetic datasets. The training dataset was then used to identify the RC models with different setups. The identification results were evaluated against separate testing datasets. The identified models were also adopted in the MPC controllers, the optimized actions of which were applied back to the emulators for control evaluation. Through experiment design, the prediction capability and control performance of a variety of RC models were evaluated. Thereby, the relationships between data, model, and control were established.

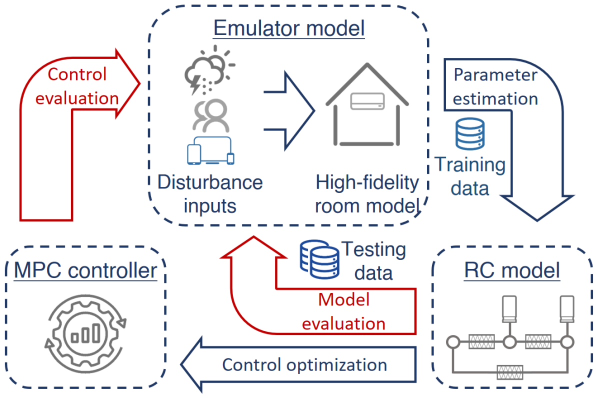

Figure 1: Structure of the simulation framework.

- **Emulator model** The experiments were carried out in an office room in Singapore. Therefore, BESTEST Case600 with light-weighted construction (ASHRAE, 2007b) was adopted to represent the typical building thermal characteristics in the tropics. The room was conditioned by a fan coil unit that supplies air at constant 13°C. The supply air flow rate was controlled against the room temperature setpoint by a PI controller.

  To better account for the uncertainty of internal heat gain in actual building operation, we collected the occupant number and plug load of an office for 3 months. The measured profiles, instead of design schedules, were applied to emulate occupant and equipment heat gain. The actual meteorological year weather data was utilized for external disturbance. The emulators are built using the Modelica Buildings Library[1] and simulated for 3 months using JModelica[2].
- **RC model** The RC models represented the simplified thermal dynamics of the room with a set of parameters θ (resistances, capacitances, and heat gain coefficients). The cooling power was estimated with the supply air flow rate m_{flow}. Emulation data in the first 3 days with the room temperature setpoint of 24 °C was used to estimate the parameters via Non-Linear Programming (NLP). Equation 1 defines the estimation problem, where x is the states, u is the inputs, d

[1] https://simulationresearch.lbl.gov/modelica/
[2] https://jmodelica.org/

is the disturbances, $t0$ and $t1$ are the start and end time of training data, and the lower and upper bounds (θ^{lb} and θ^{ub}) of parameter values are based on prior knowledge. 10 days were randomly picked from the rest days in the 3 months to evaluate the identified models. To examine the extrapolation capability of the RC models, Root Mean Squared Error (RMSE) was calculated against testing data with 22, 24, and 26 °C room temperature setpoint.

$$\theta = argmin \sum_{t0}^{t1} (T_{room} - \widehat{T}_{room})^2 \qquad (1)$$
$$s.t. \quad \widehat{T}_{room} = f(x, u, d, \theta)$$
$$\theta^{lb} \leq \theta \leq \theta^{ub}$$

- **MPC controller** The MPC optimization problem was formulated as equation 2 to maintain the room temperature setpoint by controlling the supply air flow rate. The quadratic objective function penalized the cooling power and thermal discomfort with weights q_u and q_t. Cooling power was represented by the control input m_{flow} and thermal discomfort was quantified by the room temperature deviation during the operating hours (7am to 7pm). The minimization was subject to the nominal air flow rate. The prediction and control horizon was half an hour and the internal states were estimated using the Unscented Kalman Filter. The control task was kept simple to eliminate other affecting factors and to study the effect of model mismatch. The control performance was also evaluated under the 3 room temperature setpoints (22, 24, and 26 °C) on the 10 randomly-picked days. Since the cooling load is directly related to T_{room}, lower room temperature always comes with higher cooling energy consumption. Therefore, the performance was only quantified by accumulated setpoint violation (°C.h).

$$J = \int_{t_0}^{t_0+30min} q_u(m_{flow})^2 + q_t(\widehat{T}_{room} - T_{setpoint})^2$$
$$s.t. \quad 0 \leq m_{flow} \leq 0.6 \qquad (2)$$

Design of experiments

Preliminary experiments were conducted to investigate how the temporal factors of training data affect model accuracy. It showed that time intervals smaller than 15 minutes were sufficient to capture the dynamics. Regarding the training data length, longer periods slightly reduced the error but drastically increased the computation time. Concisely, 3 days of training data with 15-minute intervals reached a balance between model accuracy and computation time. Therefore, it was adopted for further experiments.

Due to the poor envelope thermal properties, external heat gain took up a large percentage of cooling load in the BESTEST emulator. With the stricter regulations on building thermal performance, internal heat gain plays an increasingly important role in modern office buildings (Papadopoulos, 2016). Thus, another emulator was created by increasing the envelope thermal resistance (referred to as the insulated emulator in the rest of this paper). Internal heat gain constitutes 10-30% of cooling load in the BESTEST emulator and 30-50% in the insulated emulator.

The second variation in experimental design is the RC model structures. Figure 2 displays the three levels of tested complexities. R1C1 lumps the entire room into a capacitor $C1$ and a resistor $R1$ connecting the outdoor temperature node. Cooling power and internal heat gain are directly delivered to the room temperature node, so is solar heat gain but with coefficient a. R3C2 considers the wall a separate capacitor C_{wall} and two resistors R_{wi} and R_{we}. Another resistor R_{infil} is added to model the infiltration. Solar heat gain is now sent to the wall temperature node with coefficient a_{wall}. R5C3 models the heat transfer through floor with 4 extra parameters C_{floor}, R_{fi}, R_{fe}, and a_{floor}. Since the floor is exposed to a constant ground temperature, this configuration is expected to capture a separate dynamics different from the wall.

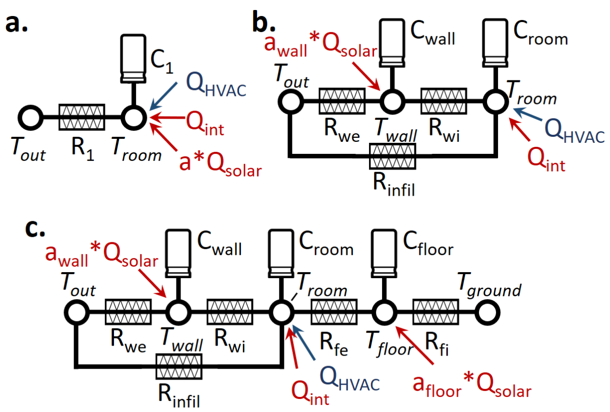

Figure 2: Structure of the R1C1 (a), R3C2 (b), and R5C3 (c) models. Arrows annotate the input of disturbances and control actions.

Table 1 summarizes the design of experiments. Five levels of measurement granularity were tested for internal heat gain: `no input`, `design schedule`, `plug load`, `CO₂ ppm`, and `ideal measurement`, `Design schedule` referred to the standard occupancy schedule of offices according to ASHRAE (2007a). The capacity Cap was left to be estimated. `Plug load` and `CO₂ ppm` were outputs of the emulator. The coefficients a_{plug} and a_{CO_2}, as well as the offset b, were to be estimated. `Ideal measurement` was the exact internal heat gain, which is barely measurable in practice. Generally, higher granularity comes with higher data acquisition cost.

Table 1: Summarized design of experiments.

Subject	Variations
Emulator	BESTEST, Insulated (higher internal load percentage)
RC model sturcture (parameters)	R1C1 ($R1, C1, a$), R3C2 (R_{wi}, $R_{we}, R_{infil}, C_{room}, C_{wall}, a_{wall}$), R5C3 ($R_{wi}, R_{we}, R_{fi}, R_{fe}$, $R_{infil}, C_{room}, C_{wall}, C_{floor}$, a_{wall}, a_{floor})
Internal heat gain input (parameters)	No input, Design schedule(Cap), Plug load(a_{plug}, b), CO_2 ppm(a_{CO_2}, b), Ideal measurement

Results

Identification and prediction

Figure 3 visualizes the RMSE results that quantify model capability. Each bar chart compares the RMSE of the 15 alternative RC models against the corresponding emulator and testing condition. The bars under certain testing conditions represent the average RMSE of the 10 randomly-picked testing days.

It can be seen that RMSE is kept lower than 1°C in most cases, which means these simplified RC structures were able to track the basic trend of the room dynamics. The change of RMSE against different datasets is similar for the two emulators: reached the lowest for the training dataset, reasonably increased under 22 and 24°C testing conditions, and further rose under 26°C. This indicates that the models have some extrapolation capability, but is weakened when the external heat gain and cooling power is smaller.

In terms of comparing RC model structures, more complex models achieved lower training error with no exemption. This is related to the stronger capability of fitting the data brought by the larger number of parameters. However, while R1C1 resulted in the largest RMSE in all testing cases, R3C2 and R5C3 showed fluctuating and not significantly different testing RMSE across the cases. This is not unexpected because heat flow through the floor is relatively insignificant, and therefore is harder to capture than heat flow through the wall (ceiling included).

Figure 4 explains the situation by disaggregating the heat flow of models respectively. According to the second subplot, the order of heat flow intensity from the most to the least significant is wall, internal heat gain, floor, and infiltration. Comparing the prediction results with the emulation data, it appears that neither model correctly predicted the disaggregated heat flow. However, both models captured the aggregated heat flow and properly predicted the room temperature.

Regarding the alternative inputs for internal heat gain, `plug load` and `CO2 ppm` led to similar results with `ideal measurement`, better than `no input` and `design schedule`. The improvement was augmented for the insulated emulator, where the percentage of internal heat gain was enlarged.

Figure 5 plots the predicted internal heat gain of R3C2 models with alternative inputs. Because the number of occupants has a very high correlation with the plug load in offices, both `plug load` and `CO2 ppm` followed the primary trend of internal heat gain. Nevertheless, it is worth noting that `plug load` missed several peaks created purely by occupants, and that `CO2 ppm` was delayed and smoothed, serving as a low pass filter. `design schedule` correctly modeled the rise in the daytime and the drop around noon, but failed to catch the minor variations and

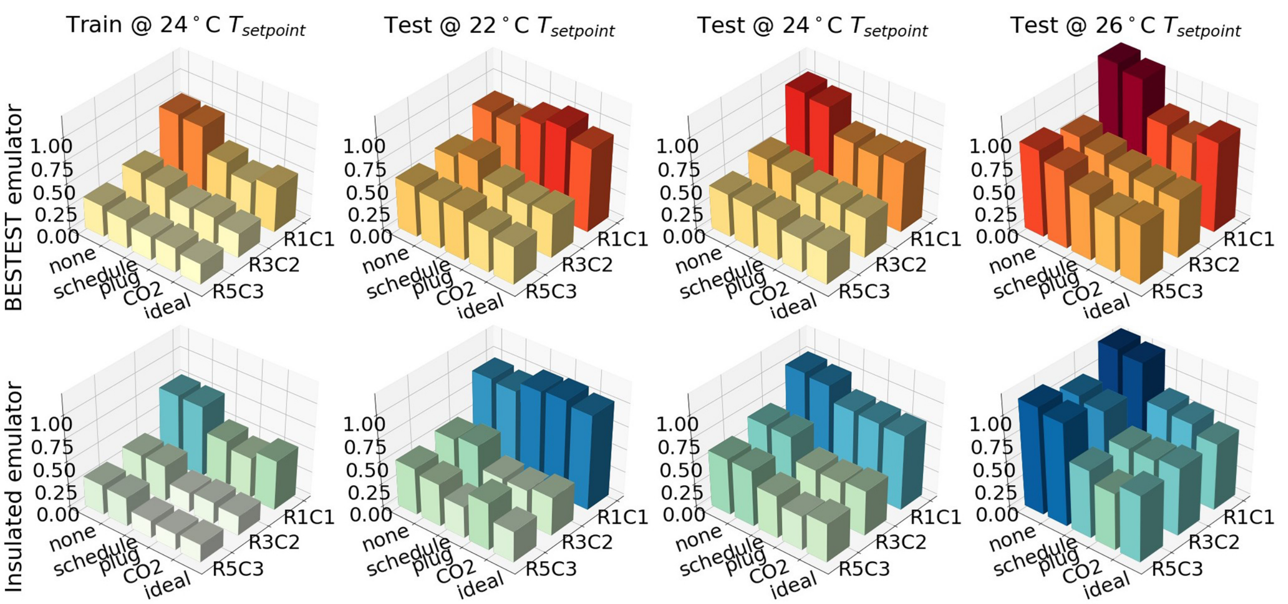

Figure 3: RMSE (°C) of different RC models for the two emulators under training and the three testing conditions. X-axis is the 5 levels of internal heat gain input, Y-axis is the 3 model complexities, and Z-axis is the RMSE. Darker colors and higher bars represent larger RMSE and worse performance.

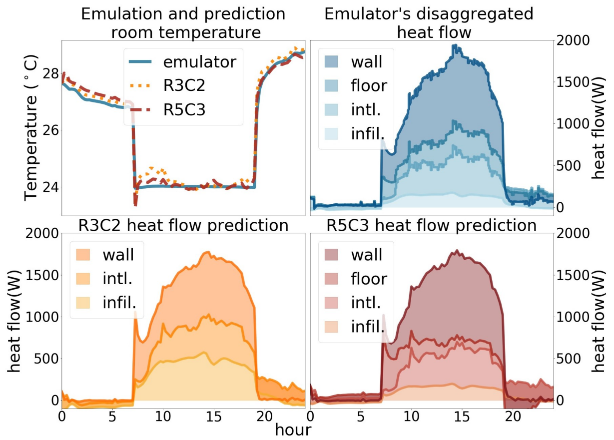

Figure 4: Comparison of R3C2 and R5C3 (with ideal internal heat gain inputs) prediction results on an exemplary testing day of the insulated emulator with 24°C $T_{setpoint}$. The disaggregated heat flow stacked up to the total heat gain of the room.

the baseload. Understandably, all three inputs underestimate the scale for the insulated emulator in response to the overall reduced cooling load. Noticeably, the `design schedule` model underestimated more intensely, which also explains the amplified difference in RMSE.

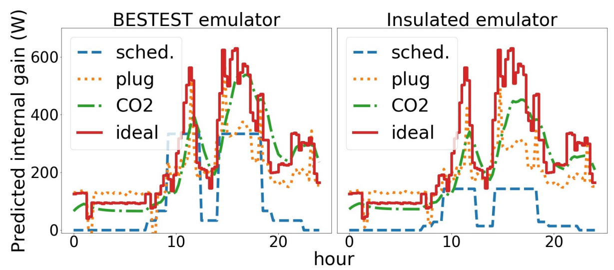

Figure 5: Predicted internal heat gain of R3C2 models with different inputs for the two emulators on one testing day.

Control experiment

Figure 6 visualizes the RC models' average accumulated setpoint violation (°C.h) from the control experiment. It is distinguishable that the 26°C $T_{setpoint}$ was violated the most compared with the other two, disregarding model complexity and input granularity. It is also distinct that the accumulated violation of R1C1 models exceeded the other two and frequently exploded. Though slightly, R5C3 models perform better than R3C2 in 16 out of 20 cases under 22 and 24°C $T_{setpoint}$. The alteration of input granularity presents no clear pattern, except mildly improved the performance of R5C3 models for the insulated emulator.

Two representative days were selected to manifest the effect of model complexities and input on control (figure 7). The first subplot displays control results of the three model complexities for the BESTEST emulator against 24°C $T_{setpoint}$. All three models had a tendency to overcool at the beginning of the day and undercool in the afternoon. While R3C2 and R5C3 quickly eliminated the offset, R1C1 resulted in over 1°C deviation. The second subplot compares the R5C3 model with five different inputs for the insulated emulator against 26°C $T_{setpoint}$. It is obvious that `no input` and `design schedule` led to larger temperature deviation than the other three. However, even in this most diverse case, the difference was only around half °C.

Discussion

Relating model and control performance

Linking figure 6 back to figure 3, it is found that RMSE reflected the major changes in control performance. For example, the severer setpoint violation of R1C1 models and at 26°C $T_{setpoint}$ both aligned with the higher RMSE. However, it is also worth noting the mismatches when investigating the impact of alternative inputs and complexities. The advantage of having more informative inputs for prediction was not prominent in the control performance. Meanwhile, the control benefit brought by more complex models was not shown in RMSE. To summarize, once the RMSE is lower than a certain level, it cannot distinguish the models' control capability. This finding somewhat agrees with Blum et al. (2019) on that RMSE is a necessary but not sufficient condition for control.

In contrast to RMSE of the entire day, Žáčeková et al. (2014) evaluated the models' prediction capability with RMSE only over the prediction horizon. Considering that this manner ties more closely to the control situation, the concept was adopted to formulate the control-oriented RMSE (CoRMSE) as in equation 3. In the equation, p is the horizon length (2 in this case), and n refers to the length of testing data. Similarly, this metric was used to evaluate the identified models under different conditions.

$$CoRMSE = \left(\frac{1}{p(n-p)} \sum_{i=1}^{n-p} \sum_{k=1}^{p} (\widehat{T}_{room,i+k|i} - T_{room,i+k})^2 \right)^{\frac{1}{2}} \tag{3}$$

Figure 8 presents the CoRMSE results. Overall, CoRMSE is considerably smaller than RMSE, indicating a relaxed requirement on the models' predictive capability. When applying MPC, the requirement is similarly lowered as the horizon is much shorter (half an hour) and there is feedback for state correction. Looking into the variation among alternative RC models, CoRMSE better matched the control results than RMSE. Taking the alteration of model inputs for illustration, no certain input stood out with remarkably smaller CoRMSE. As for R5C3 models of the insulated emulator, CoRMSE also gradually decreased as the setpoint violation did, but with an exemption of $\texttt{CO}_2$ `ppm`. Inspecting the results of R3C2

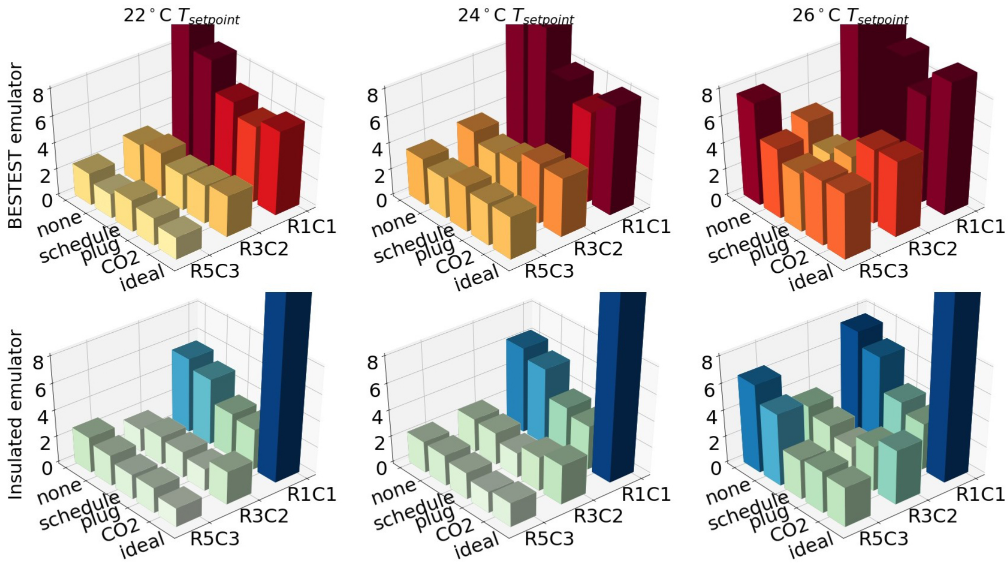

Figure 6: Accumulated setpoint violation (°C.h) of RC models for the two emulators against the three $T_{setpoint}$. The bars taller than 8 °C.h reflect unacceptable performance and are cut off to better visualize the others.

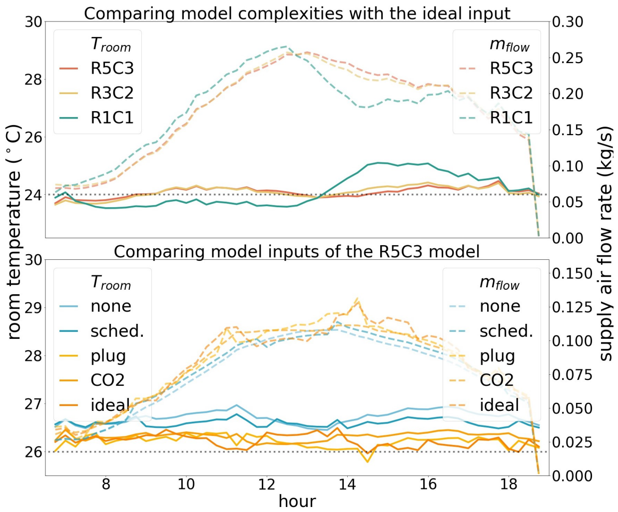

Figure 7: Control actions and results on selected testing days. The complexity comparison is for the BESTEST emulator against 24° C $T_{setpoint}$, and the input comparison is for the insulated emulator against 26° C $T_{setpoint}$.

and R5C3 models, CoRMSE showed the same relativeness as the control results in most cases, especially with 26°C $T_{setpoint}$.

Generally, CoRMSE is a more promising indicator of the control performance than RMSE. However, the explainability is still limited. For example, while the temporal lag could cause the increased errors of `CO2 ppm`, the corresponding control results are still relatively good. In addition, the exploded setpoint violation of R1C1 models with ideal input also contradicts the acceptable prediction errors (especially for the insulated emulator). The explosion is caused by huge fluctuations due to mispredictions, which means the model capability was insufficient but was not exposed by the testing datasets. These results call for further research to pursue a more robust indicator.

The importance of internal heat gain

There is evidence for the merit of having higher granularity inputs for internal heat gain, especially when predicting over a longer term. This improvement is more obvious for the insulated emulator with the more substantial role of internal heat gain. Whereas, when the prediction horizon is shortened to half an hour, the evaluation results of R3C2 and R5C3 models are less differentiated. This is because more complex models partially account for the effect of internal heat gain with other components in the model. When receding the horizon, the discrepancy is less accumulated and quickly rectified.

More representative inputs upgrade the model not just by properly predicting the internal heat gain, but also by helping identify the rest components in the model. Given more accurate information from the input, the internal and external heat gain can be decomposed. Consequently, the dynamics caused by outdoor conditions is more explicit, and parameters of the room thermal properties can be better estimated. This is another reason for the lower RMSE.

On the other hand, it is also worth noting the degree of freedom brought by the extra parameters of `design schedule` (Cap), `plug load` (a_{plug}, b), and `CO2 ppm` (a_{CO_2}, b). As the RC structures simplify the room's thermal dynamics, these parameters could be compromised to better fit the training data. This justifies why `plug load` and `CO2 ppm` sometimes got lower RMSE than `ideal measurement`. Yet, if the model has too low explainability like R1C1, the degree of freedom could harm the identification. One trace is the higher RMSE of R1C1 models with `plug load` and `CO2 ppm` against 22°C $T_{setpoint}$ in figure 3.

To sum up, the increase of measurement granularity is preferable for models' prediction capability, but is not outstanding for control performance. The relatively short prediction horizon when doing control is the main reason. Therefore, the benefit is expected to be more conspicuous if the horizon is prolonged. For MPC in typical offices, `design schedule` is a good enough estimate, but higher granularity is recommended to improve the robustness.

Directions for further investigation

Apart from the aforementioned issues, how to quantify the relationships between data requirements, model capability, and control performance remains unresolved. This paper selected internal heat gain for case study, while there are other data categories to be studied. For instance, Reynders et al. (2014) suggested including envelope conditions to improve the identifiability of RC models, which could be costly in practice. Hence, it is important to examine the trade-off between the cost and benefit regarding the ultimate goal of control.

A preliminary test is done by keeping the same experiment framework but including the floor temperature to identify R5C3 models. Equation 4 defines the new multi-input-multi-output (MIMO) problem. It is expected to improve from the single-output identification as more information is provided. On the contrary, both the model accuracy and control performance turned out to be a little worse in most cases. One speculation from the data-driven perspective is that the new output constrained the optimization more strictly so that the solution became sub-optimal for the room temperature prediction. In this sense, higher measurement granularity for room conditions does not necessarily contribute to better performance.

Further, this experiment exerts doubt on what makes a reliable control-oriented RC model. Several studies advocated the importance of RC parameters constituting the true building thermal properties (Reynders et al., 2014; Sourbron et al., 2013). The MIMO-identified models serve as a counterexample against that stand. These models better tracked the change of floor temperature and thereby modeled the disaggregated heat flow through different components. By contrast, the single-output-identified model unphysically overestimated the floor temperature but predicted the room temperature more accurately in the testing cases. Thus, it is possible for RC models to better approximate the building dynamics without having the identical physical interpretation. Along this line, a new evaluation metric other than the conventional prediction error could be developed.

$$\theta = argmin \sum_{t0}^{t1} \Big((T_{room} - \widehat{T}_{room})^2 + (T_{floor} - \widehat{T}_{floor})^2 \Big)$$
$$s.t. \quad \widehat{T}_{room} = f(x, u, d, \theta)$$
$$\theta^{lb} \leq \theta \leq \theta^{ub} \tag{4}$$

Lastly, this study is restricted to cooling an office room in the tropical climate. Other boundary conditions such as heating in a mild climate can be explored in the future. More complex control scenarios are also to be tested for better generalizability.

Conclusion

This study aims to provide a more comprehensive and in-depth understanding of data requirements and performance evaluation for RC models. A simulation framework was established to investigate different model complexities and alternative inputs for internal heat gain. The prediction capability and control performance of the identified RC models were

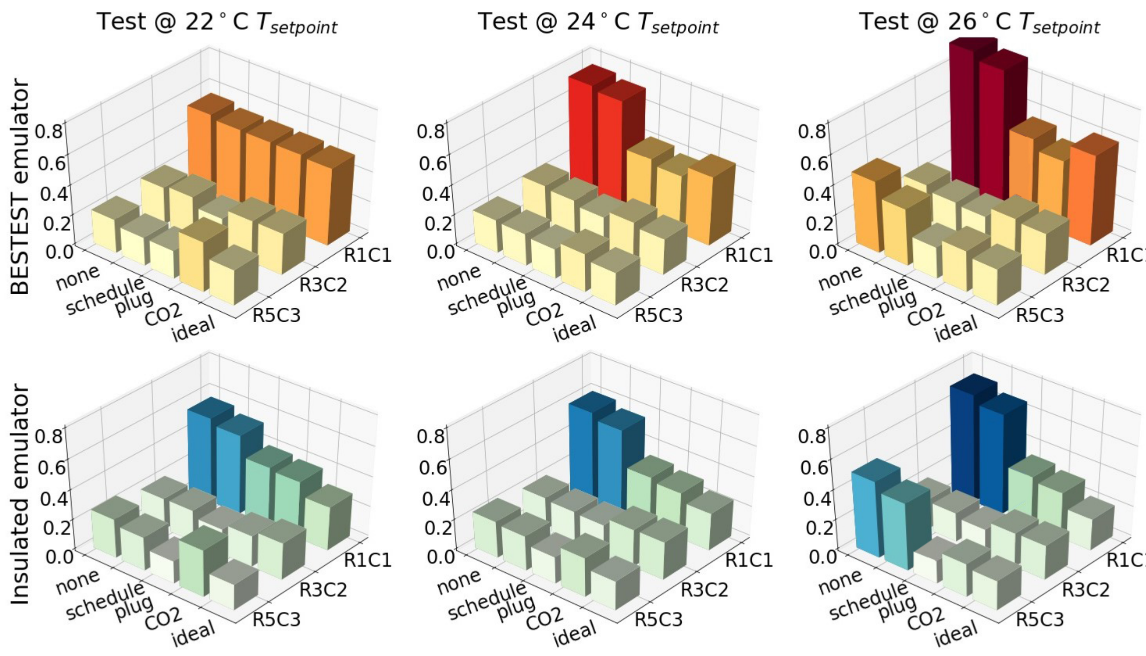

Figure 8: Control-oriented RMSE (°C) of RC models for the two emulators under the three testing conditions. Z-axis is the CoRMSE. Darker colors and higher bars represent larger CoRMSE and worse performance.

evaluated under a series of testing conditions.

Different configurations resulted in diverse model capabilities. However, the control performance did not fully reflect the predictive superiority of some RC models. This is because the variation of prediction error was not significant enough to take effect in control with the relatively short prediction horizon.

It was also demonstrated that higher input granularity enhanced model prediction capability. The benefit first came from the better representation of internal heat gain. Besides, given proper model complexities, the identification was also improved.

Accordingly, recommendations for input selection and model construction are given. Further, directions of future research in the area are pointed out.

References

Arroyo, J., F. Spiessens, and L. Helsen (2020). Identification of multi-zone grey-box building models for use in model predictive control. *Journal of Building Performance Simulation 13*(4), 472–486.

ASHRAE (2007a). *90.1. Energy Standard for Buildings Except Low-Rise Residential Buildings*.

ASHRAE (2007b). *140: Standard Method of Test for the Evaluation of Building Energy Analysis Computer Program*.

Bengea, S., V. Adetola, K. Kang, M. J. Liba, D. Vrabie, R. Bitmead, and S. Narayanan (2011). Parameter estimation of a building system model and impact of estimation error on closed-loop performance. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 5137–5143. IEEE.

Benndorf, G. A., D. Wystrcil, and N. Réhault (2018). Energy performance optimization in buildings: A review on semantic interoperability, fault detection, and predictive control. *Applied Physics Reviews 5*(4), 041501.

Blum, D., K. Arendt, L. Rivalin, M. Piette, M. Wetter, and C. Veje (2019). Practical factors of envelope model setup and their effects on the performance of model predictive control for building heating, ventilating, and air conditioning systems. *Applied energy 236*, 410–425.

De Coninck, R. and L. Helsen (2016). Practical implementation and evaluation of model predictive control for an office building in brussels. *Energy and Buildings 111*, 290–298.

Dong, B. and K. P. Lam (2014). A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting. *Building Simulation 7(1)*, 89–106.

Drgoňa, J., J. Arroyo, I. C. Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, et al. (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control 50*, 190 – 232.

Ferreira, P., A. Ruano, S. Silva, and E. Conceicao (2012). Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and buildings 55*, 238–251.

Henze, G. P., R. H. Dodier, and M. Krarti (1997). Development of a predictive optimal controller for thermal energy storage systems. *HVAC&R Research 3*(3), 233–264.

Killian, M. and M. Kozek (2016). Ten questions concerning model predictive control for energy efficient buildings. *Building and Environment 105*, 403–412.

Maasoumy, M., M. Razmara, M. Shahbakhti, and A. S. Vincentelli (2014). Handling model uncertainty in model predictive control for energy efficient buildings. *Energy and Buildings 77*, 377–392.

Papadopoulos, A. M. (2016). Forty years of regulations on the thermal performance of the building envelope in europe: Achievements, perspectives and challenges. *Energy and Buildings 127*, 942–952.

Picard, D., J. Drgoňa, M. Kvasnica, and L. Helsen (2017). Impact of the controller model complexity on model predictive control performance for buildings. *Energy and Buildings 152*, 739–751.

Reynders, G., J. Diriken, and D. Saelens (2014). Quality of grey-box models and identified parameters as function of the accuracy of input and observation signals. *Energy and Buildings 82*, 263–274.

Sourbron, M., C. Verhelst, and L. Helsen (2013). Building models for model predictive control of office buildings with concrete core activation. *Journal of building performance simulation 6*(3), 175–198.

Sturzenegger, D., D. Gyalistras, M. Morari, and R. S. Smith (2015). Model predictive climate control of a swiss office building: Implementation, results, and cost–benefit analysis. *IEEE Transactions on Control Systems Technology 24*(1), 1–12.

Tian, W., Y. Heo, P. De Wilde, Z. Li, D. Yan, C. S. Park, X. Feng, and G. Augenbroe (2018). A review of uncertainty analysis in building energy assessment. *Renewable and Sustainable Energy Reviews 93*, 285–301.

Váňa, Z., J. Cigler, J. Širokỳ, E. Žáčeková, and L. Ferkl (2014). Model-based energy efficient control applied to an office building. *Journal of Process Control 24*(6), 790–797.

Wang, Z., T. Hong, and M. A. Piette (2019). Data fusion in predicting internal heat gains for office buildings through a deep learning approach. *Applied Energy 240*, 386–398.

Žáčeková, E., Z. Váňa, and J. Cigler (2014). Towards the real-life implementation of mpc for an office building: Identification issues. *Applied Energy 135*, 53–62.

Zhan, S. and A. Chong (2021). Data requirements and performance evaluation of model predictive control in buildings: A modeling perspective. *Renewable and Sustainable Energy Reviews*, 110835.