Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Evaluating the Impact of Data Availability on Machine Learning-augmented MPC for a Building Energy Management System

Jens Engel, Thomas Schmitt, Tobias Rodemann, Jürgen Adamy

2024

Preprint:

This is an accepted article published in IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe) 2024. The final authenticated version is available online at: https://doi.org/10.1109/ISGTEUROPE62998.2024.10863763 Copyright 2024 IEEE

Evaluating the Impact of Data Availability on Machine Learning-augmented MPC for a Building Energy Management System

Jens Engel[1,2], Thomas Schmitt[1], Tobias Rodemann[1], Jürgen Adamy[2]

***Abstract*—A major challenge in the development of Model Predictive Control (MPC)-based energy management systems (EMSs) for buildings is the availability of an accurate model. One approach to address this is to augment an existing gray-box model with data-driven residual estimators. The efficacy of such estimators, and hence the performance of the EMS, relies on the availability of sufficient and suitable training data. In this work, we evaluate how different data availability scenarios affect estimator and controller performance. To do this, we perform software-in-the-loop (SiL) simulation with a physics-based digital twin using real measurement data. Simulation results show that acceptable estimation and control performance can already be achieved with limited available data, and we confirm that leveraging historical data for pretraining boosts efficacy.**

***Index Terms*—data-driven residual estimator, model predictive control, digital twin, building control, machine learning**

I. Introduction

The increasing influx of renewable energy source into the public power grid leads to an increasing demand for intelligent energy management systems (EMSs) for buildings. A promising control method for such EMSs is Model Predictive Control (MPC) [1], [2]. Yet, for MPC-based EMS control approaches to be effective, an accurate underlying model is needed.

Various approaches for model building exist, generally categorized into white-box, gray-box and black-box modeling. White-box models build using e. g. building energy performance simulation tools like EnergyPlus or TRNSYS can provide very high levels of detail and accuracy, but can generally not be directly used within an MPC's optimal control problem (OCP). Grey-box models, like state space or Resistor-Capacitor (RC) models, have a lower level of accuracy, but are well suited for the use within an OCP. Both white-box and gray-box modeling of buildings are highly complex and require expert knowledge specific to each building, i. e. models cannot be easily transferred to other buildings [1], [3], [4]. Hence, data-driven black box modeling has become a popular alternative in recent years [5]–[7], e.g., using Gaussian Processes (GPs), artificial neural networks (ANNs), reinforcement learning, or, notably, using physics-informed neural networks (PINNs) [8]–[10]. While such black-box models require only a relatively low modeling effort, a substantial amount of data is needed for training, which can be challenging to acquire [6], [7]. At the same time, the ability to incorporate specific domain knowledge, interpretability, and adaptability of the model, which are the strengths of white/gray-box models, are lost. Furthermore, simple linear gray-box models have been shown to achieve similar performance [11], [12]. Hence, hybrid approaches combining these paradigms could be a promising way to enable larger scale real world adoption of building EMSs [7].

One such hybrid approach is the augmentation of a gray-box model using a *residual estimator*. Instead of replacing or adapting the gray-box model itself, its model error is reduced by estimating the residual value using an additional data-driven model. This concept has been applied to various applications of MPC [13]–[15]. In the context of building EMSs, a common approach is to use residual estimators to directly predict heat disturbances acting on the building, caused through occupant behavior and ambient effects [16]–[21]. These exogenous effects represent one of the largest error sources in buildings [1], [22], and these approaches have shown to be able to significantly reduce the residual error [19].

In previous work, we have shown that the control performance of an MPC augmented by machine learning-based residual estimation relies strongly on the availability and performance of the estimators [21]. As such estimators rely on training data, it becomes essential to investigate how the availability of data correlates with estimator performance. When applying such an approach in a real-world setting, historical data for pre-training estimators may not be available. This motivates two questions: 1.) How much data is needed until estimators achieve a good performance? 2.) If historical data is available and new data is collected in-the-loop, how should this new data be incorporated and the estimators retrained? While many authors note that these are important issues to be addressed [17], [19], [20], to the best of our knowledge, none have evaluated this systematically. In this study, we want to evaluate these questions in the context of a MPC-based EMS of an industrial multi-zone building which integrates error compensation through data-driven heat disturbance estimation. In software-in-the-loop (SiL) simulation with a physics-based digital twin, we evaluate different data availability scenarios and retraining strategies. We analyze both the accuracy of the residual estimator in the different settings, as well as how this relates to the performance of the proposed control approach.

This paper is organized as follows: The building under study,

[1]Honda Research Institute Europe GmbH, Offenbach, Germany. E-mail: {jens.engel, thomas.schmitt, tobias.rodemann}@honda-ri.de

[2]Control Methods and Intelligent Systems Laboratory, Technical University of Darmstadt, Darmstadt, Germany E-mail: juergen.adamy@tu-darmstadt.de

979-8-3503-9042-1/24/$31.00 ©2024 IEEE

its digital twin, as well as the proposed control approach are described in Section II. The residual error estimation is described in Section III. The simulation study and the discussion of the results is presented in Section IV. Finally, in Section V, a conclusion is given.

II. Energy Management System

In this work we consider a medium-sized company building situated in Offenbach, Germany. It has a total footprint of approx. $13,000\,\mathrm{m}^2$, which is partitioned into 9 temperature zones, namely offices, halls, workshops, an emissions lab and server rooms. In addition to the public power grid connection, it has a gas-fired combined heat and power plant (CHP) for co-production of heat and electricity with $199\,\mathrm{kW_{el}}$, and a photovoltaic (PV) plant of $750\,\mathrm{kWp}$, which supply an average electrical load demand of approx. $250\,\mathrm{kW}$. Besides the heat produced by the CHP, thermal energy is supplied by gas-fired heating boilers and an electric heating, ventilation, and air conditioning (HVAC) system. Furthermore, a second-life battery with a capacity of $98\,\mathrm{kWh}$ is available [23].

As a surrogate for the real building, we use a digital twin of the facility. The digital twin is a physics-based white-box model implemented in SimulationX [24], [25]. It covers the temperature zones' connections and heat losses to both the ambient air and the ground. It takes into account internal heat gains from electrical usage and building occupancy, and features a detailed HVAC system. It makes use of real measurement data of electrical loads and weather. For the control of the building, an MPC control approach adopted from [21] is used.

A. Modeling

For this study, we use a time-discrete state-space model. While a hierarchical approach can be used to improve scalability of large buildings [26], we choose a monolithic model for simplicity. Note that the results presented in Section IV are very similar for the hierarchical approach. As states, the model considers the stationary battery's stored energy E in kWh, and the 9 zone temperatures ϑ_i in °C, where ϑ_1 – ϑ_7 are regular building zones with heating and cooling, and ϑ_8 and ϑ_9 represent server rooms, which have no heating systems.

As inputs to the system, we consider the grid power P_{grid}, the (electrical) CHP power P_{chp}, the battery's (dis-)charging power P_{bat}, the gas heating power $\dot{Q}_{\mathrm{rad}}$, and the HVAC cooling power $\dot{Q}_{\mathrm{cool}}$. The heating and cooling powers are split into $\dot{Q}_{\mathrm{heat},i}\ \forall\, i = 1, \dots, 7$ and $\dot{Q}_{\mathrm{cool},i}\ \forall\, i = 1, \dots, 9$. As disturbances, we consider the PV power P_{PV}, the building's electrical load demand P_{dem}, the ambient air temperature ϑ_{air}, and (constant) heat disturbances $\dot{Q}_{\mathrm{other},i}$, which are losses to the ground for zones for $i = 1, \dots, 7$ and internal heatings for $i = 8, 9$. The unit of all electrical and thermal powers is kW.

The *time-continuous* differential equation of a temperature zone is thus given by

$$\dot{\vartheta}_i(t) = \frac{1}{C_{\mathrm{th},i}} \left(\dot{Q}_{\mathrm{heat},i}(t) + \dot{Q}_{\mathrm{cool},i}(t) + \dot{Q}_{\mathrm{other},i}(t) \right) - \sum_{j \neq i} \frac{\beta_{ij}}{C_{\mathrm{th},i}} \left(\vartheta_i(t) - \vartheta_j(t) \right) - \frac{H_{\mathrm{air},i}}{C_{\mathrm{th},i}} \left(\vartheta_i(t) - \vartheta_{\mathrm{air}}(t) \right), \tag{1}$$

$C_{\mathrm{th},i}$ is the thermal capacity of zone i in $\frac{\mathrm{kWh}}{\mathrm{K}}$, β_{ij} the heat transfer coefficient between zones i and j in $\frac{\mathrm{kW}}{\mathrm{K}}$, and $H_{\mathrm{air},i}$ the heat transfer coefficient between zone i and the ambient air in $\frac{\mathrm{kW}}{\mathrm{K}}$. Note that $\dot{Q}_{\mathrm{heat},i}(t) = 0\ \forall\, i = 8, 9$. The numerical values of all building parameters can be found in [19]. With (2), the thermal model part can be summarized as a (time-continuous) state space model as described in [27, pp. 22], but omitted here for brevity. With the state vector $x_{\mathrm{th}} = \begin{bmatrix} \vartheta_1 & \cdots & \vartheta_9 \end{bmatrix}^\intercal$, the input vector $u_{\mathrm{th}} = \begin{bmatrix} \dot{Q}_{\mathrm{heat},1} & \cdots & \dot{Q}_{\mathrm{heat},7} & \dot{Q}_{\mathrm{cool},1} & \cdots & \dot{Q}_{\mathrm{cool},9} \end{bmatrix}^\intercal$, and the disturbance vector $d_{\mathrm{th}} = \begin{bmatrix} \vartheta_{\mathrm{air}} & \dot{Q}_{\mathrm{other},1} & \cdots & \dot{Q}_{\mathrm{other},9} \end{bmatrix}^\intercal$, it can be discretized with the sampling time $T_{\mathrm{s}} = 0.5\,\mathrm{h}$ to

$$x_{\mathrm{th}}(k+1) = A_{\mathrm{th}}(T_{\mathrm{s}}) x_{\mathrm{th}}(k) + B_{\mathrm{th}}(T_{\mathrm{s}}) u_{\mathrm{th}}(k) + S_{\mathrm{th}}(T_{\mathrm{s}}) d_{\mathrm{th}}(k) + \epsilon(k). \tag{2}$$

As initially motivated, we consider a residual error of the state-space model $\epsilon(k) = \begin{bmatrix} \epsilon_1(k) & \cdots & \epsilon_9(k) \end{bmatrix}^\intercal$ as a time-variant heat disturbance, representing exogenous influences caused by occupant behavior and ambient effects. The estimation of $\epsilon(k)$ is discussed in Section III. The electrical part of the building can directly be expressed in time-discrete form as

$$E(k+1) = E(k) + P_{\mathrm{bat}}(k) \cdot T_{\mathrm{s}}, \tag{3}$$

with the balance equation

$$0 = P_{\mathrm{grid}}(k) - P_{\mathrm{bat}}(k) + P_{\mathrm{chp}}(k) + \frac{1}{\varepsilon_{\mathrm{c}}} \dot{Q}_{\mathrm{cool}}(k) + P_{\mathrm{dem}}(k) + P_{\mathrm{PV}}(k), \tag{4}$$

where $\varepsilon_{\mathrm{c}} = 1.78$ is the energy efficiency ratio of the HVAC's cooling system. Thermal and electrical powers are coupled by

$$P_{\mathrm{chp}}(k) = c_{\mathrm{cur}} \cdot \dot{Q}_{\mathrm{chp}}(k), \tag{5}$$

$$\dot{Q}_{\mathrm{heat}}(k) = \dot{Q}_{\mathrm{rad}}(k) + \dot{Q}_{\mathrm{chp}}(k), \tag{6}$$

$$\dot{Q}_{\mathrm{heat}}(k) = \sum\nolimits_{i=1}^{i=7} \dot{Q}_{\mathrm{heat},i}(k), \tag{7}$$

$$\dot{Q}_{\mathrm{cool}}(k) = \sum\nolimits_{i=1}^{i=9} \dot{Q}_{\mathrm{cool},i}(k), \tag{8}$$

where $c_{\mathrm{cur}} = 0.677$ is the CHP's power-to-heat ratio.

B. Control Approach

The model described in Section II-A is utilized to control the digital twin within a SiL setting with MPC. The optimization objective includes three cost functions,

$$J_{\mathrm{comf}}(k) = \sum_{n=1}^{N_{\mathrm{p}}} \sum_{i=1}^{7} w_{\mathrm{b},i} \cdot \left(\vartheta_i(n|k) - 22°\mathrm{C} \right)^2, \tag{9}$$

$$J_{\mathrm{server}}(k) = \sum_{n=1}^{N_{\mathrm{p}}} \sum_{i=8}^{9} w_{\mathrm{s},i} \Big(\max\left(15°\mathrm{C} - \vartheta_i(n|k),\, 0\right) + \max\left(\vartheta_i(n|k) - 21°\mathrm{C},\, 0\right) \Big), \tag{10}$$

$$J_{\mathrm{mon}}(k) = \sum_{n=0}^{N_{\mathrm{p}}-1} \ell_{\mathrm{mon}} \Big(P_{\mathrm{grid}}(n|k), P_{\mathrm{chp}}(n|k), \dot{Q}_{\mathrm{heat}}(n|k) \Big), \tag{11}$$

where $w_{\mathrm{b},i}$ and $w_{\mathrm{s},i}$ are weights derived from the zones' thermal capacities, i. e.

$$w_{\mathrm{b},i} = \frac{C_{\mathrm{th},i}}{\sum_{j=1}^{7} C_{\mathrm{th},j}} \; \forall\, i = 1, \ldots, 7, \tag{12}$$

$$w_{\mathrm{s},i} = \frac{C_{\mathrm{th},i}}{\sum_{j=8}^{9} C_{\mathrm{th},j}} \; \forall\, i = 8, 9, \tag{13}$$

and ℓ_{mon} describes the costs arising from gas usage and buying/selling electrical energy from/to the public grid, including peak costs. The notation $a(n|k)$ refers to the value of $a(k+n)$ as predicted at time step k, and will be used from here on. The overall objective function is then given by

$$\begin{aligned} J_{\mathrm{opt}}(k) = w_{\mathrm{comf}} \cdot J_{\mathrm{comf}}(k) + w_{\mathrm{server}} \cdot J_{\mathrm{server}}(k) \\ + w_{\mathrm{mon}} \cdot J_{\mathrm{mon}}(k), \end{aligned} \tag{14}$$

where $w_{\mathrm{comf}} = 0.99$, $w_{\mathrm{server}} = 1$, $w_{\mathrm{mon}} = 0.01$ are chosen such that the emphasis is completely on the thermal regulation, which makes the interpretation of the simulation results easier. Note that multi-objective MPC can be used to either determine appropriate fixed weights [28], [29], or to dynamically choose the Pareto optimal solution of your preference [30].

Summarizing all decision variables as $u_{\mathrm{all}} = [u_{\mathrm{th}}^{\intercal}, P_{\mathrm{grid}}, P_{\mathrm{chp}}, P_{\mathrm{bat}}, \dot{Q}_{\mathrm{cool}}, \dot{Q}_{\mathrm{heat}}]$, the MPC's OCP is given by

$$\min_{\boldsymbol{u}_{\mathrm{all}}(k)} \; J_{\mathrm{opt}}(k) \tag{15a}$$

$$\text{s. t. } (2)\text{–}(8), \; \forall\, n = 0, \ldots, N_{\mathrm{p}} - 1, \tag{15b}$$

and subject to box constraints on both states and control inputs. $\boldsymbol{u}_{\mathrm{all}}(k) = (u_{\mathrm{all}}(0|k), \ldots, u_{\mathrm{all}}(N_{\mathrm{p}} - 1|k))$ denotes the sequence of control inputs. The time step notation (k) and $(k+1)$ in (2) – (8) are to be read as $(n|k)$ and $(n+1|k)$, respectively. As the prediction horizon, we use $N_{\mathrm{p}} = 48$ steps of $T_{\mathrm{s}} = 0.5\,\mathrm{h}$ each, i. e. 1 day in total.

III. Residual Error Estimation

As motivated in the introduction, and described in the previous section, we consider a residual model error $\epsilon(k)$ in the controller. We assume the time-variant error to represent unknown exogenous influences. While these influences are not directly measurable and difficult to model explicitly, they are strongly correlated with measurable and predictable data. Therefore, we use machine learning regression models to learn a data-driven estimator $\tilde{\epsilon}(k)$. As features, we propose the total building load $P_{\mathrm{dem}}(k)$, the time of day $\mathrm{ToD}(k) : \mathbb{N} \to [0, 24)$ in h, the day of week $\mathrm{DoW}(k) : \mathbb{N} \to \{0, \ldots, 6\}$ in d, (lags of) the ambient temperature $\vartheta_{\mathrm{air}}(k-n)$, as well as the electrical loads in the server zones $P_{\mathrm{server},1}(k), P_{\mathrm{server},2}(k)$, which are part of the overall building load $P_{\mathrm{dem}}(k)$ and can be metered separately. We split the estimation into two parts: one regressor estimates the errors of the building zones $\epsilon_1(k), \ldots, \epsilon_7(k)$, a second those of the server zones $\epsilon_8(k), \epsilon_9(k)$. As input features for the regressors we use:

- Building zones: $P_{\mathrm{dem}}(k)$, $\vartheta_{\mathrm{air}}(k)$, $\vartheta_{\mathrm{air}}(k-1)$, $\vartheta_{\mathrm{air}}(k-2)$, $\mathrm{ToD}(k)$, $\mathrm{DoW}(k)$,
- Server zones: $P_{\mathrm{dem}}(k)$, $\vartheta_{\mathrm{air}}(k)$, $P_{\mathrm{server},1}(k)$, $P_{\mathrm{server},2}(k)$.

The target variable of the regression models is the observed model error $\epsilon(k)$. It can be calculated as the difference between measured and predicted state. In general, the predicted state $x_{\mathrm{th}}(1|k)$ may already contain an estimate $\tilde{\epsilon}(k)$, so the true target variable is calculated as

$$\epsilon(k) = \big(x_{\mathrm{th}}(k+1) - x_{\mathrm{th}}(1|k)\big) + \tilde{\epsilon}(k). \tag{16}$$

In general, this difference will also include prediction errors of the disturbances. As we only want to estimate exogenous influences, this can be corrected by first recalculating $x_{\mathrm{th}}(1|k)$ with measurements of the disturbances.

IV. Simulation Study

We want to evaluate how different data availability scenarios affect the performance of the previously introduced residual estimation, and how this relates to controller performance. To do this, we carry out a SiL simulation study using the digital twin model and real measurement data for the years 2022 and 2023. All simulations were implemented in Python, with the digital twin connected through a functional mock-up unit (FMU), and using Gurobi as a solver. In a first step, we simulated the controller without any error compensation, i. e. $\tilde{\epsilon}(k) = 0$, for the year 2022. We then use this data to evaluate 4 different training data availability scenarios in simulation of the year 2023:

- S0: No error compensation is applied, i. e. $\tilde{\epsilon}(k) = 0$.
- S1: Monthly retraining without any prior data, i.e. we incrementally simulate each month, after which we train estimators on the collected data, and apply them in the following month.
- S2: Analogous monthly retraining but with prior data, i.e. additionally historical training data for all of 2022 is appended to the collected data.
- S3: Analogous to S2, but a rolling 12 calendar month window is applied to the training data.

We compare the results against two baseline scenarios:

- S4: Estimators trained on historical full year 2022 data are applied to 2023 simulation without any retraining.
- S5: Estimators trained on S0, re-simulating 2023 with active compensation.

We performed the evaluation with different candidate regression models, namely XGBoost [31], a multilayer perceptron (MLP), and multiple linear regression. Out of these, XGBoost performed best for estimating the building zone errors, while the MLP yielded very similar performance. For the server zone estimator, a linear regressor was chosen. Since we want to focus on assessing the general estimation performance of the intrinsic model error, we use perfect predictions for PV power, the building's loads and ambient air temperature. Figure 1 shows the weighted mean absolute residual error (WMARE), i. e.

$$\mathrm{WMARE} = \frac{1}{N} \sum_{k=0}^{N-1} \underbrace{\sum_{i=1}^{9} w_{\mathrm{th},i} \cdot |\epsilon_i(k)|}_{\mathrm{WARE}(k)}, \tag{17}$$

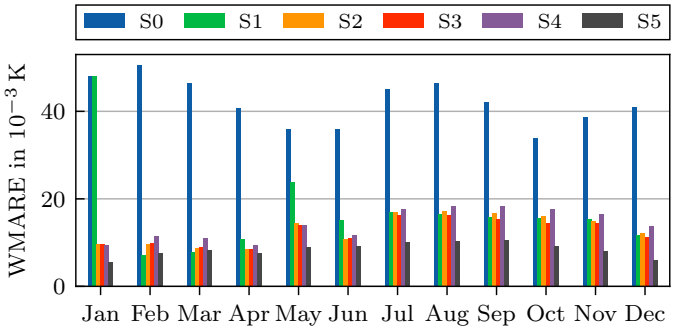

Fig. 1. Weighted mean absolute residual error (WMARE) over the course of the year 2023 in the different scenarios.

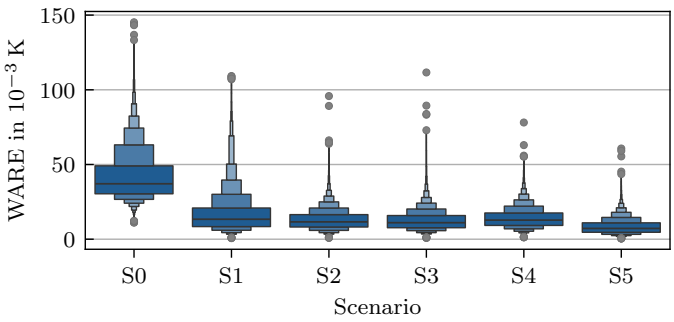

Fig. 2. Letter-value plot showing the distribution of the weighted absolute residual error (WARE) compared between the different scenarios. The central quantile represents 50 % of the data, which are then halved at every next quantile, i. e. 25 %, 12.5 %, and so forth.

with $w_{\mathrm{th},i} = \frac{C_{\mathrm{th},i}}{\sum_{j=1}^{9} C_{\mathrm{th},j}}$ $\forall\, i = 1, \ldots, 9$ over the course of the year for all scenarios. It shows that even without any prior training data in scenario S1, good estimation performance can be achieved in the first months of the year, even very slightly outperforming the pretraining scenarios S2–S5. During May, the estimation performance drops significantly, recovering in June. From July on, S1 achieves similar performance to S2 and S3, indicating that half a year of training data is a critical amount for achieving consistent performance. Scenario S3 seems to slightly outperform S2, suggesting that forgetting data from the previous year may be advantageous. Scenario S4, i. e. when no retraining of the previous year estimator is applied, exhibits similar but decreased performance to scenarios S2 and S3. Overall, scenarios S1–S4 have good but, compared to the perfect information scenario S5, significantly decreased performance. These findings are confirmed by the distribution of the weighted absolute residual error (WARE) over the whole year 2023 as shown in Figure 2. Without available prior training data, the worst compensation performance is achieved, but the overall error is still significantly reduced (S1). The errors of S2 and S3 are almost equally distributed, showing no significant difference. Even though the performance of S4 is worse than S2 and S3, the overall error is still significantly reduced, indicating that the applied regressor models generalize reasonably well. A key criterion of the thermal control performance is the root mean square error (RMSE) of the temperature tracking, which indicates how far individual zones deviate from the reference

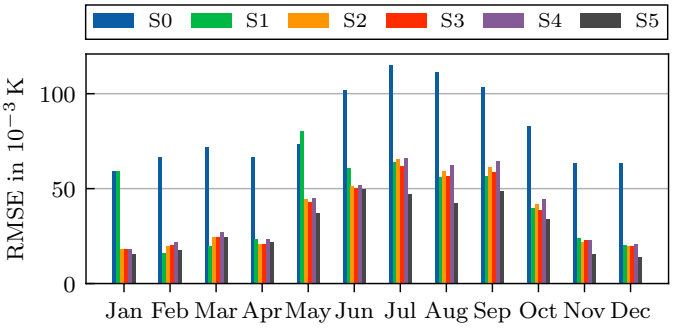

Fig. 3. Root mean square error (RMSE) of the temperature tracking over the course of the year 2023 in the different scenarios.

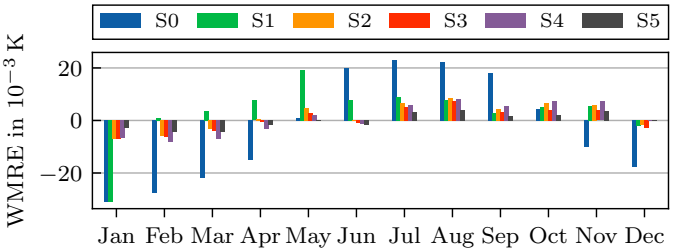

Fig. 4. Weighted mean residual error (WMRE) over the course of the year 2023 in the different scenarios.

temperature,

$$\mathrm{RMSE} = \sum_{i=1}^{7} w_{\mathrm{b},i} \cdot \sqrt{\frac{1}{N} \sum_{k=1}^{N} \left(\vartheta_i(k) - 22°\mathrm{C}\right)^2}. \quad (18)$$

Figure 3 shows this metric over the course of the year 2023. It can be seen that the difference in tracking performance between scenarios mostly matches that in compensation performance. While the latter is relatively consistent from June onward, the tracking performance has a seasonal trend. This is due to the controller weighing the tracking error against monetary costs (especially peak costs). This is also the reason for why there is almost no difference in tracking performance in June: under certain conditions, to avoid peaks, better compensation performance may not significantly improve the tracking error. Notably, in May, in S1 a worse tracking performance than in S0 is achieved, even though the absolute residual error was significantly reduced. This is due to the fact that while the mean *absolute* error was reduced, the mean error has a significant bias, as shown in Figure 4. Without historical data, the estimator cannot account for the vanishing negative bias in May. This causes the temperature to drift, as the controller will consistently receive biased estimates of the heat disturbances.

To confirm the overall findings, we performed the same simulation study for the year 2022 (by simulating 2021 first). Table I shows the full-year WMARE and the RMSE of the temperature tracking for both 2022 and 2023. These data confirm the previous findings that incremental learning without prior data can quickly achieve good performance and that leveraging historical data is beneficial. Furthermore, the data indicate that there is a slight advantage if windowed historical data is used for training.

Table I
Weighted mean absolute residual error (WMARE) and root mean square error (RMSE) of the tracking performance for all scenarios in 2022 and 2023.

Metric	Year	S0	S1	S2	S3	S4	S5
WMARE	2022	41.21	12.96	11.42	11.38	12.63	7.58
in 10^{-3} K	2023	40.38	14.98	13.69	13.09	14.89	8.85
RMSE	2022	97.35	53.79	48.92	48.68	51.16	44.62
in 10^{-3} K	2023	89.47	51.07	46.40	44.56	48.15	37.56

V. Conclusion & Outlook

In this study we evaluated how different data availability scenarios affect the performance of residual estimation of exogenous influences on an industrial multi-zone building and how this relates to controller performance. We found that even without historical data available, data-driven residual estimation can quickly achieve good performance on a very conservative amount of data. If historical data is available, building on this data in an incremental fashion improves estimation performance. We further showed that thermal control performance of our proposed approach benefits from estimation accuracy, where higher accuracy leads to better performance. Overall, we find that the proposed residual error estimation approach can achieve good results on a conservative amount of training.

References

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

[2] T. Péan, R. Costa-Castelló, and J. Salom, "Price and carbon-based energy flexibility of residential heating and cooling loads using model predictive control," *Sustainable Cities and Society*, vol. 50, p. 101579, Oct. 2019.

[3] S. Prívara, J. Cigler, Z. Váňa, F. Oldewurtel, C. Sagerschnig, and E. Žáčeková, "Building modeling as a crucial part for building predictive control," *Energy and Buildings*, vol. 56, pp. 8–22, 2013.

[4] W. Tian, Y. Heo, P. de Wilde, Z. Li, D. Yan, C. S. Park, X. Feng, and G. Augenbroe, "A review of uncertainty analysis in building energy assessment," *Renewable and Sustainable Energy Reviews*, vol. 93, pp. 285–301, 2018.

[5] Z. Wang and Y. Chen, "Data-driven modeling of building thermal dynamics: Methodology and state of the art," *Energy and Buildings*, vol. 203, p. 109405, 2019.

[6] E. T. Maddalena, Y. Lian, and C. N. Jones, "Data-driven methods for building control — a review and promising future directions," *Control Engineering Practice*, vol. 95, p. 104211, 2020.

[7] G. Halhoul Merabet, M. Essaaidi, M. Ben Haddou, B. Qolomany, J. Qadir, M. Anan, A. Al-Fuqaha, M. R. Abid, and D. Benhaddou, "Intelligent building control systems for thermal comfort and energy-efficiency: A systematic review of artificial intelligence-assisted techniques," *Renewable and Sustainable Energy Reviews*, vol. 144, p. 110969, Jul. 2021.

[8] G. Gokhale, B. Claessens, and C. Develder, "Physics informed neural networks for control oriented thermal modeling of buildings," *Applied Energy*, vol. 314, p. 118852, May 2022.

[9] L. Di Natale, B. Svetozarevic, P. Heer, and C. Jones, "Physically consistent neural networks for building thermal modeling: Theory and analysis," *Applied Energy*, vol. 325, p. 119806, Nov. 2022.

[10] X. Wang and B. Dong, "Physics-informed hierarchical data-driven predictive control for building hvac systems to achieve energy and health nexus," *Energy and Buildings*, vol. 291, p. 113088, Jul. 2023.

[11] F. Bünning, B. Huber, A. Schalbetter, A. Aboudonia, M. Hudoba de Badyn, P. Heer, R. S. Smith, and J. Lygeros, "Physics-informed linear regression is competitive with two machine learning methods in residential building mpc," *Applied Energy*, vol. 310, p. 118491, Mar. 2022.

[12] P. Stoffel, M. Berktold, and D. Müller, "Real-life data-driven model predictive control for building energy systems comparing different machine learning models," *Energy and Buildings*, vol. 305, p. 113895, Feb. 2024.

[13] E. Picotti, F. Bianchin, and M. Bruschetta, "Real-time learning-based nonlinear model predictive control of a virtual motorcycle employing grey-box modeling through gaussian processes," *Control Engineering Practice*, vol. 144, no. Germany, p. 105837, Mar. 2024.

[14] D. Scheurenberg, S. Stemmler, and D. Abel, "Evaluation of data enhanced model predictive control for a coupled tank system," in *2023 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 2023, pp. 79–84.

[15] F. M. Gray and M. Schmidt, "A hybrid approach to thermal building modelling using a combination of Gaussian processes and grey-box models," *Energy and Buildings*, vol. 165, pp. 56–63, apr 2018.

[16] C. A. Thilker, H. Madsen, and J. B. Jørgensen, "Advanced forecasting and disturbance modelling for model predictive control of smart energy systems," *Applied Energy*, vol. 292, no. Denmark, p. 116889, Jun. 2021.

[17] M. J. Ellis, "Machine learning enhanced grey-box modeling for building thermal modeling," in *2021 American Control Conference (ACC)*. IEEE, may 2021.

[18] P. Kumar, J. B. Rawlings, M. J. Wenzel, and M. J. Risbeck, "Grey-box model and neural network disturbance predictor identification for economic MPC in building energy systems," *Energy and Buildings*, 2023.

[19] T. Schmitt, J. Engel, and T. Rodemann, "Regression-based model error compensation for hierarchical MPC building energy management system," in *2023 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 2023.

[20] W. Liang, H. Li, S. Zhan, A. Chong, and T. Hong, "Energy flexibility quantification of a tropical net-zero office building using physically consistent neural network-based model predictive control," *Advances in Applied Energy*, p. 100167, Feb. 2024.

[21] J. Engel, T. Schmitt, T. Rodemann, and J. Adamy, "Hierarchical MPC for building energy management: Incorporating data-driven error compensation and mitigating information asymmetry," *Applied Energy*, vol. 372, p. 123780, Oct. 2024.

[22] F. Oldewurtel, D. Sturzenegger, and M. Morari, "Importance of occupancy information for building climate control," *Applied Energy*, vol. 101, pp. 521–532, jan 2013.

[23] M. Stadie, T. Rodemann, A. Burger, F. Jomrich, S. Limmer, S. Rebhan, and H. Saeki, "V2B vehicle to building charging manager," in *EVTeC: 5th International Electric Vehicle Technology Conference 2021*. Japanese Society of Automotive Engineers, May 2021.

[24] R. Unger, B. Mikoleit, T. Schwan, B. Bäker, C. Kehrer, and T. Rodemann, "Green building - modeling renewable building energy systems with eMobility using Modelica," in *Proceedings of Modelica 2012 conference*. Munich, Germany: Modelica Association, September 2012.

[25] A. Castellani, S. Schmitt, and S. Squartini, "Real-world anomaly detection by using digital twin systems and weakly supervised learning," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 7, pp. 4733–4742, jul 2021.

[26] J. Engel, T. Schmitt, T. Rodemann, and J. Adamy, "Hierarchical economic model predictive control approach for a building energy management system with scenario-driven EV charging," *IEEE Transactions on Smart Grid*, pp. 3082–3093, 2022.

[27] T. Schmitt, "Multi-objective building energy management optimization with model predictive control," Ph.D. dissertation, Technische Universität Darmstadt, Darmstadt, 2022.

[28] T. Schmitt, J. Engel, T. Rodemann, and J. Adamy, "Application of Pareto optimization in an economic model predictive controlled microgrid," in *28th Mediterranean Conference on Control and Automation (MED)*. IEEE, 2020, pp. 868–874.

[29] T. Schmitt, T. Rodemann, and J. Adamy, "Multi-objective model predictive control for microgrids," *at - Automatisierungstechnik*, vol. 68, no. 8, pp. 687–702, 2020.

[30] T. Schmitt, M. Hoffmann, T. Rodemann, and J. Adamy, "Incorporating human preferences in decision making for dynamic multi-objective optimization in Model Predictive Control," *Inventions*, vol. 7, no. 3, 2022.

[31] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, aug 2016.