Proceedings of Machine Learning Research vol. 242:117–129, 2024 6th Annual Conference on Learning for Dynamics and Control

Linearised Data-Driven LSTM-Based Control of Multi-Input HVAC Systems

Andreas Hinderyckx AHINDERYCKX@YAHOO.CO.UK
KU Leuven, EluciDATA Lab, Sirris (Brussels, Belgium)

Florence Guillaume FLORENCE.GUILLAUME@SIRRIS.BE
EluciDATA Lab, Sirris (Brussels, Belgium)

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

The pursuit of sustainability has paved the way for building management systems (BMSs) that can steer buildings in an energy-efficient way. In this article, a deep learning approach is proposed to control multi-input heating, ventilation, and air conditioning (HVAC) systems in order to minimise both thermal discomfort and operational cost. More particularly, a long short-term memory (LSTM)-based encoder-decoder process model, trained on historical weather data and control sequences generated while the building was steered by a modern rule-based controller (RBC), is fed into an optimisation problem, to which a change of variables is applied to efficiently model the effect of interdependent control inputs. Both the nonlinear LSTM process model and the cost function of the optimisation problem are linearised to formulate the control problem as a mixed integer linear programming (MILP) problem, which ensures that the controller can operate in near real-time and with limited computational power. Moreover, to avoid resorting to model extrapolation and to improve the model's predictive performance, the set of allowed control signal values is restricted using a percentile-based approach. In addition to the purely data-driven controller (DDC), a hybrid controller is designed to leverage the strengths of the RBC and the DDC. The performance of both controllers is benchmarked against the RBC's performance using the BOPTEST simulation environment under various experiment settings, highlighting how the hyperparameters affect the controller's performance. Compared to the RBC, we show that the proposed controllers realise substantial improvements in terms of both thermal comfort and operational cost while controlling a single zone or two zones simultaneously.

Keywords: HVAC Control, Encoder-Decoder LSTM Model, PNMPC, Interdependent Inputs

1. Introduction

The soaring prices of energy and ever-increasing use thereof make efficient use of this resource more important than ever before. Out of all sectors, the building sector is one of the major consumers, as it is responsible for both 30% of the global annual greenhouse gas emissions and 20 to 40% of the global energy consumption (Hidalgo-León et al. (2019)). Moreover, the European Union has made a reduction in net greenhouse gas emissions one of its primary goals, aiming to reduce it by at least 55% by 2030 and to make climate neutrality legally binding by 2050 (European Parliament (2022)). Additionally, many heating, ventilation and air conditioning (HVAC) systems are not yet fully automated, and require human expertise to continually tune or control them (Royapoor et al. (2018), Stluka et al. (2018)). These factors motivate the search for physics-based or data-driven approaches that guarantee thermal comfort in a cost-efficient manner using HVAC systems.

In the past decades, digital and electronic equipment have decreased in price and risen in popularity. This has caused the category of model predictive control (MPC) algorithms to gain increased

attention. These algorithms require a process model to predict the evolution of the building temperature. One approach is to use a physics-based model to approximate the building thermal dynamics mathematically. However, this approach is time consuming and expensive, as an accurate physical model must be conceived for each new building. On top of that, the calculations that make use of this model are computationally taxing, which gives rise to problems when applying this technique to controllers with limited computational capabilities (Foucquier et al. (2013)). Alternatively, grey box modelling leverages partial physical knowledge and combines it with data to calibrate the parameters of a simplified physical model. Although effective in practice, this type of model requires extensive domain knowledge and data to be tuned correctly. To address the issues of the physics-based approaches, and given the increased availability of building data, data-driven modelling has risen in popularity during the last decade. This black box approach constructs an approximate model of a building, solely based on weather data and sensor data collected from that specific building. Therefore, the building does not need to be specifically modelled and development times are reduced, as black box models can be designed and calibrated more efficiently compared to their physics-based counterparts. Recently, Terzi et al. (2020) have shown that long short-term memory (LSTM) networks can be effective when used as process model in the context of modelling building dynamics. This motivates further research to be conducted in this area.

The next sections of this paper are organised as follows. Work that forms the foundation of the approach and how the proposed methodology further builds upon it, is summarised in Section 2. Section 3 describes the data-driven methodology, and in particular the process model and the minimisation problem forming the MPC control loop. Section 4 evaluates the methodology on both a single-zone and a multi-zone Building Optimization Testing (BOPTEST) test case (Arroyo et al. (2021)). Finally, the main conclusions and possible directions for future work are outlined in Section 5.

2. Related Work and Contributions

The proposed approach builds upon five pieces of related work, with a main focus on designing a data-driven controller for a building that is already controlled via an rule-based controller (RBC), as is often encountered in real-world applications. Firstly, the practical nonlinear model predictive control (PNMPC) algorithm proposed by Plucenio et al. (2007) enables to use a nonlinear process model in the classical MPC control loop, without having to incur the computational cost of inferencing a nonlinear model at each time step. To do so, the output of the model is approximated by a linear function of the control increments around the model's base response. Secondly, Terzi et al. (2020) and Schwedersky et al. (2019) propose to use an LSTM-based process model to model complex system dynamics. In Terzi et al. (2020), the system to be steered consists of a large building, but the process model is used to control solely the number of active coolers, without controlling temperature set points or interdependent control inputs. Schwedersky et al. (2019) integrate the LSTM-based process model in the PNMPC framework to control a pH-neutralisation process. This article builds further on this integration both by relaxing the strong assumption that all operating conditions are simulated in the training data, and by extending the process model to an LSTM-based encoder-decoder architecture. Thirdly, Asghari et al. (2022) discuss a transformation technique to eliminate nonlinear products in the constraints and the objective function of a minimisation problem. This technique, combined with the PNMPC algorithm, will be used to rewrite the control problem as a minimisation problem with both a linear cost function and linear constraints. This allows to use linear, rather than nonlinear, mixed integer programming solvers to

solve the minimisation problem at each time step, which greatly reduces computational overhead. Fourthly, Jain et al. (2020) propose a similar approach to our work, but we extend their approach by linearising the optimisation problem by modeling interdependent inputs as the product of a continuous and a discrete input, and by using an encoder-decoder LSTM (rather than a feed-forward neural network)-based process model to perform multi-step forecasting. The proposed approach leverages the five previous pieces of work by combining them to construct an encoder-decoder-based PNMPC controller that handles interdependent inputs effectively.

The contributions of this article are fourfold. Firstly, it is the first work to use an LSTM-based encoder-decoder neural network as process model in the PNMPC algorithm to naturally integrate past and future values of disturbances and control actions. To that end, the work of Schwedersky et al. (2019) is used as a basis and modified accordingly. Secondly, we focus on designing a data-driven controller for a building that is already controlled via an RBC, as it is often encountered in real-word applications. It follows that empirical control actions for a given context (i.e. disturbances and previous control sequences) do not exhibit high variability, which might lead to untrustworthy predictions due to extrapolation of the process model. To avoid such extrapolations in unexplored areas of the feature space, a percentile-based technique is proposed, which improves the model's predictive performance by limiting the range of allowed control signal values. Thirdly, we extend the control problem to handle interdependent control actions, which introduce extra nonlinearities in the optimisation problem. To eliminate nonlinearities that arise from these interdependencies and from the nonlinear process model, we combine the PNMPC algorithm with a product-based transformation of variables, as discussed in Asghari et al. (2022). Fourthly, the control approach is evaluated in a single-zone and a multi-zone BOPTEST test case, both of which are characterised by different inputs and thermal properties. To validate the approach, we explore what savings in terms of thermal discomfort and operational cost can be realised when using a data-driven controller (DDC) rather than an RBC to steer an HVAC system. To this aim, the BOPTEST simulation environment is used to extensively compare the performance of the RBC and the DDC, and to uncover how the latter is affected by the hyperparameter values.

3. Methodology

The control approach is based on the MPC control loop as described in Camacho et al. (2004). We highlight its main components below.

3.1. Process Model

Building on the work of Terzi et al. (2020) and of Schwedersky et al. (2019) and motivated by its strong predictive capabilities for time series (Schmidhuber et al. (2005), Yu et al. (2019)), we use an LSTM-based process model to forecast the zone temperature from time $+1$ up to time $+N$, where time 0 denotes the current time. However, instead of the classic LSTM architecture, we consider an LSTM-based encoder-decoder network, which is better suited for the control task under consideration and results in a lower root mean squared error (RMSE). Indeed, the features that are input of the model can be divided into three categories: (1) variables known at past time steps (*past covariates*), (2) variables known at both past and future time steps (*future covariates*), and (3) the values of the target variable known at past time steps (*past target variable values*). To predict the target variable value at time $+i$, $i \in \{1, 2, \ldots, N\}$, the past covariates and past target feature values are available up to time 0, along with the future covariate values up to time step $+i$. Therefore,

when using a direct, rather than a recursive, forecasting approach, the input size varies depending on which time step is predicted: this poses an issue for the conventional LSTM network, as it has a fixed input size. We alleviate this issue by using an LSTM-based encoder-decoder network as process model (see Figure 1). This model consists of two parts: (1) an encoder that processes the past covariates and the past values of the target variable into a lower dimensional context vector, $\boldsymbol{c}$, and (2) a decoder that processes the context vector in combination with the set of future covariates to produce the next target feature values. This architecture enables the model to produce a sequence of output values, whose length is independent of that of the input sequence.

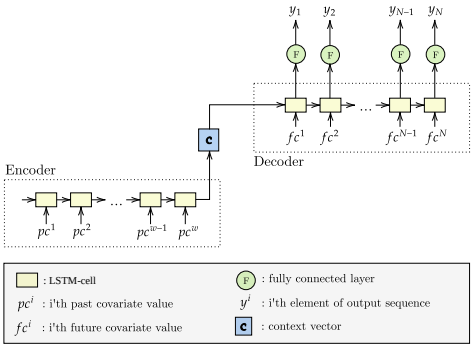

Figure 1: Graphical overview of the LSTM-based encoder-decoder process model. w denotes the number of past covariate features and N the prediction horizon. Figure based on Angioi (2022).

To formulate the control task as a linear, rather than a nonlinear, mixed integer programming problem, the PNMPC algorithm is applied to linearise the process model at each time step. In doing so, the process model's output at the next N time steps is approximated by a linear function of the control increments, $\boldsymbol{\Delta u}$, within the control horizon:

$$\boldsymbol{y}(\boldsymbol{u}) = \begin{bmatrix} y_{+1}(u_{1,+1}, \ldots, u_{k,+1}) \\ \vdots \\ y_{+N}(u_{1,+1}, \ldots, u_{k,+1}, \ldots, u_{1,+N}, \ldots, u_{k,+N}) \end{bmatrix} \approx \boldsymbol{f}(\boldsymbol{u}^{(0)}) + \boldsymbol{G}(\boldsymbol{u}^{(0)})\boldsymbol{\Delta u}, \quad (1)$$

where $y_{+i}, i \in \{1, 2, \ldots, N\}$ is the process model output at time $+i$ and $\boldsymbol{u} = [u_{1,+1} \ldots u_{k,+1} \ldots u_{1,+N} \ldots u_{k,+N}]^T$, where $u_{\ell,+i}$, $\ell \in \{1, 2, \ldots, k\}$, $i \in \{1, 2, \ldots, N\}$ is the value of the ℓth control action at time step $+i$. Moreover, $\boldsymbol{f} = \boldsymbol{y}(\boldsymbol{u}^{(0)})$ denotes the process model's free response, $\boldsymbol{u}^{(0)}$ the optimal control actions vector obtained at the previous control time and $\boldsymbol{\Delta u} = \boldsymbol{u} - \boldsymbol{u}^{(0)}$ the control increment vector. Generalising Equation (1) in vector notation for time step t, we obtain:

$$\hat{\boldsymbol{y}}(\boldsymbol{u}^{(t)}) \approx \begin{bmatrix} \hat{y}_{+1}(\boldsymbol{u}^{(t-\Delta t)}) \\ \hat{y}_{+2}(\boldsymbol{u}^{(t-\Delta t)}) \\ \vdots \\ \hat{y}_{+N}(\boldsymbol{u}^{(t-\Delta t)}) \end{bmatrix} + \begin{bmatrix} \frac{\partial \hat{y}_{+1}}{\partial (\boldsymbol{u})_1}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} & 0 & \cdots & 0 \\ \frac{\partial \hat{y}_{+2}}{\partial (\boldsymbol{u})_1}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} & \frac{\partial \hat{y}_{+2}}{\partial (\boldsymbol{u})_2}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \hat{y}_{+N}}{\partial (\boldsymbol{u})_1}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} & \frac{\partial \hat{y}_{+N}}{\partial (\boldsymbol{u})_2}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} & \cdots & \frac{\partial \hat{y}_{+N}}{\partial (\boldsymbol{u})_{kN}}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} \end{bmatrix} \begin{bmatrix} (\boldsymbol{\Delta u}^{(t)})_1 \\ (\boldsymbol{\Delta u}^{(t)})_2 \\ \vdots \\ (\boldsymbol{\Delta u}^{(t)})_{kN} \end{bmatrix} \quad (2)$$

where $(\boldsymbol{u})_l$ denotes the lth component of vector $\boldsymbol{u}$ and

$$(\boldsymbol{\Delta u}^{(t)})_l = (\boldsymbol{u}^{(t)})_l - (\boldsymbol{u}^{(t-\Delta t)})_l$$

and

$$\frac{\partial \hat{y}_{+j}}{\partial (\boldsymbol{u})_l}\Big|_{\boldsymbol{u}^{(t-\Delta t)}} \approx \frac{\hat{y}_{+j}(\boldsymbol{u}^{(t-\Delta t)} + \boldsymbol{\varepsilon}_{(l)}) - \hat{y}_{+j}(\boldsymbol{u}^{(t-\Delta t)})}{\varepsilon} \tag{3}$$

denotes the partial derivative of the jth element of the output vector with respect to the lth future control signal value within the control horizon, evaluated at $\boldsymbol{u}^{(t-\Delta t)}$ (Yang et al. (2016)). Here, $\boldsymbol{\varepsilon}_{(l)}$ represents a zero-vector containing a small perturbation ε at position l. The hyperparameters used for the process model are depicted in Table 1.

LSTM cells	16	Epochs	8	Batch Size	64
Input Length	8	Loss Function	MSE	Training Optimiser	COCOB
Early Stopping Delta	1e-5	Patience	3		

Table 1: The process model's hyperparameter values.

3.2. Minimisation Problem

The minimisation problem, that is embedded in the MPC control loop, is of the following form:

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{minimize}} \quad & \boldsymbol{c}^T \boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} \preccurlyeq \boldsymbol{h}, \end{aligned} \tag{4}$$

where $\boldsymbol{c}^T \boldsymbol{x}$ denotes the cost function that is to be minimised, $\boldsymbol{A}\boldsymbol{x} \preccurlyeq \boldsymbol{h}$ denotes the set of constraints, and $\preccurlyeq$ represents the element-wise inequality operator. The cost function encodes preferences with respect to the different evaluation criteria:

$$\mathbf{1}\Big[\underbrace{f_c(\boldsymbol{u})^T}_{\text{operational cost}} \quad \underbrace{\omega\boldsymbol{\zeta}^T \quad \omega\boldsymbol{\xi}^T}_{\text{thermal discomfort}} \Big]^T, \tag{5}$$

where $\mathbf{1} = [1 \dots 1] \in \mathbb{R}^{(k+2)N}$, f_c represents the operational cost as a linear function of the control signal values within the control horizon, $\boldsymbol{u}$, and where $\boldsymbol{\zeta} = [\zeta_{+1} \dots \zeta_{+N}]^T \succcurlyeq \mathbf{0}$, $\boldsymbol{\xi} = [\xi_{+1} \dots \xi_{+N}]^T \succcurlyeq \mathbf{0}$. Here, ζ_{+i} and ξ_{+i}, $i \in \{1, 2, \dots, N\}$ are slack variables that indicate how much the zone temperature deviates from the lower and upper comfort bounds, respectively, for time step $+i$. These slack variables are introduced to convert the hard thermal constraints to soft constraints, to ensure that a solution always exists. Finally the weight $\omega > 0$ allows to specify the relative importance of the thermal discomfort compared to the operational cost. The set of constraints can be decomposed into three types: (1) constraints that enforce the zone temperatures to stay within the (softened) comfort bounds, (2) constraints that ensure admissible values for the control signals, slack variables and possible auxiliary variables, and (3) constraints that link the values of the control signals to the process model's output temperatures.

Real-world control problems are often characterised by interdependent inputs, which are inputs that cannot be considered separately as they jointly affect the process' output. In building control, binary on/off activation $\boldsymbol{b}_\ell = [b_{\ell,+1} \dots b_{\ell,+N}]^T$ and continuous supply temperature set point signals $\boldsymbol{v}_\ell = [v_{\ell,+1} \dots v_{\ell,+N}]^T$ of a given HVAC component ℓ are examples of interdependent control signals. To model their joint effect, the problem can be written as a mixed integer linear programming

(MILP) in which the product of the variables is used as feature of the process model: $\boldsymbol{b}_\ell \odot \boldsymbol{v}_\ell$, where $\odot$ denotes the element-wise multiplication operator. However, this expression is nonlinear, which entails that nonlinear solvers are needed to solve the optimisation problem. To circumvent this computational overhead, an auxiliary vector $\tilde{\boldsymbol{u}}_\ell = \boldsymbol{b}_\ell \odot \boldsymbol{v}_\ell$ as well as extra constraints are introduced for each pair of interdependent control signals, which allows the problem to be rewritten in a linear form and to be solved using MILP solvers (Asghari et al. (2022)); in this case, CVXOPT is used to solve the MILP problem (Andersen et al. (2024)). Combining Equations (4) and (5), the vector of unknowns $\boldsymbol{x}$ is given by the block vector $\begin{bmatrix}\boldsymbol{u}^T & \boldsymbol{\zeta}^T & \boldsymbol{\xi}^T & \boldsymbol{b}^T & \boldsymbol{v}^T\end{bmatrix}^T$, where $\boldsymbol{u} = \begin{bmatrix}\tilde{\boldsymbol{u}}_1^T \dots \tilde{\boldsymbol{u}}_k^T\end{bmatrix}^T$, $\boldsymbol{b} = \begin{bmatrix}\boldsymbol{b}_1^T \dots \boldsymbol{b}_k^T\end{bmatrix}^T$ and $\boldsymbol{v} = \begin{bmatrix}\boldsymbol{v}_1^T \dots \boldsymbol{v}_k^T\end{bmatrix}^T$.

3.3. Percentile Cutoff

In real-world control applications, one of the main challenges when building a purely data-driven process model is the lack of variability in the training data, and in particular in the control actions, which often follow a (semi-)manual control strategy. This might lead to a large discrepancy between the control actions considered during the control phase and those encountered during the training phase, and hence to extrapolation of the process model. To address this issue, there are two options. Either the training data could be generated in such a way that it sufficiently covers all possible values of the model's input space, or the control approach could be limited to only use values that occur sufficiently often in the training data. We opt for the latter approach, as using the former approach would result in significant costs or thermal discomfort for the occupants, which is to be avoided in real-life applications. To restrict the range of control values, a form of robust optimisation is used, which we call percentile cutoff. This approach limits the range of values allowed during the control phase to those values included between the 100αth and $100(1-\alpha)$th percentiles, where $\alpha \in (0, 0.5)$ is a hyperparameter. The higher the value of α, the lower the extrapolation of the process model, but the more constrained the control actions. To mitigate the risk of model extrapolation, other robust optimisation techniques could be considered, such as adjusting the cost function to penalise control actions that rarely occur in the training data.

4. Implementation and Results

To validate the methodology, two building emulators available in the BOPTEST framework are used: (1) a single-zone building with low thermal inertia and equipped with an emission heating system, called the `bestest_hydronic` test case, and (2) a two-zone building with higher thermal inertia, the `multizone_office_simple_hydronic` test case, within which each zone has a dedicated emission subsystem and an air handling unit (AHU) (BOPTEST, icupeiro (2022)). For each test case, a dataset is generated by controlling the building emulator for a simulated duration of one year using an advanced RBC, which consists of a hysteresis controller with a night set-back and a predetermined room temperature compensation. For this period, the weather data, the control signal values, the zone temperatures and the HVAC system sensor data are saved. The first nine months of data constitute the process model's training data while the last three months are split into two equally sized sets that make up the validation and testing sets, respectively. The past covariates consist of data that is captured by HVAC system sensors, while the future covariates consist of the control actions and of variables that can be predicted (such as weather forecasts).

To assess the controllers' performance, we consider two criteria: (1) the thermal discomfort, and (2) the operational cost, or energy cost[1], which are displayed in Figure 2. Equation (6) shows the thermal discomfort, which is defined as the (cumulative) deviation of the zone temperature from the thermal comfort range, where $[t_0, t_f]$ is the period during which the HVAC system is controlled, $s_z^{(t)}$ is the deviation from the comfort range in zone z at time t, and M is the number of zones. Equation (7) shows the operational cost, which is defined in terms of the magnitude of the auxiliary variable, where $\tilde{u}_{z,i}^{(t)}$ is the value of the ith auxiliary variable of zone z at time t, and $\tilde{U}_z$ is the set of auxiliary control signals of zone z. Figure 3 shows that for the `bestest_hydronic` test case, the auxiliary variables are strongly correlated to the value of the BOPTEST cost key performance indicator (KPI). This motivates the choice of a linear function for the (unknown) operational cost in Equation (5) and the use of (7) as evaluation metric.

$$D(t_0, t_f) = \frac{1}{M} \sum_{z=1}^{M} \int_{t_0}^{t_f} \| s_z^{(t)} \| \, dt \qquad (6)$$

$$\widetilde{C}(t_0, t_f) = \frac{1}{M} \sum_{z=1}^{M} \int_{t_0}^{t_f} \sum_{i \in \widetilde{U}_z} \tilde{u}_{z,i}^{(t)} \, dt \qquad (7)$$

Figure 2: Thermal discomfort (top) and operational cost (bottom) metrics.

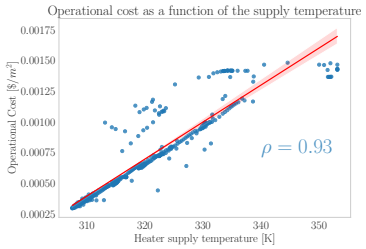

Figure 3: The BOPTEST cost KPI is strongly correlated to the auxiliary variable for the `bestest_hydronic` test case.

4.1. Experiments

To assess the controllers' performance, receding horizon control is used for both test cases, where the constrained optimisation problem is repeatedly solved for a period of one week, ranging from 23 December 2021 until 30 December 2021.

4.1.1. SINGLE-ZONE BUILDING

The left panel of Figure 4 shows the evolution of the zone temperature obtained by both the RBC (red) and the DDC (blue). The DDC manages to keep the temperature within the comfort bounds more often, although the temperature generated by the DDC strongly oscillates. This is explained by the fact that the simulated building has a low thermal inertia and that the DDC switches the HVAC system on and off more often. The RBC avoids this oscillatory behaviour by applying more gradual transitions in the HVAC system's activation regime[2]. Predicting such volatile temperatures poses an additional challenge for the process model, as is shown in the right panel of Figure 4. An overview of the process model's accuracy in terms of its RMSE for each of the test cases is shown in Table 2. Figure 5 shows that the data-driven approach manages to eliminate the sharp peaks in thermal discomfort of up to 5 K that are generated by the RBC, thanks to its ability to anticipate the need for heating shortly before the building starts to be occupied. However, the DDC does result in higher thermal discomfort compared to the RBC at times where the comfort bounds are constant.

1. The operational cost and energy cost criteria coincide, up to a proportional factor, since simulations are based on a constant pricing scenario.
2. More details on the control signals for the RBC and DDC can be found in Hinderyckx (2023).

	Test Set	Control Set
`bestest_hydronic`	0.26 K	0.45 K
`multizone_office_simple_hydronic`	0.24 K	0.27 K

Table 2: RMSE of the process model during the training phase and the control phase for both test cases (in Kelvin).

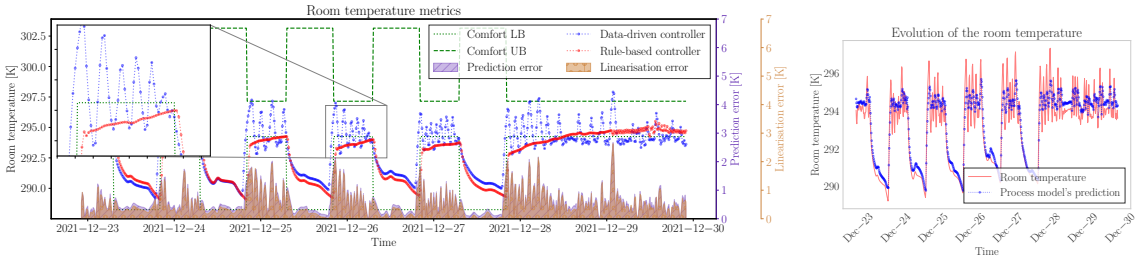

Figure 4: Left panel: Zone temperature resulting from the DDC and RBC, and forecasting errors of the DDC. The comfort bounds are shown in green and the linearisation and prediction errors are depicted by the shaded regions, whose values can be read on the right-hand axis. Right panel: Comparison between the predicted zone temperature and the true zone temperature during the control phase.

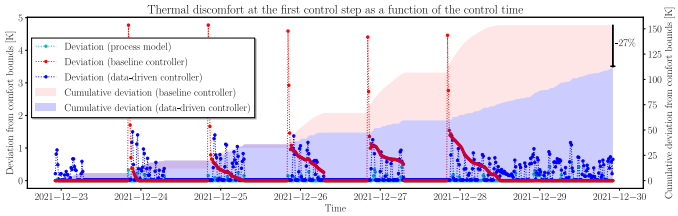

Figure 5: Thermal discomfort $D(t_0, t_f)$ as a function of time. The peaks in discomfort resulting from the RBC occur at time steps at which the comfort bounds change. The values of the cumulative thermal discomfort that are visualised by the shaded regions can be read on the right-hand axis.

Given that the DDC outperforms the RBC when the comfort bounds change, and that the RBC outperforms the DDC when the comfort bounds are constant, we expect to see improved performance when the two controllers are combined. To leverage the advantages of both controllers, we introduce a hybrid controller, which dynamically — based on the shape of the past and future comfort bounds — decides whether to use the DDC or the RBC. More concretely, if there is a change in the comfort bounds in a window of size W centered around the control time, the DDC is selected, while the RBC is selected if the comfort bounds are constant within the window. Using a window size W equal to 12 time steps, the temperatures shown in the left panel of Figure 6 are obtained. The hybrid controller is able to anticipate heating needs, as is shown by the figure's enlarged view, while limiting the oscillatory behaviour, which mainly occurs when the DDC mode is activated. To reduce the magnitude of oscillations, an alternative to using a hybrid controller is to penalise high variations of the control signals in the cost function. However, this comes at the expense of worsening the two other dimensions: guaranteeing comfort and minimising cost. Moreover, the oscillations could be dampened by reducing the linearisation error. This could be achieved by applying an iterative version of the PNMPC algorithm, PNMPCi, which runs the PNMPC algorithm

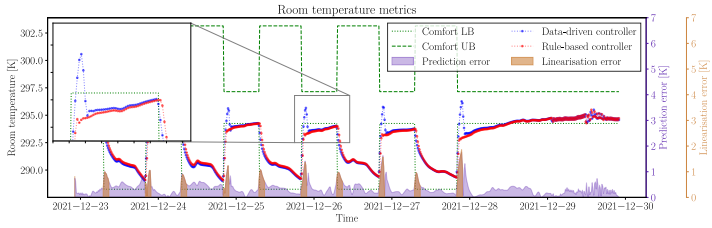

	DDC	Hybrid DDC
Operational cost	-9%	-3%
Thermal discomfort	-27%	-44%

Figure 6: Zone temperature obtained by the hybrid controller (left) and relative improvement of the DDC and of the hybrid DDC compared to the RBC for the two evaluation metrics (right).

for multiple iterations to reduce the mismatch between the process model's output and its linear approximation (Schwedersky and Flesch (2022)).

The right panel of Figure 6 shows the relative improvement of the DDC and of the hybrid DDC compared to the RBC. Globally speaking, the two controllers outperform the RBC in terms of both thermal comfort and cost. The DDC realises larger cost savings due to its oscillatory control behaviour. Although theoretically optimal, this behaviour is not desirable in practice, as it could cause significant wear and tear on the HVAC system. Moreover, the hybrid DDC reduces the thermal discomfort significantly compared to the DDC.

To assess the robustness of the control approach, the one week control experiment is repeated for different times throughout the year. Figure 7 shows the performance of the hybrid DDC, expressed as relative improvement compared to the RBC's performance, and the forecast RMSE as a function of the control week offset, which is defined as the number of weeks in between the end of the training phase and the beginning of the control phase. The RMSE of the process model does not fluctuate much with the control week offset, which indicates that the model accuracy is robust to different times of the year. The hybrid DDC shows an overall improvement in terms of both evaluation metrics, except for a single peak in the relative discomfort at an offset of six weeks. However, this relative increase translates into a negligible increase of the absolute discomfort $D(t_0, t_f)$.

4.1.2. MULTI-ZONE BUILDING

To simulate a more realistic scenario, the control approach is applied to a two-zone building, in which each zone contains a dedicated HVAC system, composed of both a heating and a ventilation system that have to be steered simultaneously. To extend the approach to a multi-zone building, decentralised MPC is used by deploying a separate instance of the controller in each zone. Table 3 summarises the results of the decentralised MPC approach applied to the `multizone_office_simple_hydronic` test case. One apparent outlier in the results of the hybrid DDC is the RMSE value of the south zone temperature. This is due to the significant discrepancy between the data encountered during the training phase and the control phase, as is shown in the left panel of Figure 8. For all experiments, a default value of $\alpha = 10^{-3}$ was used in the percentile cutoff procedure. However, increasing α to $5 \cdot 10^{-2}$ leads to both a better predictive and control performance, as is depicted in the right panel of Figure 8. This shows that managing the associated trade-off is important to improve the performance: larger values of α leads to a better predictability, at the cost of limiting the freedom of the controller, as the range of allowed control values is limited.

	KPIs		RMSE	
	Thermal discomfort	Operational cost	Training	Control
North zone	-36%	-3%	0.24 K	0.27 K
South zone	-64%	-3%	0.14 K	**1.14 K**
Average	-50%	-3%		

Table 3: Evaluation metrics and RMSE obtained by the hybrid DDC in the multi-zone, multi-subsystem test case.

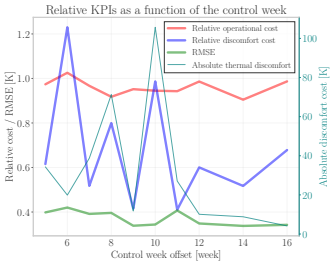

Figure 7: Impact of the control week on the RMSE (green) and evaluation metrics for the hybrid DDC (red and blue).

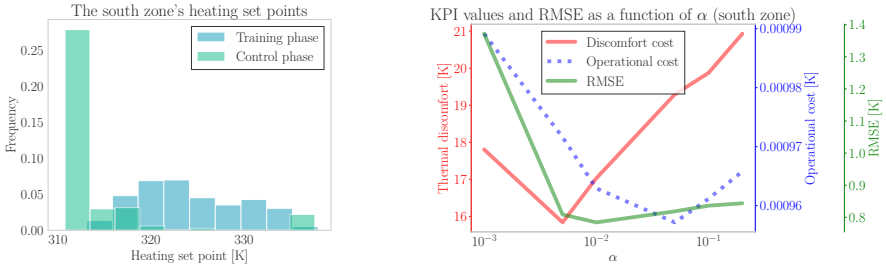

Figure 8: Discrepancy between the distribution of the training data and the control data for the south zone (left) and effect of α on the RMSE (green) and on the evaluation metrics (red and blue) of the south zone (right).

5. Conclusion

This work presents a linearised encoder-decoder LSTM controller that combines past and future covariates in a natural way and extends previous work by considering the case of interdependent control actions, while preserving the affine characteristic of the cost function. To cope with the limited variability observed in real training datasets, the range of allowed control signal values is restricted using a percentile-based approach, allowing to improve the process model's predictive accuracy during the control phase. Making use of the BOPTEST framework, the proposed approach is shown to globally outperform the performance of advanced RBCs in terms of both the thermal comfort and the operational cost. Furthermore, it is demonstrated that dynamically combining the data-driven and rule-based controllers allows to anticipate heating needs while maintaining more stable zone temperatures under constant comfort bounds, by leveraging the advantages of both controllers. The robustness of the approach was shown in terms of the number of weeks separating the training and the control phase, as well as in terms of the complexity of the building emulator. The proposed hybrid controller can be further extended to take dynamic pricing into account by building on the work of Bindels and Nelemans (2022).

Acknowledgments

This research is funded by the Brussels-Capital Region — Innoviris (Brussels Public Organisation for Research and Innovation), under the grant number 2021-RDIR-53.

References

Martin S. Andersen, Joachim Dahl, and Lieven Vandenberghe. CVXOPT: Python Software for Convex Optimization, 2024. URL https://cvxopt.org/.

Alessandro Angioi. Time series forecasting with an lstm encoder/decoder in tensorflow 2.0 — Alessandro Angioi, my personal site/blog. mainly about science, technology, and coding. https://www.angioi.com/time-series-encoder-decoder-tensorflow/, 2022. [Online; accessed 16-December-2022].

Javier Arroyo, Carlo Manna, Fred Spiessens, and Lieve Helsen. An OpenAI-Gym environment for the Building Optimization Testing (BOPTEST) framework. In Proceedings of the 17th IBPSA Conference, Bruges, Belgium, September 2021.

Mohammad Asghari, Amir M. Fathollahi-Fard, S. M. J. Mirzapour Al-e hashem, and Maxim A. Dulebenets. Transformation and linearization techniques in optimization: A state-of-the-art survey. Mathematics, 10(2):283, Jan 2022. ISSN 2227-7390. doi: 10.3390/math10020283. URL http://dx.doi.org/10.3390/math10020283.

Martijn Bindels and Bart Nelemans. Cost reduction of heat pump assisted membrane distillation by using variable electricity prices. Desalination, 530:115669, 2022. ISSN 0011-9164. doi: https://doi.org/10.1016/j.desal.2022.115669. URL https://www.sciencedirect.com/science/article/pii/S0011916422001242.

BOPTEST. Bestest hydronic heat pump (detailed documentation). https://ibpsa.github.io/project1-boptest/testcases/ibpsa/testcases_ibpsa_bestest_hydronic_heat_pump/. [Online; accessed 23-May-2023].

E.F. Camacho, C. Bordons, and C.B. Alba. Model Predictive Control. Advanced Textbooks in Control and Signal Processing. Springer London, 2004. ISBN 9781852336943. URL https://books.google.be/books?id=Sc1H3f3E8CQC.

European Parliament. Reducing carbon emissions: Eu targets and measures. https://www.europarl.europa.eu/news/en/headlines/society/20180305STO99003/reducing-carbon-emissions-eu-targets-and-measures, 2022. [Online; accessed 1-February-2023].

Aurélie Foucquier, Sylvain Robert, Frédéric Suard, Louis Stéphan, and Arnaud Jay. State of the art in building modelling and energy performances prediction: A review. Renewable and Sustainable Energy Reviews, 23:272–288, 2013. ISSN 1364-0321. doi: https://doi.org/10.1016/j.rser.2013.03.004. URL https://www.sciencedirect.com/science/article/pii/S1364032113001536.

Ruben Hidalgo-León, Jaqueline Litardo, Javier Urquizo, Daniel Moreira, Pritpal Singh, and Guillermo Soriano. Some factors involved in the improvement of building energy consumption: A brief review. In 2019 IEEE Fourth Ecuador Technical Chapters Meeting (ETCM), pages 1–6, 2019. doi: 10.1109/ETCM48019.2019.9014890.

Andreas Hinderyckx. Linearised lstm-based data-driven model predictive control of multi-input hvac systems. Master's thesis, KU Leuven, Kasteelpark Arenberg 1 bus 2200, BE-3001 Leuven, June 2023. Supervised by Florence Guillaume and Wannes Meert.

icupeiro. Model documentation — IBPSA Project 1 BOPTEST, Building Optimization Performance Tests. https://github.com/ibpsa/project1-boptest/blob/63aa06fb6d9c3f0e3e9ae05df6afd9fa15ecc72f/testcases/multizone_office_simple_hydronic/doc/index.html, August 2022. [Online; accessed 23-May-2023].

Achin Jain, Francesco Smarra, Enrico Reticcioli, Alessandro D'Innocenzo, and Manfred Morari. Neuropt: Neural network based optimization for building energy management and climate control, 2020.

A. Plucenio, D.J. Pagano, A.H. Bruciapaglia, and J.E. Normey-Rico. A practical approach to predictive control for nonlinear processes. IFAC Proceedings Volumes, 40(12):210–215, 2007. ISSN 1474-6670. doi: https://doi.org/10.3182/20070822-3-ZA-2920.00035. URL https://www.sciencedirect.com/science/article/pii/S1474667016355288. 7th IFAC Symposium on Nonlinear Control Systems.

Mohammad Royapoor, Anu Antony, and Tony Roskilly. A review of building climate and plant controls, and a survey of industry perspectives. Energy and Buildings, 158:453–465, 2018. ISSN 0378-7788. doi: https://doi.org/10.1016/j.enbuild.2017.10.022. URL https://www.sciencedirect.com/science/article/pii/S0378778817318522.

Jürgen Schmidhuber, Daan Wierstra, and Faustino Gomez. Evolino: Hybrid neuroevolution/optimal linear search for sequence learning. pages 853–858, 01 2005.

Bernardo B. Schwedersky and Rodolfo C.C. Flesch. Nonlinear model predictive control algorithm with iterative nonlinear prediction and linearization for long short-term memory network models. Engineering Applications of Artificial Intelligence, 115:105247, 2022. ISSN 0952-1976. doi: https://doi.org/10.1016/j.engappai.2022.105247. URL https://www.sciencedirect.com/science/article/pii/S0952197622003177.

Bernardo B. Schwedersky, Rodolfo C.C. Flesch, and Hiago A.S. Dangui. Practical nonlinear model predictive control algorithm for long short-term memory networks. IFAC-PapersOnLine, 52(1):468–473, 2019. ISSN 2405-8963. doi: https://doi.org/10.1016/j.ifacol.2019.06.106. URL https://www.sciencedirect.com/science/article/pii/S2405896319301922. 12th IFAC Symposium on Dynamics and Control of Process Systems, including Biosystems DYCOPS 2019.

Petr Stluka, Girija Parthasarathy, Steve Gabel, and Tariq Samad. Architectures and Algorithms for Building Automation — An Industry View, pages 11–43. Springer International Publishing, Cham, 2018. ISBN 978-3-319-68462-8. doi: 10.1007/978-3-319-68462-8_2. URL https://doi.org/10.1007/978-3-319-68462-8_2.

E. Terzi, T. Bonetti, D. Saccani, M. Farina, L. Fagiano, and R. Scattolini. Learning-based predictive control of the cooling system of a large business centre. Control Engineering Practice, 97:104348,

2020. ISSN 0967-0661. doi: https://doi.org/10.1016/j.conengprac.2020.104348. URL https://www.sciencedirect.com/science/article/pii/S0967066120300307.

Bryant Picon Yang, Agustinho Plucenio, and Dpto de Automaçao e Sistemas. Practical non-linear model predictive control pnmpc: Algorithm implementations. In CBA. 2016.

Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. A Review of Recurrent Neural Networks: LSTM Cells and Network Architectures. Neural Computation, 31(7):1235–1270, 07 2019. ISSN 0899-7667. doi: 10.1162/neco_a_01199. URL https://doi.org/10.1162/neco_a_01199.