PROCEEDINGS OF ECOS 2022 - THE 35TH INTERNATIONAL CONFERENCE ON
EFFICIENCY, **C**OST, **O**PTIMIZATION, **S**IMULATION AND ENVIRONMENTAL IMPACT OF ENERGY SYSTEMS
1-3 JULY, 2022, COPENHAGEN, DENMARK

Automating the data-driven predictive control design process for building thermal management

Edward O'Dwyer[a], Paola Falugi[b], Nilay Shah[c] and Eric Kerrigan[d]

[a] *Imperial College London, United Kingdom, e.odwyer@imperial.ac.uk, CA*
[b] *Imperial College London, United Kingdom, p.falugi@imperial.ac.uk*
[c] *Imperial College London, United Kingdom, n.shah@imperial.ac.uk,*
[d] *Imperial College London, United Kingdom, e.kerrigan@imperial.ac.uk*

Abstract:

Decarbonisation of the energy sector has brought about a need for intelligent building energy management strategies that can respond optimally to changing external conditions and user requirements. Data-driven modelling methods that can reduce the implementation effort of such strategies have recently gained significant attention, but the inherent variability of building design can mean that a non-trivial parameter selection and tuning process remains. These parameters include continuous and granular variables. This paper proposes a method for automating the control design process for building energy management, enabling predictive control implementation without bespoke system identification or extensive controller tuning effort. This is achieved by combining Data-enabled Predictive Control (DeePC) methods with a parameter selection and tuning procedure using the Mesh Adaptive Direct Search (MADS) algorithm. This selection and tuning procedure is carried out prior to the implementation of the DeePC using data measurements from the building under standard control. The impact of parameter selection on the control performance is shown through a set of case studies using both real data and a simulation environment in which predictive control is used to maintain thermal comfort in a building. Using the autotuned DeePC approach, a reduction of 74% in thermal comfort violation is achieved compared to a classical PI approach, while a reduction of 28% is achieved relative to a DeePC control without autotuning.

Keywords:

Data-driven control, Mesh-adaptive direct search, building energy, automated deployment.

1. Introduction

The need for intelligent building control that can offer flexible operation so as to contribute to a decarbonised energy landscape is becoming ever more urgent [1]. Model Predictive Control (MPC) for buildings has been well established as a key technology in this transition [2] due to its ability to incorporate predictions, constraints and external incentives into an optimisation based framework. A significant bottleneck in the rollout of the technology is the need for suitable building models. The varying nature of building designs and user behaviour makes off-the-shelf solutions challenging. For example, it has been noted [3] that creating and calibrating models accounts for 70% of the implementation effort of MPC in buildings.

Driven by this need to expedite the implementation process, data-driven predictive control methods have become increasingly popular in the academic literature on building control [4, 5]. Rather than leaning on building modelling expertise, data-driven methods can exploit the increasing ubiquity of data and data-handling platforms in the built environment to learn the relevant dynamics from historical data [6].

Aside from the built environment, the wider control community has seen significant developments in data-driven methods in recent times. Of particular note has been the upsurge of methods founded on Willems' fundamental lemma [7]. These methods rely on the outcome of the fundamental lemma that, under certain straightforward excitation conditions, historical data structures can be used to project all permissible trajectories of a Linear Time-Invariant (LTI) system. Data-driven control methods analogous to various model-based control methods have been developed based on this [8], notably including a data-driven counterpart to MPC, referred to as Data-enabled Predictive Control (DeePC) [9]. Many extensions have been proposed in recent times, often centred around robustness requirements [10], nonlinearity in the system under control [11] and noisy data [12].

Due to the desire for data-driven predictive control in the built environment, building energy management can be commonly found as an application of DeePC and its alternatives in literature [13, 14]. The key advantage of these methods is the ease with which implementation can be carried out. A gap that remains, however, is that although model parameterisation is not required, controller design is still needed. Regularisation and relaxation weights, as well as data trajectory lengths, can greatly impact the behaviour of the controller. Different systems require different control parameters and, once again, given the varying nature of buildings, a one-size-fits-all approach is unlikely to be sufficient. Since replacing a model design problem with a control design problem is not a desirable outcome, methods for automating the process of control parameter selection are highly desirable in this context. Given the data-driven nature of the controller, using data for controller parameter selection would be of interest.

In this paper, a method is proposed for tuning and configuring a DeePC strategy in an automated manner using a Derivative-Free Optimisation (DFO) approach, namely the Mesh Adaptive Direct Search (MADS) algorithm as proposed in [15]. Using a similar structure to the data-driven controller, a data-driven simulation method is used and a blackbox system is formed in which controller parameters are the inputs and a simulation inaccuracy metric is the output. This blackbox system is optimised using MADS - i.e., parameters are found to minimise simulation inaccuracy - and the optimal parameters are then used to define a DeePC strategy which is used for building control. The data-driven simulation behaviour is illustrated using data from a real dwelling. The behaviour of controllers in which autotuning has been carried out is then analysed using a detailed building modelling framework in which two buildings with six zones in each are simulated. The performance of the autotuned DeePC controllers is compared to that of DeePC controllers without bespoke tuning in terms of their ability to maintain comfort in the simulated buildings. A comparison is also made with a non-predictive PI strategy.

In Section 2., the DeePC and blackbox optimisation methodology is outlined. This is followed by the set of case studies in Section 3., in which real data and a simulation environment are used to analyse the performance of the approach.

2. Automating the deployment of data-driven building energy management

2.1. DeePC for building energy management

In this section, a DeePC formulation is outlined. It should be noted that variations to this formulation are possible. However, this section provides a background for subsequent sections which rely on similar concepts and notation.

An LTI system of unknown parameters can be given as

$$\begin{aligned} x[k+1] &= Ax[k] + Bu[k] \\ y[k] &= Cx[k] + Du[k], \end{aligned} \tag{1}$$

where $u[k] \in \mathbb{R}^m$ and $y[k] \in \mathbb{R}^p$ are the input and output vectors, respectively, $x[k] \in \mathbb{R}^n$ is the system state-vector, and $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$, $C \in \mathbb{R}^{p\times n}$ and $D \in \mathbb{R}^{p\times m}$ are parameter matrices.

Matrices formed of $T_0 \in \mathbb{Z}_{>0}$ persistently exciting historical data sequences measured from this system are constructed, most commonly in a Hankel structure. A trajectory w, for example, is defined as persistently exciting of order $L_0 \in \mathbb{Z}_{>0}$, if the Hankel matrix $\mathscr{H}_{L_0}(w)$ is of full row-rank with

$$\mathscr{H}_{L_0}(w) := \begin{bmatrix} w_1 & \cdots & w_{T_0-L_0+1} \\ \vdots & \ddots & \vdots \\ w_{L_0} & \cdots & w_{T_0} \end{bmatrix}. \tag{2}$$

Input and output sequences collected from a system can respectively be defined as $u_{tr} = \left[u_1^T, \dots, u_{T_0}^T\right]^T \in \mathbb{R}^{mT_0}$ and $y_{tr} = \left[y_1^T, \dots, y_{T_0}^T\right]^T \in \mathbb{R}^{pT_0}$. The Hankel matrix structures, composed of these trajectories can then be formed and partitioned as follows:

$$\begin{aligned} \begin{bmatrix} U_p \\ U_f \end{bmatrix} &:= \mathscr{H}_{T_{ini}+N}(u_{tr}), \\ \begin{bmatrix} Y_p \\ Y_f \end{bmatrix} &:= \mathscr{H}_{T_{ini}+N}(y_{tr}), \end{aligned} \tag{3}$$

where N is the prediction horizon, T_{ini} is user-defined, U_p and Y_p have T_{ini} rows while U_f and Y_f have N rows. These data structures can be used to define permissible trajectories from the system defined in (7) by finding solutions that satisfy the following:

$$\begin{bmatrix} U_p \\ U_f \\ Y_p \\ Y_f \end{bmatrix} g \quad = \quad \begin{bmatrix} u_{ini} \\ u_f \\ y_{ini} \\ y_f \end{bmatrix}, \tag{4}$$

where $g \in \mathbb{R}^{T_0 - T_{ini} - N + 1}$, u_{ini} and y_{ini} denote the most recent T_{ini} input and output measurements from the system, respectively, and u_f and y_f define future input and output trajectories. By using this condition as a constraint in an optimisation formulation, a predictive control strategy can be formed. This is the central concept behind the DeePC strategy proposed by [9]. Successful application of the strategy requires the objective function of the optimisation formulation to include certain constraint relaxations as well as regularisation of the optimisation variable g to account for real-world imperfections in the data and nonlinearity in the system under control.

The DeePC approach has been applied to the problem of building thermal management [13,14,16] whereby inputs are heat flows and outputs are room temperatures. As noted in the introduction, building modelling can be time-consuming and, given the diverse nature of the global building stock and the need to quickly deliver methods for improved energy management, expediting the control implementation process is an attractive prospect. Nonetheless, in removing the need for model parameterisation, a controller parameterisation problem remains. Regularisation weights, relaxation weights and trajectory lengths must be chosen. Suitable choices of these parameters can vary from system to system. As the key advantage of the approach in this application is to reduce the implementation effort, methods are required to automate the parameter selection process.

2.2. Optimal selection of control parameters

A method for optimising the DeePC parameters using historical data is proposed here. Following this concept, the building is operated in a standard rule-based or classical manner for a period of time prior to implementation of the DeePC to generate data for parameter optimisation. This data is used to construct Hankel matrices in the manner described in Section 2.1.. Equation 4 is then used to form a data-driven simulation formulation as will be described in this section. The accuracy of the data-driven simulation is dependent on the same parameters as the DeePC controller. By comparing the data-driven simulation output to the historical measurements, the simulation performance can be quantified. A blackbox system can then be formed whereby the inputs are the controller parameters and the output is the simulation inaccuracy. A suitable blackbox optimisation method is then required to optimise the inputs to minimise the output.

Parameters affecting the behaviour of the DeePC algorithm include the trajectory initialisation length T_{ini} and the regularisation and relaxation weights. The former is an integer variable, while the latter two are continuous. The optimisation approach chosen must then be able to handle both continuous and integer variables. The Derivative-Free Optimisation (DFO) approach proposed in [15] is one such method. This method employs a Mesh Adaptive Direct Search (MADS) algorithm to determine a local optimum to a blackbox optimisation problem by sampling from a grid that gets successively refined. The blackbox system is next defined.

Assume at time k, where $k > T_0 + T_{ini} + N$, measured input data $u^{meas} = \left[u_1^T, \ldots, u_k^T\right] \in \mathbb{R}^{mk}$ and output data $y^{meas} = \left[y_1^T, \ldots, y_k^T\right] \in \mathbb{R}^{pk}$ have been collected. Picking a starting point k_0 and defining T_0-length sequences $u_{tr}^{sim} = \left[u_{k_0}^{meas} \ldots u_{k_0+T_0-1}^{meas}\right]$ and $y_{tr}^{sim} = \left[y_{k_0}^{meas} \ldots y_{k_0+T_0-1}^{meas}\right]$, the Hankel matrices of (3) are formed. An N-length simulated output y_f^{sim} can be found from time (k_1) to $(k_1 + N - 1)$ for any $k_1 > k_0 + T_0 + T_{ini}$ by solving $y_f^{sim} = Y_f g_{sim}^*$ where g_{sim}^* is the solution to:

$$g_{sim}^* = \arg\min_{g_{sim}} \sum_{i=0}^{T_{ini}-1} \left(\varepsilon_i^{ini}\right) + a_0 g_{sim}^T g_{sim} + a_1 \sum_{i=0}^{N-1} \left(\varepsilon_i^{f}\right) \tag{5}$$

s.t.

$$\begin{aligned} U_p g &= u_{ini}^{sim} \\ U_f g - u_f^{sim} &\leq \varepsilon^f \\ u_f^{sim} - U_f g &\leq \varepsilon^f \\ Y_p g - y_{ini}^{sim} &\leq \varepsilon^{ini} \\ y_{ini}^{sim} - Y_p g &\leq \varepsilon^{ini}. \end{aligned} \tag{6}$$

In this formulation, $u_{ini}^{sim} = [u_{k_1-T_{ini}}^{meas} \ldots u_{k_1-1}^{meas}]$, $u_f^{sim} = [u_{k_1}^{meas} \ldots u_{k_1+N-1}^{meas}]$, $y_{ini}^{sim} = [y_{k_1-T_{ini}}^{meas} \ldots y_{k_1-1}^{meas}]$ and a_0 and a_1 are the regularisation and relaxation weights. The non-negative decision variables ε^f and ε^{ini} are used to relax the equality conditions. The simulation inaccuracy is then defined as the Root-Mean-Squared (RMS) difference between the simulated trajectory y_f^{sim} and the measured data $[y_{k_1}^{meas} \ldots y_{k_1+N-1}^{meas}]$. This can be repeated for multiple selections of k_0 and k_1 with the average RMS outcome taken as the performance measure denoted bb_{err} (the blackbox output) for a given set of control parameters $\theta = [T_{ini}, a_1, a_2]$ (the blackbox input). The MADS algorithm is used here to sample different values of θ to determine the parameters that produce the minimal bb_{err}. These parameters are then used to define the DeePC control problem. The blackbox system is summarised in Algorithm 1, where the sets defining the test points for k_0 and k_1 are defined as K_0 and K_1 respectively.

Algorithm 1 Black box for parameter tuning

for $k_0 \in K_0$ **do**
  $u_{tr}^{sim} = \left[u_{k_0}^{meas} \ldots u_{k_0+T_0-1}^{meas}\right]$
  $y_{tr}^{sim} = \left[y_{k_0}^{meas} \ldots y_{k_0+T_0-1}^{meas}\right]$
  Update U_p, U_f, Y_p, Y_f
  for $k_1 \in K_1$ **do**
    Update y_{ini}^{sim}, u_{ini}^{sim}, u_f^{sim}
    Solve (5) to find g_{sim}^*
    $y_f^{sim} = Y_f g_{sim}$
    $y_{err}(k_0, k_1) = RMS(y_f^{sim} - [y_{k_1}^{meas} \ldots y_{k_1+N-1}^{meas}])$
  end for
end for
$bb_{err}(\theta) = \sum_{k_0 \in K_0, k_1 \in K_1} y_{err}(k_0, k_1)$

3. Case studies

3.1. Parameter optimisation using MADS

To illustrate the use of MADS and the simulation strategy of (5), real heat supply and internal air temperature data were taken from a dwelling and a data-driven simulation was carried out with parameters tuned using the proposed blackbox optimisation approach. The performance of the data-driven simulation was then compared to the measured data. Two choices of k_0 and ten choices of k_1 were used in the blackbox optimisation. At each 15 minute timestep over a 25-day period (a different period to that used for parameter tuning), a data-driven 10-hour prediction trajectory of the air temperature was generated using (5) and knowledge of the measured heat supply trajectory for the period, shown in Fig. 1.

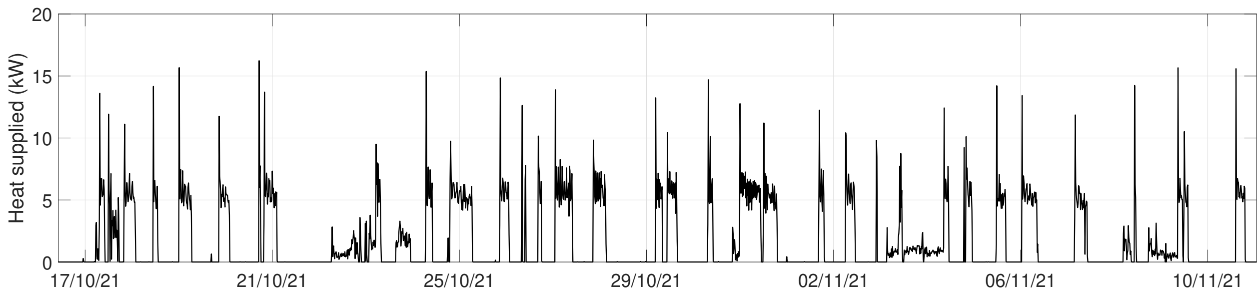

Figure 1: Measured heat supply

The resulting data-driven simulation performance is shown here for a 26-day period. The 1-step-ahead (15 minutes), 10-step-ahead (2.5 hours) and 20-step-ahead (5 hours) predictions are plotted against the measured data in Fig. 2.

It can be seen that the simulation accuracy attained is quite high. The 15-minute-ahead prediction is indistinguishable from the measured values, while a small over-estimation of the temperature can be seen in the 5-hour-ahead and 10-hour-ahead predictions. This level of accuracy was achieved despite the fact that no external disturbances such as weather or occupancy were used in the data-driven prediction. This is a promising outcome as it indicates that useful prediction can be made without the need for bespoke modelling and without the need for weather forecasting tools. Of course, if forecasts of external conditions (or any other influencing factors) were available, an improved predictive performance could be achieved. Nonetheless, the results suggest that improved performance over a non-predictive control baseline could be achieved with very

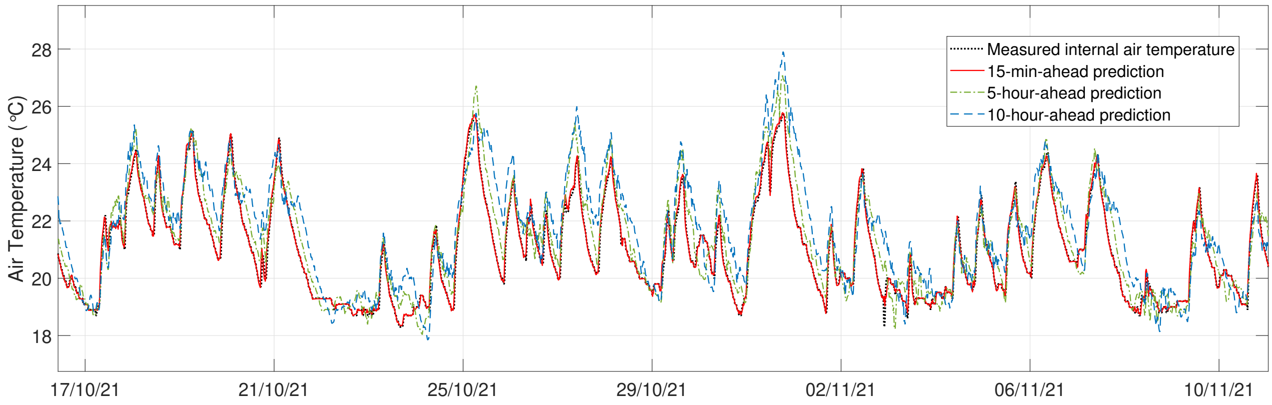

Figure 2: Data-driven simulated performance against measured data

low deployment effort.

While real, measured data can be used to show the ability of MADS to tune the data-driven simulation approach, the control performance can only be assessed by using a detailed simulation environment to analyse the closed-loop behaviour. The development of such an environment is described in the next section.

3.2. Simulation framework

To demonstrate the concept and the impact of this tuning approach, an analysis of the control performance is required. In particular, a comparison is needed between the DeePC approach with and without the use of the proposed tuning strategy. A simulation environment is needed to enable a consistent comparison between scenarios. Here, a detailed physics-based building simulation framework is employed. The simulation framework is designed to incorporate external weather conditions as well as internal gains to ensure realistic behaviour. The framework generates data which are fed to the data-driven controller at 15-minute intervals. At each interval, the controller uses these data to determine an optimal heat supply trajectory which is then passed back to the simulator.

Models of two buildings were generated using Energyplus [17] and translated to Matlab using the BRCM toolbox [18]. The building dimensions are identical with 5 rooms in each, but the first has low levels of insulation while the second is well-insulated. In the modelled heating systems, heat pumps generate heat which is delivered to the buildings via radiators which are sized to account for the different insulation levels present. DeePC controllers are applied in each room to manage the radiator temperature setpoints to ensure that the air temperature in each room remains as low as possible without violating pre-defined comfort bounds where possible. Two buildings are used to emphasise the point that different building properties will impact the choice of DeePC tuning parameters necessitating bespoke tuning. A diagram of the building design is shown in Fig. 3, while the building properties are summarised in Table 1 and Table 2.

Table 1: 1-bedroom flat construction details: low insulation

	Thermal transmittance W/m^2K	Heat capacitance (per m^2 surface area) kJ/m^2K	Solar Factor (G-value)	Area m^2
Roof	0.285	53.3	-	-
Floor	0.388	301.6	-	-
Ext. walls	0.987	460.9	-	-
Internal walls	2.581	33.1	-	-
Glazing	1	-	0.5	6.2

The resulting building models can be represented as linear state-space models of 102 states, 6 controlled inputs (the heat supply to each room), 6 outputs (room air temperatures) and 13 disturbances (internal gains for each room from occupants and appliances, ground temperature, air temperature and solar gains for each sun-facing surface). An additional random perturbation is applied to the states as an additional source of uncertainty. The simulation models can then be represented in the form:

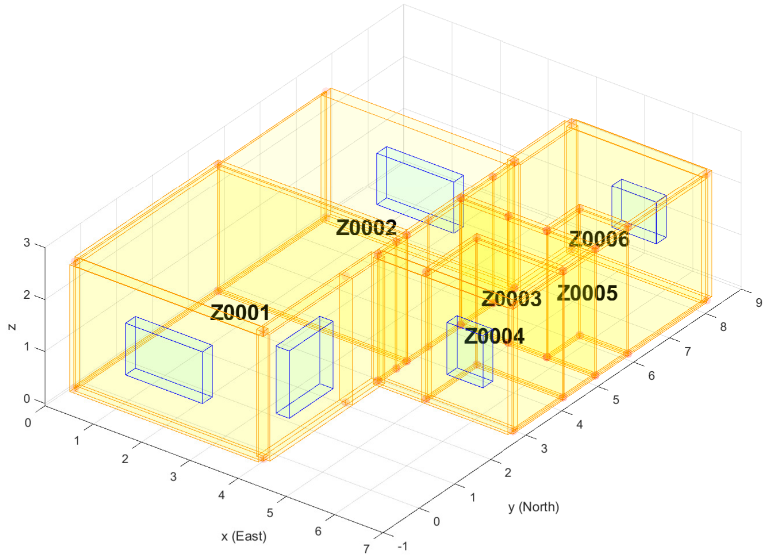

Figure 3: 5-room dwelling modelled using Energyplus and the BRCM toolbox

Table 2: 1-bedroom flat construction details: high insulation

	Thermal transmittance W/m^2K	Heat capacitance (per m^2 surface area) kJ/m^2K	Solar Factor (G-value)	Area m^2
Roof	0.165	80.9	-	-
Floor	0.153	305.1	-	-
Ext. walls	0.187	446.1	-	-
Internal walls	2.581	33.1	-	-
Glazing	1	-	0.5	6.2

$$x[k+1] = A_b x[k] + B_b u[k] + E_b d[k] + \omega[k] \tag{7}$$
$$y[k] = C_b x[k] + \nu[k], \tag{8}$$

where $A_b \in \mathbb{R}^{106\times106}$, $B_b \in \mathbb{R}^{106\times6}$, $E_b \in \mathbb{R}^{106\times13}$ and $C_b \in \mathbb{R}^{6\times106}$ are parameter matrices, $x \in \mathbb{R}^{106}$, $u \in \mathbb{R}^{6}$, $y \in \mathbb{R}^{6}$ and $d \in \mathbb{R}^{13}$ are states, inputs, outputs and disturbances respectively and $\omega \in \mathbb{R}^{106}$ and $\nu \in \mathbb{R}^{6}$ are Gaussian process and measurement noise terms with zero mean and a variance of 0.04, chosen to account for other sources of heat flow in the system as well as sensor noise.

3.3. Scenario description

Using the simulation environment outlined in Section 3.2., three separate 30-day scenarios were simulated using the same weather data. During weekdays, the occupancy periods are from 07:15-08:15, and 18:30-23:15 while at weekends, the building is occupied from 05:15-19:30 and 21:45-00:00. These schedules are taken from the outcomes of the approach of Buttitta et al. [19]. In the first scenario, a standard Proportional-Integral (PI) controller is used in each room to control the heat flowing from the radiators. The minimum temperature comfort bound is used as the setpoint for the PI controllers. This comfort bound is set at 20°C during occupied periods and 16°C during unoccupied periods. This is the baseline scenario, representing the standard performance that could be expected without a predictive strategy.

In the second and third scenarios, a DeePC controller is used in each room to send setpoints to the lower level PI controllers at 15-minute intervals. The DeePC controllers in the second scenario are parameterised using the method proposed in Section 2., while in the third scenario, the optimised parameters from room 5 in building 2 were applied to all rooms in both buildings. The purpose of this is to represent a scenario in which bespoke tuning is not carried out, but a generic pre-defined set of parameters is used in all rooms instead.

The DeePC controllers are configured to find the lowest temperature setpoint trajectory that remains within the comfort bounds over a prediction horizon of 5-hours. As well as the lower comfort bound described for the PI scenario, an upper comfort bound of 22°C for occupied periods and 24°C for unoccupied times is also applied. Typically, energy consumption would also be minimised in such a strategy, however, the resulting trade-off between comfort and energy would complicate the comparison between scenarios and so there is no penalty on energy use or energy cost applied here. A segmented DeePC formulation as described in [14] is used to provide more consistent behaviour. A prediction horizon of 20 steps, corresponding to 5 hours was chosen.

In the DeePC scenarios, PI control is used for the first 5 days of simulation so as to gather data for the autotuning procedure. The DeePC controller is then used for 30 days. The trajectories used to build the Hankel matrices are taken from the most recent data measurements, with a check carried out to ensure the persistent excitation condition is met. Three trajectories are averaged together to better handle noise. The length of the trajectories is dictated by T_0, which is chosen here to ensure the full Hankel structure is square, following the observations of performance in power system [20] and quadcopter [21] applications. For the DeePC strategy with autotuning, two choices of k_0 and ten choices of k_1 are used in the blackbox optimisation. These choices were arbitrary, limited in this case study by the quantity of data available and the time taken to run the optimisation. MADS was implemented in Matlab using NOMAD v3.9.1. The average time taken for the tuning optimisation in each room was 21 mins, with a maximum of 26 mins and a minimum of 9 mins on a desktop with a 2.9GHz processor.

3.4. Performance analysis

The comfort violations measured in each room from the 30-day simulation were collected for the three scenarios. In Fig. 4a and 4b these violations are shown in units of °C.h whereby a temperature deviation of 1°C outside the comfort bounds for 1 hour would contribute 1°C.h to the total. It can be seen that the controllers tuned using the MADS-based strategy outperform the alternatives in every room, apart from room 5 of building 2, for which the violations of both DeePC scenarios are identical. This is expected since the parameters for this room were the same for both DeePC scenarios. The PI controlled rooms had the largest violations in all rooms.

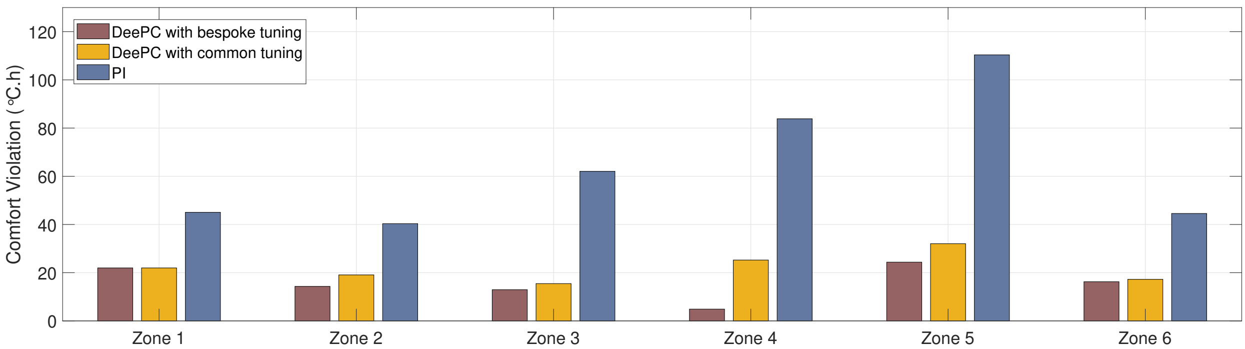

(a) Comfort violations for each scenario in each room of building 1

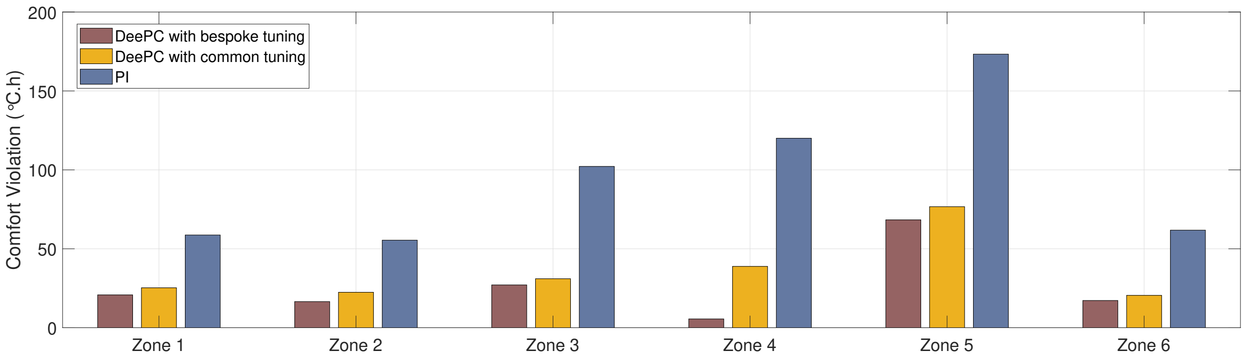

(b) Comfort violations for each scenario in each room of building 2

To better understand the behaviour of the different scenarios, Fig. 5 zooms in on the simulated room temperatures for a period of approximately 30 hours, averaged across all zones and both buildings. The dashed black line represents the lower setpoint bound. With no predictive capability, the PI controller commences heating only after the setpoint changes. The lack of pre-heating leads to comfort violations at the beginning of the occupied period. The DeePC strategies can predict the need for pre-heating, however, the pre-heating period of the DeePC controller without bespoke tuning is too short in some instances to heat the room sufficiently.

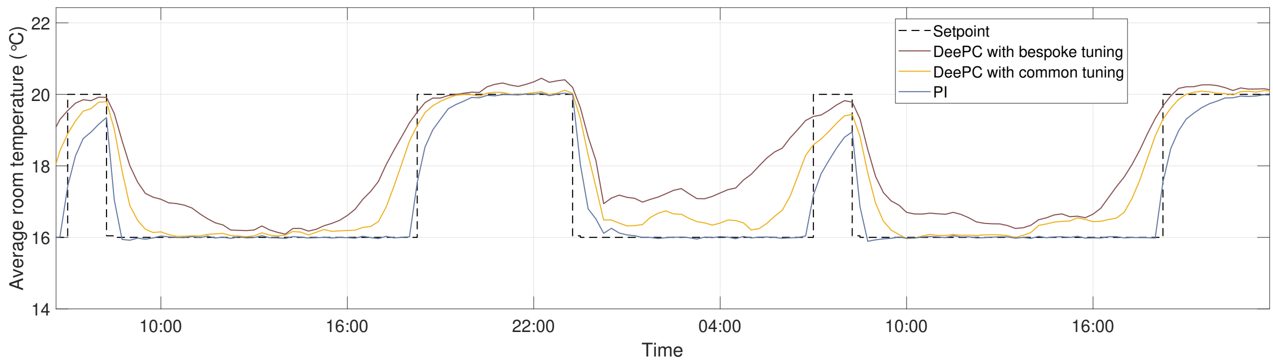

Figure 5: Snapshot of simulated air temperature averaged across all zones and buildings under different control scenarios

The total comfort violations across the two buildings are shown in Table 3. Summed over both buildings, the autotuned DeePC strategy reduced violations by 74% compared to the PI strategy and by 28% compared to the DeePC strategy with common tuning across all rooms.

Table 3: Total comfort violations for 30 simulated days

	DeePC with autotuning (°C.h)	DeePC no autotuning (°C.h)	PI (°C.h)
Building 1	94.7	131.0	386.2
Building 2	155.6	214.9	571.6

The results emphasise the need for appropriate parameterisation to be carried out. Designing a controller for one system and applying it to another without adapting the parameters may not ensure good performance.

3.5. Further discussion

Some additional points are worth noting here. Firstly, the exact implementation details used in this study are not crucial to the overall concept. Different blackbox optimisation methods may be used as well as different performance criteria, rather than the averaged RMS error. Indeed, minimising the worst-case error may offer more robust control and is a topic for future investigation. Alternative controller formulations may also be applied. Many versions of the DeePC approach now exist with new options regularly proposed in recent times.

Furthermore, the performance of the approach is dependent on the training data gathered. In this study, persistent excitation was checked in the training data by ensuring the satisfaction of the Hankel matrix rank condition, but a more thorough examination of the data quality may be fruitful. It may also be prudent to update the parameters online as new data is collected. Additional work is needed to better analyse the relationship between the simulation performance and control performance. Although it was shown in this study that choosing parameters to achieve good simulation performance tends to lead to good control performance, it may not necessarily be the best parameter set for control in all cases.

Finally, an interesting point to note is that the DeePC results show significant improvement over the baseline, despite the fact that external conditions are not considered in the predictions made by the controller. Certainly, it could be expected that by incorporating forecasts of weather, internal gains etc., the performance could be improved further, however, the results indicate that even without such knowledge, improved control performance can be achieved. This is of particular significance if the goal is to develop controllers that are quick and easy to implement.

4. Conclusions

Data-driven predictive control methods such as DeePC can expedite the roll-out of intelligent control in the built environment by removing the need for building-specific modelling to be done as part of the control design process. This has been noted as a significant barrier. Nonetheless, while the modelling component of the implementation process is removed, a control design component remains. Specifically, parameters and weights must be selected to provide good performance in any given system and the behaviour of DeePC can be quite sensitive to poor choices of these parameters. This paper proposes a method for automating this parameter selection process. By formulating a data-driven simulation problem with the same structure as the data-driven controller, a blackbox optimisation problem is formed in which the inputs to the blackbox are the control parameters and the output is the simulation inaccuracy. A blackbox optimisation method, namely MADS, is used to find the optimal parameter set in an automated fashion. Based on these parameters, DeePC controllers are designed to control the air temperature in a building.

By testing the data-driven optimisation approach on real data, it was found that short term predictions were achievable, though as external influencing factors such as weather were omitted, certain inaccuracies appear. The complete parameter optimisation and control strategy was tested using a simulation environment in which two buildings, each with 6 rooms were modelled. DeePC controllers in each room were used to maintain thermal comfort. The controllers were compared with and without autotuning, with an additional comparison to a non-predictive PI-based approach also carried out. By autotuning the controllers to suit their relevant rooms, a reduction in comfort violation of 28% was achieved in total across the two buildings compared to the controllers that had not been autotuned. A 74% reduction in comfort violation was achieved compared to the PI-based approach. The results emphasised the significant influence of the control parameter choice on the overall performance, indicating that bespoke tuning may be needed and off-the-shelf controllers with pre-defined parameters may not perform well across different systems.

Acknowledgments

This work has received funding from the EPSRC (Engineering and Physical Sciences) under the Active Building Centre project (reference number: EP/V012053/1).

References

[1] Lowes R, Rosenow J, Qadrdan M, Wu J. Hot stuff: Research and policy principles for heat decarbonisation through smart electrification. Energy Research and Social Science. 2020;70(March):101735.

[2] Drgoňa J, Arroyo J, Cupeiro Figueroa I, Blum D, Arendt K, Kim D, et al. All you need to know about model predictive control for buildings. Annual Reviews in Control. 2020;50:190-232.

[3] Henze GP. Model predictive control for buildings: a quantum leap? Journal of Building Performance Simulation. 2013;6(3):157-8.

[4] Maddalena ET, Lian Y, Jones CN. Data-driven methods for building control — A review and promising future directions. Control Engineering Practice. 2020;95:104211.

[5] Kathirgamanathan A, De Rosa M, Mangina E, Finn DP. Data-driven predictive control for unlocking building energy flexibility: A review. Renewable and Sustainable Energy Reviews. 2021;135:110120.

[6] O'Dwyer E, Pan I, Acha S, Shah N. Smart energy systems for sustainable smart cities: Current developments, trends and future directions. Applied Energy. 2019 January;237:581-97.

[7] Willems JC, Rapisarda P, Markovsky I, De Moor BLM. A note on persistency of excitation. Systems and Control Letters. 2005;54(4):325-9.

[8] De Persis C, Tesi P. Formulas for Data-Driven Control: Stabilization, Optimality, and Robustness. IEEE Transactions on Automatic Control. 2020;65(3):909-24.

[9] Coulson J, Lygeros J, Dorfler F. Data-enabled predictive control: In the shallows of the deepc. In: 2019 18th European Control Conference, ECC 2019. EUCA; 2019. p. 307-12.

[10] Huang L, Zhen J, Lygeros J, Dörfler F. Robust Data-Enabled Predictive Control: Tractable Formulations and Performance Guarantees. arXiv preprint arXiv:210507199. 2021.

[11] Dörfler F, Coulson J, Markovsky I. Bridging direct & indirect data-driven control formulations via regularizations and relaxations. arXiv preprint arXiv:210101273. 2021.

[12] Berberich J, Köhler J, Müller MA, Allgöwer F. Data-Driven Model Predictive Control With Stability and Robustness Guarantees. IEEE Transactions on Automatic Control. 2021;66(4):1702-17.

[13] Lian Y, Shi J, Koch MP, Jones CN. Adaptive Robust Data-driven Building Control via Bi-level Reformulation: an Experimental Result. arXiv preprint arXiv:210605740. 2021:1-12.

[14] O'Dwyer E, Kerrigan EC, Falugi P, Zagorowska MA, Shah N. Data-driven predictive control with reduced computational effort and improved performance using segmented trajectories. arXiv preprint arXiv:210810753. 2021:1-11.

[15] Audet C, Le Digabel S, Tribes C. The mesh adaptive direct search algorithm for granular and discrete variables. SIAM Journal on Optimization. 2019;29(2):1164-89.

[16] Lian Y, Jones CN. From System Level Synthesis to Robust Closed-loop Data-Enabled Predictive Control. In: 2021 60th IEEE Conference on Decision and Control (CDC). IEEE; 2021. p. 1478-83.

[17] US Department of Energy. EnergyPlus; 2018 [Online]. Available from: https://energyplus.net/.

[18] Sturzenegger D, Gyalistras D, Semeraro V, Morari M, Smith RS. BRCM Matlab Toolbox: Model generation for model predictive building control. Proceedings of the American Control Conference. 2014:1063-9.

[19] Buttitta G, Turner WJN, Neu O, Finn DP. Development of occupancy-integrated archetypes: Use of data mining clustering techniques to embed occupant behaviour profiles in archetypes. Energy and Buildings. 2019 sep;198:84-99.

[20] Huang L, Coulson J, Lygeros J, Dörfler F. Decentralized Data-Enabled Predictive Control for Power System Oscillation Damping. IEEE Transactions on Control Systems Technology. 2021:1-13.

[21] Elokda E, Coulson J, Beuchat PN, Lygeros J, Dörfler F. Data-enabled predictive control for quadcopters. International Journal of Robust and Nonlinear Control. 2021;31(18):8916-36.