

Data-driven predictive control with improved performance using segmented trajectories

E. O'Dwyer*, E. C. Kerrigan, P. Falugi, M. A. Zagorowska and N. Shah

Abstract—A class of data-driven control methods has recently emerged based on Willems' fundamental lemma. Such methods can ease the modelling burden in control design but can be sensitive to disturbances acting on the system under control. In this paper, we extend these methods to incorporate segmented prediction trajectories. The proposed segmentation enables longer prediction horizons to be used in the presence of unmeasured disturbance. Furthermore, a computation time reduction can be achieved through segmentation by exploiting the problem structure, with computation time scaling linearly with increasing horizon length. The performance characteristics are illustrated in a set-point tracking case study in which the segmented formulation enables more consistent performance over a wide range of prediction horizons. The computation time for the segmented formulation is approximately half that of an unsegmented formulation for a horizon of 100 samples. The method is then applied to a building energy management problem, using a detailed simulation environment, in which we seek to minimise the discomfort and energy of a 6-room apartment. With the segmented formulation, a 72% reduction in discomfort and 5% financial cost reduction is achieved, compared to an unsegmented formulation using a one-day-ahead prediction horizon.

Index Terms—Data-driven predictive control, optimal control, building energy management, Willems' fundamental lemma

I. INTRODUCTION

THE increased focus on digital technology in recent times has drawn attention to data-driven control methods, with applications ranging from building control [1], to autonomous vehicles [2]. By reducing the modelling burden in the control design phase, the deployment of advanced control can be streamlined, with control decisions directly obtained from data measurements [3]. In data-rich environments, such methods may then be advantageous. Nonetheless, these approaches can often lack the theoretical foundations of both model-based and

indirect data-driven methods, such as system-identification, whereby data is used to derive a model for control [4].

Willems' fundamental lemma [5], has been used as a foundation for a large body of data-driven control research, with predictive control representations recently developed that can offer stability and certain robustness guarantees [4], without requiring the derivation of a parametric model. For example, a data-enabled predictive control formulation was proposed in [6] and shown to be competitive with Model Predictive Control in the deterministic case. In [7], an equivalence was then shown between this direct data-driven approach and an alternative indirect approach, in which the parameters of a multi-step prediction model are derived using the same training data criteria. This was expanded upon in [8], where further analysis of the performance of these direct and indirect formulations was carried out for different systems using various relaxation and regularisation techniques. The results suggested that noisy data have a greater impact on direct formulations, while system nonlinearities have a greater impact on indirect formulations. Additionally, rather than relying on a single training period for data acquisition, a strategy was developed in [9] by which multiple, potentially short, datasets can be used instead, without compromising the theoretical foundation of the fundamental lemma.

Uncertainty in data measurements will impact the performance of a data-driven approach, thus several methods have been developed to ensure viability in stochastic settings. In [10], the authors supplement a data-driven controller with a data-driven extended Kalman filter to reduce sensitivity to noise. Robust formulations have also been developed to enable performance guarantees under certain conditions of system stochasticity, such as the robust modification proposed in [11], ensuring exponential stability in the presence of measurement noise. In [12], a chance-constrained distributionally robust formulation was developed for stochastic linear, time-invariant (LTI) systems, providing probabilistic guarantees on performance. Tractable, robust formulations are proposed to ensure performance guarantees under uncertainty in [13], while in [14], a correspondence was found between the fundamental lemma perspective and that of System-Level Synthesis (SLS), which was then exploited to formulate a robust closed-loop data-predictive controller. A robust building-level controller was implemented in [15].

Apart from LTI systems, approaches based on the fundamental lemma have been extended to include systems with time-varying parameters [16]. Furthermore, by combining

* Corresponding author.

E. O'Dwyer is with the Department of Chemical Engineering, Imperial College London, email: e.odwyer@imperial.ac.uk

P. Falugi is with the Department of Electrical and Electronic Engineering, Imperial College London, email: p.falugi@imperial.ac.uk

E. C. Kerrigan is with the Department of Electrical and Electronic Engineering, Department of Aeronautics, Imperial College London, email: e.kerrigan@imperial.ac.uk

M. A. Zagorowska is with the Department of Electrical and Electronic Engineering, Imperial College London, email: m.zagorowska@imperial.ac.uk

N. Shah is with the Department of Chemical Engineering, Imperial College London, email: n.shah@imperial.ac.uk

these approaches with the perspective of Koopman operator theory, it was shown in [17] that control of nonlinear systems can be achieved within a data-driven framework. Integral action was additionally incorporated into approaches proposed in [18] and [19], thus ensuring offset-free tracking.

In the presence of noise or unmeasured disturbances, data-driven models can be prone to overfitting, particularly if they are over-parameterised [20]. An under-explored consideration in the development of data-driven predictive controllers of this type is the link between the prediction trajectory length and the number of parameters, and how this relates to the control performance. Though parameters are not explicitly identified in direct formulations, an implicit model identification step is carried out in the regularised form as discussed in [7] and [8], leading to an equivalent, implicit, parametric representation. Prediction of longer trajectories may lead to an impaired performance under uncertainty, since the number of parameters in this representation is related to the length of the prediction horizon in the standard data-enabled predictive controller formulations.

In this paper, we present a formulation in which the prediction trajectory is divided into multiple shorter trajectories (denoted segments). These segments can be identified in the same manner as the unsegmented formulation, with less training data and less computational effort by exploiting the problem structure. This novel segmentation approach decouples the link between the number of implicit model parameters and the prediction horizon length, thus making the formulation less sensitive to noise and disturbance in the training data. The method is analysed and compared to the unsegmented version in a set-point tracking case study with different regularisation parameters and horizon lengths. A performance improvement is shown in terms of set-point tracking error and computational time requirement. Following this, a building energy management case study is implemented, based on a detailed building simulation environment with realistic disturbances. Comfort and energy cost objectives are solved in a prioritised manner. As the segmented-trajectory approach behaves more consistently than the unsegmented approach for longer horizon lengths, a reduction in both cost and energy consumption is achieved for a one-day-ahead prediction horizon.

In Section II, a background to the unsegmented data-driven predictive approach is provided based on the fundamental lemma, followed by the proposed modifications that result in a segmented formulation. In Section III, a set-point tracking case study is presented, with an analysis provided of the control performance and computational time associated with the segmented and unsegmented formulations. This is followed in Section IV by the building energy management case study, which is used to illustrate the benefits of the segmented formulation in a relevant application. The paper ends with conclusions in Section V.

II. MODIFIED DATA-DRIVEN PREDICTIVE CONTROL FORMULATION

A. Data-driven predictive control preliminaries

As noted in the introduction, different variations of data-predictive control have been proposed. We provide a brief

overview of the direct data-enabled predictive controller of [6] and the indirect multi-step prediction approach of [7] and [8] in this section.

A discrete-time n^{th} -order LTI state-space system can be represented at sample instant k by:

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

where $x[k] \in \mathbb{R}^n$ is the system state-vector, $u[k] \in \mathbb{R}^m$ and $y[k] \in \mathbb{R}^p$ are the input and output vectors, respectively, 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, the parameters of which are assumed to be unknown. Following the terminology of [21], we define $\mathcal{B}$ as the *behaviour* of (1), where the behaviour is defined as the set of possible outcomes of the system. The lag of the system is denoted ℓ , defined as the smallest integer for which the observability matrix $\mathcal{O}_\ell(A, C) := [C, CA, \dots, CA^{\ell-1}]$ has full rank. By Willems' fundamental lemma [5], arbitrary input and output sequences can be derived from a sufficiently long set of input/output data without explicitly estimating the parameters of (1). The input output sequences will be called *trajectories*. A representation of the system can then be found and used for predictive control only in terms of measured data.

An offline data collection procedure is carried out to achieve this in which $T_0 \in \mathbb{Z}_{>0}$ sequences of *persistently exciting* input and output data measurements are given as $u_{tr} = [u_1^T, \dots, u_{T_0}^T]^T \in \mathbb{R}^{mT_0}$ and $y_{tr} = [y_1^T, \dots, y_{T_0}^T]^T \in \mathbb{R}^{pT_0}$ respectively with $\mathbb{Z}_{>0}$ denoting the set of positive integers. A trajectory w is defined as persistently exciting of order L_0 , $L_0 \in \mathbb{Z}_{>0}$, if the Hankel matrix $\mathcal{H}_{L_0}(w)$ is of full row-rank with

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

Note that non-square Hankel matrices are permitted in this definition. From [5], for a controllable, observable $\mathcal{B}$, if $w \in \mathcal{B}$ is a persistently exciting, T_0 -samples-long trajectory of order $t+n$, then any t -samples long trajectory in $\mathcal{B}$ can be described as a linear combination of the columns of $\mathcal{H}_t(w)$, and any $\mathcal{H}_t(w)g$ is a trajectory of $\mathcal{B}$ where $g \in \mathbb{R}^{t+1}$. For persistent excitation, $T_0 \geq (m+1)(t+n)-1$. Here we seek to construct trajectories of length $N+T_{ini}$, where $N \in \mathbb{Z}_{>0}$ is the prediction horizon and $T_{ini} \in \mathbb{Z}$ is some initialisation length. Following [22, Lem. 1], by fixing the first T_{ini} samples of a trajectory, the subsequent N samples are uniquely specified if $T_{ini} \geq \ell$.

The training data sequences u_{tr} and y_{tr} are arranged in the Hankel form of (2) with $T_0 \geq (m+1)(T_{ini} + N + n) - 1$ and $L_0 = T_{ini} + N$. The training data structures can then be defined at this point as $U_{tr} := \mathcal{H}_{T_{ini}+N}(u_{tr})$ and $Y_{tr} := \mathcal{H}_{T_{ini}+N}(y_{tr})$. These matrices are then partitioned such that the first T_{ini} block rows of U_{tr} and Y_{tr} are denoted by the subscript α and are referred to as initialisation data, with the remaining rows denoted by β and referred to as prediction

data. The partitioned data matrices are thus defined as

$$\begin{aligned} \begin{bmatrix} U_\alpha \\ U_\beta \end{bmatrix} &:= \mathcal{H}_{T_{ini}+N}(u_{tr}), \\ \begin{bmatrix} Y_\alpha \\ Y_\beta \end{bmatrix} &:= \mathcal{H}_{T_{ini}+N}(y_{tr}). \end{aligned} \quad (3)$$

Defining initialisation sequences $u_{ini} \in \mathbb{R}^{mT_{ini}}$ and $y_{ini} \in \mathbb{R}^{pT_{ini}}$ as the T_{ini} most recent measurements, any future trajectories $u_f \in \mathbb{R}^{mN}$ and $y_f \in \mathbb{R}^{pN}$ can be found as the solution to

$$\begin{bmatrix} U_\alpha \\ U_\beta \\ Y_\alpha \\ Y_\beta \end{bmatrix} g = \begin{bmatrix} u_{ini} \\ u_f \\ y_{ini} \\ y_f \end{bmatrix}, \quad (4)$$

where $g \in \mathbb{R}^{T_0-T_{ini}-N+1}$.

This leads to the insight that u_f and y_f , the future trajectories of $\mathcal{B}$, can be found for a given training data set and given initialisation trajectories u_{ini} and y_{ini} . From this, a Data-enabled Predictive Control (DeePC) formulation was proposed in [6], whereby the following optimisation is carried out:

$$\min_{g, u_f, y_f} V(g, u_f, y_f) \quad (5)$$

s.t.

$$\begin{bmatrix} U_\alpha \\ U_\beta \\ Y_\alpha \\ Y_\beta \end{bmatrix} g = \begin{bmatrix} u_{ini} \\ u_f \\ y_{ini} \\ y_f \end{bmatrix} \quad (6)$$

$$u_f \in \mathcal{U} \quad (7)$$

$$y_f \in \mathcal{Y} \quad (8)$$

with $V(\cdot)$ representing an objective to be minimised and $\mathcal{U}$ and $\mathcal{Y}$ representing the input and output constraint sets, respectively.

Whereas model parameters are not explicitly derived in this formulation, an equivalence was identified in [7] and [8] between this form and a predictive control formulation based on a multi-step prediction model derived from data. This multi-step model version is referred to in [8] as an indirect data-driven formulation, in contrast to direct data-driven formulations in which no model is identified such as in (5)–(8).

Using the indirect formulation, the same training data structures can be used, but here they are used to derive a multi-step predictor P^* by the least-squares method as

$$P^* = \arg \min_P \left\| P \begin{bmatrix} U_\alpha \\ U_\beta \\ Y_\alpha \end{bmatrix} - Y_\beta \right\|_F^2, \quad (9)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Using the Moore-Penrose inverse (denoted $\dagger$), this can be expressed explicitly as

$$P^* := Y_\beta \begin{bmatrix} U_\alpha \\ U_\beta \\ Y_\alpha \end{bmatrix}^\dagger. \quad (10)$$

This predictor can then be used to derive future trajectories as

$$y_f = P^* \begin{bmatrix} u_{ini} \\ u_f \\ y_{ini} \end{bmatrix}. \quad (11)$$

To examine the model defined by (11) in Section II-B, it is useful here to define a partitioned version of P^* , given as $[P_1^* \ P_2^* \ P_3^*]$ where $P_1^* \in \mathbb{R}^{pN \times mT_{ini}}$ is associated with the initialisation input trajectory, $P_2^* \in \mathbb{R}^{pN \times mN}$ is associated with the future input trajectory and $P_3^* \in \mathbb{R}^{pN \times mT_{ini}}$ is associated with the initialisation output trajectory.

An indirect data-driven predictive control formulation can then be represented by replacing (4) with (11). Additional conditions can be applied to enforce causality in the model.

In the following section, both the direct and indirect methods summarised here will be used to illustrate the rationale of a modified version of the data-predictive control approach, which is the main contribution of this work.

B. Segmentation of prediction trajectory

Over-parameterisation of a model can impede prediction performance in the presence of uncertainties, such as noise, nonlinearities, or unmeasured disturbances in the system under control. In the context of the data-driven controller described in Section II-A, relaxations of the initialisation constraints and regularisation of the optimisation variables can be introduced to improve performance in this regard in the direct formulation. Similarly, slack variables can be introduced to the indirect form.

Nonetheless, the number of parameters of $P_2^* \in \mathbb{R}^{pN \times mN}$ in the indirect form increases with the horizon length. For example, the final entry of the predicted output sequence, $y_f[N]$, is a function of N preceding inputs ($u_f[1], \dots, u_f[N]$). Thus, for longer prediction horizons, the impact of over-parameterisation may become more pronounced, as will be shown in the illustrative example in Section III. In the direct formulation, model parameters are not explicitly derived; however, in [7] it is reasoned that the same model is implicitly identified in the direct form as the indirect form leading to the same performance drop for longer horizons. It should be noted that with perfect data, the fixing of u_{ini} and y_{ini} , ensures a unique representation of u_f and y_f and thus the issue does not arise.

A modification is proposed here whereby the prediction trajectory is divided into segments of length T_{ini} (with $T_{ini} \leq N$) to decouple the relationship between horizon length and the number of parameters (implicit parameters in the direct formulation, explicit parameters in the indirect formulation), thereby ensuring better scalability to problems with longer prediction horizons. The key insight here is that by assuming the system does not change over the prediction horizon, we can construct the full horizon using shorter trajectories. Each *prediction* trajectory acts as the *initialisation* trajectory for its subsequent segment.

The shorter trajectories used in this formulation necessitate a change in the training data matrix definitions. In Section II, T_0 dictates the training length, where the conditions $T_0 \geq$

$(m+1)(T_{ini}+N+n)-1$ and $T_{ini} \geq \ell$ were imposed. For the segmented form, we replace T_0 with T_a . Since each segment is at most of length T_{ini} , rather than N (the final segment may be shorter), we now impose $T_a \geq (m+1)(2T_{ini}+n)-1$. Notably, for $T_{ini} < N$ a shorter training period is then sufficient.

The updated training sequences are then defined as $u_{tr_s} = [u_1^T, \dots, u_{T_a}^T]^T \in \mathbb{R}^{mT_a}$ and $y_{tr_s} = [y_1^T, \dots, y_{T_a}^T]^T \in \mathbb{R}^{pT_a}$ and the associated Hankel matrices are defined as

$$\begin{aligned} \begin{bmatrix} U_{\alpha_s} \\ U_{\beta_s} \end{bmatrix} &:= \mathcal{H}_{2T_{ini}}(u_{tr_s}), \\ \begin{bmatrix} Y_{\alpha_s} \\ Y_{\beta_s} \end{bmatrix} &:= \mathcal{H}_{2T_{ini}}(y_{tr_s}), \end{aligned} \quad (12)$$

with $U_{\alpha_s} \in \mathbb{R}^{mT_{ini} \times (T_a - 2T_{ini} + 1)}$, $Y_{\alpha_s} \in \mathbb{R}^{pT_{ini} \times (T_a - 2T_{ini} + 1)}$.

The trajectories u_f and y_f are partitioned into F segments given as $[u_{f_1}^T, \dots, u_{f_F}^T]^T = u_f$ and $[y_{f_1}^T, \dots, y_{f_F}^T]^T = y_f$ respectively, where $u_{f_i} \in \mathbb{R}^{mT_{ini}}$ and $y_{f_i} \in \mathbb{R}^{pT_{ini}}$, $\forall i \in \{1, \dots, F-1\}$, and the final segments $u_{f_F} \in \mathbb{R}^{m(N-(F-1)T_{ini})}$ and $y_{f_F} \in \mathbb{R}^{p(N-(F-1)T_{ini})}$. Equation 4 can then be replaced by the following equations with u_{ini} and y_{ini} replaced by u_{f_0} and y_{f_0} , respectively, for notational brevity:

$$\begin{bmatrix} U_{\alpha_s} \\ U_{\beta_s} \\ Y_{\alpha_s} \\ Y_{\beta_s} \end{bmatrix} g_i = \begin{bmatrix} u_{f_{i-1}} \\ u_{f_i} \\ y_{f_{i-1}} \\ y_{f_i} \end{bmatrix}, \quad \forall i \in \{1, \dots, F-1\}, \quad (13)$$

where $g_i \in \mathbb{R}^{T_a - 2T_{ini} + 1}$, $\forall i \in \{1, \dots, F\}$. The equation associated with the final segment is given as

$$\begin{bmatrix} U_{\alpha_s} \\ \tilde{U}_{\beta_s} \\ Y_{\alpha_s} \\ \tilde{Y}_{\beta_s} \end{bmatrix} g_F = \begin{bmatrix} u_{f_{F-1}} \\ u_{f_F} \\ y_{f_{F-1}} \\ y_{f_F} \end{bmatrix}, \quad (14)$$

where $\tilde{U}_{\beta_s}$ and $\tilde{Y}_{\beta_s}$ represent the first $N - (F-1)T_{ini}$ block rows of U_{β_s} and Y_{β_s} , respectively. A diagram illustrating the segmentation concept is shown in Fig. 1 for a prediction trajectory divided into three segments.

Using (13) and (14) to predict future input and output trajectories of the system, we can formulate a predictive controller in which we seek to minimise a cost function given as $V_s(\cdot)$ by solving

$$\min_{g_1, \dots, g_F} \sum_{i=1}^F V_s(g_i) \quad (15)$$

s.t.

$$\begin{bmatrix} U_{\alpha_s} \\ Y_{\alpha_s} \end{bmatrix} g_1 = \begin{bmatrix} u_{f_0} \\ y_{f_0} \end{bmatrix} \quad (16)$$

$$\begin{bmatrix} -U_{\beta_s} & U_{\alpha_s} \\ -Y_{\beta_s} & Y_{\alpha_s} \end{bmatrix} \begin{bmatrix} g_{i-1} \\ g_i \end{bmatrix} = \mathbf{0}_{2T_{ini}}, \quad \forall i \in \{2, \dots, F\} \quad (17)$$

$$U_{\beta_s} g_i \in \mathcal{U}, \quad \forall i \in \{1, \dots, F\} \quad (18)$$

$$Y_{\beta_s} g_i \in \mathcal{Y}, \quad \forall i \in \{1, \dots, F\} \quad (19)$$

where $\mathbf{0}_a$ denotes a column of zeros of length a . No penalty on the input is included. It should be noted that this problem




Fig. 1: Illustration of data-driven control approach (a) without and (b) with segmentation of prediction trajectory

is partially separable, unlike the unsegmented version. With a suitable choice of optimisation solver, a linear increase in computation time can be achieved for an increasing F . This is shown empirically in Section III.

In a similar manner, this segmented perspective can be applied to the indirect formulation. We now denote the multi-step predictor matrix as P_s^* , which can be found as

$$P_s^* = Y_{\beta_s} \begin{bmatrix} U_{\alpha_s} \\ U_{\beta_s} \\ Y_{\alpha_s} \end{bmatrix}^\dagger. \quad (20)$$

The first $F-1$ segments of the prediction trajectory can then be found as

$$y_{f_i} = P_s^* \begin{bmatrix} u_{f_{i-1}} \\ u_{f_i} \\ y_{f_{i-1}} \end{bmatrix}, \quad \forall i \in \{1, \dots, F-1\}. \quad (21)$$

The final segments u_{f_F} and y_{f_F} are of length $m(N - (F-1)T_{ini})$ and $p(N - (F-1)T_{ini})$ respectively. To obtain these segments, P_s^* is partitioned as $[P_{s1}^* \ P_{s2}^* \ P_{s3}^*]$ where $P_{s1}^* \in \mathbb{R}^{pT_{ini} \times mT_{ini}}$, $P_{s2}^* \in \mathbb{R}^{pT_{ini} \times mT_{ini}}$ and $P_{s3}^* \in \mathbb{R}^{pT_{ini} \times mT_{ini}}$. If N is not a multiple of T_{ini} , this final segment will be shorter than all preceding segments and the final $p(FT_{ini} - N)$ rows must be omitted from P_{s1}^* , P_{s2}^* and P_{s3}^* and $m(FT_{ini} - N)$ columns must be omitted from P_{s2}^* . These reduced matrices are represented by the accent $\tilde{\cdot}$, and the multi-step model of this segment is then given as

$$y_{f_F} = [\tilde{P}_1^* \ \tilde{P}_2^* \ \tilde{P}_3^*] \begin{bmatrix} u_{f_{F-1}} \\ u_{f_F} \\ y_{f_{F-1}} \end{bmatrix}. \quad (22)$$




Fig. 2: Two-mass system with springs and dampers

A key feature of this segmented representation can be seen by comparing the dimensions of P_{s2}^* with P_2^* . Without segmentation, the number of parameters grows with an increasing prediction horizon. By decoupling the number of parameters from the prediction horizon length, we propose that the segmented formulation will perform better than the unsegmented formulation in problems with longer prediction horizons in the presence of unmeasured disturbance. The same underlying model is derived for all horizon lengths. The performance impact is examined in detail in the following section in which different versions of the formulation are applied to a simple example under a wide range of conditions.

III. ILLUSTRATIVE EXAMPLE: TWO-MASS SYSTEM

A. System description

A two-mass-spring-damper example is used to illustrate the performance of the segmented predictive controller compared with the unsegmented version. The code needed to reproduce these examples is available on Code Ocean. The system comprises two masses, two springs and two dampers, and is described in the following equations:

$$\dot{x}(t) = Ax(t) + B(u(t) + d(t)), \quad (23)$$

where $x = (y_1, y_2, \dot{y}_1, \dot{y}_2)$ with y_1 and y_2 representing the displacement of masses m_1 and m_2 respectively (shown in Fig 2) and $u(t)$ is the input force applied to the mass m_1 . An additional disturbance $d(t)$ can be applied to m_1 . The parameter matrices A and B are given as

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\frac{(k_1+k_2)}{m_1} & \frac{k_2}{m_1} & -\frac{(c_1+c_2)}{m_1} & \frac{c_2}{m_1} \\ \frac{k_2}{m_2} & -\frac{c_2}{m_2} & \frac{c_2}{m_2} & -\frac{c_2}{m_2} \end{bmatrix}, B = \begin{bmatrix} 0 \\ 0 \\ \frac{1}{m_1} \\ 0 \end{bmatrix}, \quad (24)$$

where the masses are defined $m_1 = 0.5$ and $m_2 = 1.5$, the spring constants are defined as $k_1 = 2$ and $k_2 = 2$ and the damping constants are defined as $c_1 = 1$ and $c_2 = 1$. The system is shown in Fig. 2.

We investigate a case study using this system whereby we seek to control the displacement y_2 to track a set-point y_{sp} , by calculating an input trajectory u using a data-driven predictive controller. We adopt a direct data-driven formulation and compare segmented and unsegmented versions of the strategy in scenarios with and without a time-varying unmeasured disturbance d applied to m_1 . The system is analysed for various hyper-parameter choices, specifically the prediction horizon and the regularisation weight. A one-second sample time is

used for the predictive controller, with input and disturbance signals held constant for the duration of the sample. The input force is constrained to the interval $[-1, 1]$ and results are compiled from a 100-second run.

B. Prioritised objective formulation

To handle the regularisation, relaxation and set-point deviation penalties, a prioritised framework is used as described in [23]. The problem is solved in two stages, first a feasibility stage followed by a set-point deviation minimisation. The first optimisation minimises the initialisation slacks given as $\varepsilon_{f_i} \in \mathbb{R}^{T_{ini}}$, $\forall i \in \{1, \dots, F\}$. The objective and constraints of this linear problem are defined as follows:

$$J_1^* := \min_{\substack{g_1, \dots, g_F, \\ \varepsilon_{f_1}, \dots, \varepsilon_{f_F}}} \sum_{i=1}^F \sum_{j=1}^{T_{ini}} \varepsilon_{f_{i,j}} \quad (25)$$

s.t.

$$U_{\alpha_s} g_1 = u_{f_0} \quad (26)$$

$$\begin{bmatrix} Y_{\alpha_s} \\ -Y_{\alpha_s} \end{bmatrix} g_1 - \begin{bmatrix} \varepsilon_{f_1} \\ \varepsilon_{f_1} \end{bmatrix} \leq \begin{bmatrix} y_{f_0} \\ -y_{f_0} \end{bmatrix} \quad (27)$$

$$\begin{bmatrix} -U_{\beta_s} & U_{\alpha_s} \end{bmatrix} \begin{bmatrix} g_{i-1} \\ g_i \end{bmatrix} = \mathbf{0}_{T_{ini}}, \forall i \in \{2, \dots, F\} \quad (28)$$

$$\begin{bmatrix} -Y_{\beta_s} & Y_{\alpha_s} \end{bmatrix} \begin{bmatrix} g_{i-1} \\ g_i \end{bmatrix} - \varepsilon_{f_i} \leq \mathbf{0}_{T_{ini}}, \forall i \in \{2, \dots, F\} \quad (29)$$

$$\begin{bmatrix} Y_{\beta_s} & -Y_{\alpha_s} \end{bmatrix} \begin{bmatrix} g_{i-1} \\ g_i \end{bmatrix} - \varepsilon_{f_i} \leq \mathbf{0}_{T_{ini}}, \forall i \in \{2, \dots, F\} \quad (30)$$

$$-\varepsilon_{f_i} \leq \mathbf{0}_{T_{ini}}, \forall i \in \{1, \dots, F\} \quad (31)$$

$$U_{\beta_s} g_i \in \mathcal{U}, \forall i \in \{1, \dots, F\} \quad (32)$$

$$Y_{\beta_s} g_i \in \mathcal{Y}, \forall i \in \{1, \dots, F\}. \quad (33)$$

The second optimisation objective is composed of a penalty on the sum of the absolute deviation of the output from the set-point, given as $\varepsilon_y \in \mathbb{R}^N$, and a regularisation penalty on g with the relative weight between the two penalties set by choice of $\lambda_g > 0$. The quadratic objective and linear constraints of this problem are given as

$$J_2^* := \min_{\substack{g_1, \dots, g_F, \varepsilon_y, \\ \varepsilon_{f_1}, \dots, \varepsilon_{f_F}}} \sum_{j=1}^N \varepsilon_{y_j} + \lambda_g \sum_{i=1}^F g_i^T g_i \quad (34)$$

s.t.

$$(26)-(33)$$

$$\sum_{i=1}^F \sum_{j=1}^{T_{ini}} \varepsilon_{f_{i,j}} \leq J_1^* \quad (35)$$

$$\begin{bmatrix} I_F \otimes Y_{\beta_s} \\ -I_F \otimes Y_{\beta_s} \end{bmatrix} \begin{bmatrix} g_1 \\ \vdots \\ g_F \end{bmatrix} - \begin{bmatrix} \varepsilon_y \\ \varepsilon_y \end{bmatrix} \leq \begin{bmatrix} y_{sp} \\ -y_{sp} \end{bmatrix} \quad (36)$$

$$-\varepsilon_y \leq \mathbf{0}_N, \quad (37)$$

where $\otimes$ denotes the Kronecker product and I_F denotes the identity matrix of size F .




Fig. 3: Training input signal and unmeasured disturbance signal applied to two-mass system




Fig. 4: Set-point tracking using the segmented formulation with and without unmeasured disturbance present and the unsegmented formulation with unmeasured disturbance present

In the unsegmented formulation, $F = 1$, constraints (28), (29) and (30) are omitted and α_s and β_s subscripts are replaced by α and β respectively.

To generate training data, the system was simulated in open-loop with the input u varied at 10-second intervals by drawing a sample from a uniform distribution in the interval $[-1, 1]$ to generate a persistently exciting training set of input and output data. A disturbance signal, which acts as the unmeasured disturbance, was also generated and applied to the system. This disturbance was composed of a sinusoidal component and a uniformly distributed random noise component. Fig. 3 shows the training input and disturbance signals.

Fig. 4 shows the set-point tracking performances of the unsegmented and segmented versions of the data predictive controller as well as a scenario in which no unmeasured disturbance was present. In all cases, $N = 30$, $T_{ini} = 5$ and $\lambda = 0.5$. The set-point to be tracked is plotted as a dashed black line. Excellent tracking performance can be observed for the undisturbed scenario. In the presence of the disturbance, the performance degrades, but the segmented version outperforms the unsegmented version. The following sections explore this in greater detail.

C. Performance analysis: No disturbance present

Since the formulation can be sensitive to the choice of hyper-parameters, a more detailed analysis was next carried out whereby different hyper-parameter choices (specifically the prediction horizon N , the regularisation weight λ_g and the initialisation/segment length T_{ini}) were used, with the tracking performance quantified for each as the sum of the absolute set-point deviations at each time step. This analysis was done for scenarios with and without the presence of disturbance, but we first present the undisturbed case with $d = 0$.

Fig. 5a shows the tracking performance of both segmented and unsegmented in the disturbance-free case under a range of prediction horizon from 5 to 100 seconds. In the unsegmented case, slight performance degradation can be seen with longer prediction horizons. This trend is not present in the segmented version. By varying the regularisation parameter λ_g , it can be seen in Fig. 5b that in both segmented and unsegmented cases, scenarios with smaller values of λ_g outperform those with larger values. This implies, as expected, that high regularisation penalties are not needed with perfect data and no disturbance. As λ_g increases, the performance of both formulations degrade as the relative weight on set-point deviation decreases. The choice of T_{ini} dictates the initialisation length of both cases and the segment length of the segmented version. Scenarios with $T_{ini} \leq 5$ perform well in both segmented and unsegmented cases as shown in Fig. 5c.

D. Performance analysis: Unmeasured disturbance present

With unmeasured disturbances now present, the same scenarios are simulated, with a greater performance gap emerging between the segmented and unsegmented formulations. In all cases, the set-point error is larger for disturbed scenarios than for undisturbed scenarios, as expected. Fig. 6a shows the performance of both formulations for a range of prediction horizons. The performance of the segmented version remains consistent across all horizon lengths. This is not true for the unsegmented case, for which significantly larger errors can be seen with horizons longer than 20 seconds.

Fig. 6b shows the performances achieved for different regularisation penalties. Unlike in the undisturbed case, a certain degree of regularisation is necessary, with very small values of λ_g corresponding to large set-point errors. Furthermore, the performance of the segmented version exceeds that of the unsegmented version across a wide range of values. A well chosen λ_g leads to improved performance of the segmented approach, as seen for $\lambda \approx 0.5$. In contrast, the performance of the unsegmented approach remains poor regardless of the choice of λ_g . We can also see in Fig. 6c that the choice of T_{ini} has a similar impact in both segmented and unsegmented cases. In both cases, the best performance is achieved for $T_{ini} = 5$ while for $T_{ini} < 4$ and $T_{ini} > 10$, the performance drops significantly.

The key outcome of these examples is that the performance of the segmented formulation exceeds that of the unsegmented version when unmeasured disturbances are present, without compromising performance when they are not.




Fig. 5: Comparison between segmented and unsegmented formulations for undisturbed two-mass spring-damper system varying (a) prediction horizons, (b) regularisation weight and (c) initialisation length

E. Performance analysis: Computation time

As the problem structure of the segmented version is now different to the unsegmented case, it is worth considering the computational time required for each. The structure of the segmented problem allows for the problem sparsity to be exploited, leading to a potential reduction in computation time. This was examined by finding the average computation time (over 100 runs) needed for the second-level quadratic optimisation, with various prediction horizons (with $\lambda_g = 0.5$ and $T_{ini} = 5$ for all cases). All scenarios were computed using the *quadprog* function, with the interior-point-convex algorithm, using the sparse setting for the internal linear solver in Matlab on a 2.9 GHz processor, with the results shown in Fig. 7. For $N = 100$, it can be seen that the unsegmented version computation time is almost double that of the segmented version.

A wider range of prediction horizons are plotted on a log-log scale plot in Fig. 8, in which we can see the slope of the computation time increase associated with the unsegmented version is approximately one, implying a linear time increase with horizon length. For the unsegmented case, the slope approaches three for longer horizons, implying a cubic time increase for increasing horizon length.




Fig. 6: Comparison between segmented and unsegmented formulations for two-mass spring-damper system with unmeasured disturbance varying (a) prediction horizon, (b) regularisation weight and (c) initialisation length




Fig. 7: Computational time for segmented and unsegmented formulations




Fig. 8: Log-log plot of computational time for segmented and unsegmented formulations

Analysis of the simple case study shows the benefits of splitting the time horizon into segments. A more realistic case study is next examined.

IV. APPLICATION TO BUILDING ENERGY MANAGEMENT

A. The building energy management challenge

An active area of research in recent times concerns the use of predictive control for building energy management. Modern energy systems require more flexibility to handle the combined influences of increased renewable generation and increased electrification of heating and transport. Making use of buildings as active, flexible components in such an energy landscape is a key requirement in global decarbonisation efforts [24]. Despite the pressing need for advanced control technologies, the underlying model complexity of a building and the wide variation in building designs has led to the model development process acting as a significant barrier to technology uptake [25]. Consequently, data-driven predictive control techniques have recently received attention for the application of building energy management [26].

The segmented formulation proposed in this paper is suited to this domain. Diurnal building usage and energy tariff patterns, along with the slow thermal dynamics of well-insulated buildings, make longer prediction horizons advantageous. Furthermore, many disturbances tend to impact the energy demand of a building. Measurements of these may not be available. External temperature, solar radiation and internal gains will influence the building's behaviour, potentially corrupting the ability of a data-driven algorithm to identify input/output behaviour from a given data-set. A simulated case study was carried out to investigate the performance of the segmented formulation in this setting, using state-of-the-art EnergyPlus [27] building simulation software and comparing the performance of the unsegmented and segmented formulations.

B. Building simulation environment

A popular technique for building thermal simulation is to represent the structure as a Resistance Capacitance (RC) network [28], particularly when knowledge of the physical composition of the building is available. The materials making up the walls, floors, ceilings, and windows are represented as configurations of resistances and capacitances whereby current flows through the circuit are analogous to heat flows through building components. Here, an EnergyPlus model of a six-room apartment was created based on standard building materials and thermal behaviour characteristics taken from the Tabula Webtool [29], and the underlying thermal model was extracted using the Building Resistance-Capacitance Modelling (BRCM) toolbox [30]. This resulting thermal model can be represented as a 102-state, linear, state-space system with six inputs (radiators in each room) and six outputs (the room temperatures). A schematic of the apartment layout can be seen in Fig. 9.

The building model is influenced by the ambient temperature and solar irradiance from different orientations. For this, weather data from a London-based weather station was




Fig. 9: Schematic of six-room apartment used for the building energy case study

obtained from the CEDA archive [31]. The occupancy profile used in the simulation was taken from the occupancy-integrated archetype approach of [32]. During occupied periods, a comfort set-point band between 20°C and 22°C was desired, while in unoccupied times, the temperatures were allowed to vary between 16°C and 26°C. The input in each room was constrained between 0 and the upper heat supply limit of the radiator in the room. The radiators were sized to emit a maximum of 100W per m² of floor area. The simulation ran with a 10-second sample time.

A separate data-driven predictive controller in each room with a sample time of 15 minutes was used to dictate the heat flow from the radiator to the room. The future set-point requirements were known to the controllers, as well as the current and previous room temperature and heat flow measurements. No measurements or forecasts of the weather were available to the controllers. A training period was carried out in which the radiators attempted to track a set-point varying between the upper and lower set-point bounds, using a PI controller. Note that this approach implies that the comfort set-point bounds should not be violated during the training period. The length of the training period depended on the formulation used (segmented or unsegmented) and the prediction horizon chosen for a particular scenario.

A set of scenarios were designed to compare the performance of the segmented and unsegmented formulations using different prediction horizons in this simulation environment. For these scenarios, we seek to minimise the deviation of the room temperatures outside the comfort bounds at a minimal cost. It was assumed that a heat pump supplies heat to the radiators with a Coefficient of Performance (COP) of 2.5, with electricity purchased via a time-varying tariff. For this, wholesale electricity price data was used with the Octopus Agile pricing tariff mechanism applied [33].

The formulation of Section III was modified slightly to incorporate an energy cost in the objectives. Once again, a prioritised framework was used, with the slack variables minimised first, followed by discomfort minimisation in a second




Fig. 10: Cost and comfort objectives for different prediction horizons using segmented and unsegmented formulations

optimisation, before finally minimising energy cost. The first two optimisation levels are formulated as in (25)–(37). The financial cost is considered in the third optimisation problem. The predicted electricity price for the period from $k + 1$ to $k + N$ is given as $C_{elec} = (c[k + 1], \dots, c[k + N]) \in \mathbb{R}^N$. The predicted electricity cost for heat pump consumption associated with the room over the prediction horizon was then included in the third-level objective as follows:

$$\min_{\substack{g_1, \dots, g_F, \varepsilon_y, \\ \varepsilon_{f_1}, \dots, \varepsilon_{f_F}}} \eta C_{elec} I_F \otimes U_{\beta_s} \begin{bmatrix} g_1 \\ \vdots \\ g_F \end{bmatrix} + \lambda_g \sum_{i=1}^F g_i^T g_i, \quad (38)$$

where η denotes the heat pump COP.

The constraints for this third-level problem are the same as for the second-level problem, with an additional constraint needed to enforce the optimal comfort performance, given as

$$\|\varepsilon_y\|_1 \leq \|\varepsilon_y^*\|_1, \quad (39)$$

where ε_y^* is the optimal ε_y computed in the second-level optimisation problem.

A decentralised architecture was used, in which each room has a separate controller and no communication between controllers occurs.

C. Performance analysis of data-driven controllers

Simulations were carried out to analyse the performance of the controllers for a three-week period using different prediction horizon lengths with the segmented and unsegmented formulations. Prediction horizons from 10 samples (2.5 hours) to 95 samples (1 day) are investigated. In all cases, $T_{ini} = 5$ and $\lambda_g = 1$ as these values were found to perform best for both segmented and unsegmented formulations. The results are summarised in Fig. 10, where the total heating cost for the apartment is plotted on the Y-axis and a discomfort metric is plotted on the X-axis. This discomfort metric is defined as the summation of absolute deviations from the comfort temperature set-point band, summed across each zone, scaled appropriately to achieve units of $^\circ\text{C}\cdot\text{hr}$.

Because the comfort objective has a higher priority than the financial objective, it should be expected that a longer prediction horizon would result in improved comfort, since pre-heating can be better exploited. This improvement in comfort may come at a financial cost, however. In the segmented formulation, this trend is clearly visible in Fig. 10, Table I

TABLE I: Discomfort metric for unsegmented and segmented formulations with different horizons (3-week simulation)

	$N = 10$	$N = 20$	$N = 40$	$N = 60$	$N = 95$
Unsegmented ($^\circ\text{C}\cdot\text{hr}$)	338	290	260	695	628
Segmented ($^\circ\text{C}\cdot\text{hr}$)	290	237	186	176	174

TABLE II: Heating cost for unsegmented and segmented formulations with different horizons (3-week simulation)

	$N = 10$	$N = 20$	$N = 40$	$N = 60$	$N = 95$
Unsegmented (£)	58.3	58.7	61.3	57.6	62.8
Segmented (£)	58.7	58.5	59.1	59.3	59.4

and Table II. Each increase in horizon length leads to a slight comfort improvement, with diminishing returns. The unsegmented formulation produces a similar (with slightly higher discomfort) performance for horizon lengths of 10 and 20 samples. Beyond this, the comfort performance varies without a clear pattern. The underlying values are summarised in Table I.

From a financial cost perspective, higher discomfort levels can imply an insufficient heat supply and thus lower heating cost. This can be seen in the unsegmented scenario with $N = 60$ in Fig. 10. Although the lowest financial cost is achieved, this is accompanied with a significant increase in discomfort, indicating that the controllers are not carrying out their primary objective and the low financial cost is a consequence of poor control performance (i.e., under-heating). The segmented formulation performance is far more consistent, with slight improvements in comfort accompanied by slight increases in heating cost as the horizon lengths increase. In the $N = 95$ case, the segmented approach reduced discomfort relative to the unsegmented approach by 72%, with a 5% cost reduction. Table II summarises the financial results.

A one-week window of the average apartment temperatures using $N = 95$ with the segmented and unsegmented formulations is plotted in Fig. 11 to illustrate the differing control performance. The electricity price for the same period is also shown. The unsegmented formulation overheats the apartment during the unoccupied periods compared to the segmented formulation. In both formulations, the controllers tend to pre-heat the apartment in advance of an electricity price spike. This behaviour is more effective in the segmented version, as the unsegmented case does not allow the temperatures to drop low enough to fully make use of the thermal energy inherently stored in the building.

As in the examples of Section III, the performance of the unsegmented strategy breaks down as the prediction horizon length increases, while the segmented formulation is more consistent in a wider range of operational strategies.

V. CONCLUSIONS

This paper proposes an extension to a data-driven predictive control formulation for linear systems with unmeasured




Fig. 11: Average room temperature in building with segmented and unsegmented formulations, plotted for one week of the simulation period ($N = 95$), and corresponding external weather conditions and electricity price profile for the period

disturbances. The proposed formulation modifies an existing data-enabled predictive control approach by segmenting the prediction horizon to decouple the number of implicit model parameters from the prediction horizon length. By doing so, the formulation performs better than the unsegmented formulation in the presence of unmeasured disturbance with longer trajectories.

The method was analysed here first using a set of case studies based on a two-mass-spring-damper system. Under various hyper-parameter choices, the segmented formulation outperformed the unsegmented formulation in terms of set-point tracking when disturbances were present, particularly with longer prediction horizons. Comparable performance was achieved in disturbance-free cases. The computation time associated with the proposed segmented formulation scales linearly with horizon length, improving on the time increase observed for the unsegmented formulation. For a scenario with a 100-sample horizon, computation time is approximately halved by using a segmented formulation.

The segmented formulation was applied to a building energy management case study to demonstrate the importance of these performance characteristics in a more realistic setting, using a state-of-the-art building simulation environment with realistic weather profiles acting as unmeasured disturbances. The segmented formulation performed more consistently with horizon length variation in terms of occupant comfort levels and energy consumption. For a scenario with a one-day-ahead prediction horizon, the segmented approach reduced discomfort relative to the unsegmented approach by 72%, with a 5% cost reduction.

Further work is needed to assess the impact of segmentation on the various extensions of the data-predictive controller that have been developed, such as formulations with robustness

guarantees and formulations for time-varying parameters and nonlinear systems. Additionally, methods for offset-free control in the presence of disturbance would also be beneficial to the data-driven context. Computational efficiency and hyper-parameter selection are also key aspects that require further focus to ensure algorithms that are tailored appropriately to a given context.

ACKNOWLEDGEMENT

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] 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, no. June 2019, p. 104211, 2020.
- [2] V. Bhattacharyya, A. F. Canosa, and B. HomChaudhuri, "Fast Data-Driven Model Predictive Control Strategy for Connected and Automated Vehicles," *ASME Letters in Dynamic Systems and Control*, vol. 1, no. 4, pp. 1–5, 2021.
- [3] J. Wang, S. Li, H. Chen, Y. Yuan, and Y. Huang, "Data-driven model predictive control for building climate control: Three case studies on different buildings," *Building and Environment*, vol. 160, no. March, p. 106204, 2019.
- [4] C. De Persis and P. Tesi, "Formulas for Data-Driven Control: Stabilization, Optimality, and Robustness," *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 909–924, 2020.
- [5] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. De Moor, "A note on persistency of excitation," *Systems and Control Letters*, vol. 54, no. 4, pp. 325–329, 2005.
- [6] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," *2019 18th European Control Conference, ECC 2019*, pp. 307–312, 2019.
- [7] F. Fiedler and S. Lucia, "On the relationship between data-enabled predictive control and subspace predictive control," 2020. [Online]. Available: <http://arxiv.org/abs/2011.13868>
- [8] F. Dörfler, J. Coulson, and I. Markovsky, "Bridging direct & indirect data-driven control formulations via regularizations and relaxations," 2021. [Online]. Available: <http://arxiv.org/abs/2101.01273>
- [9] H. J. Van Waarde, C. De Persis, M. K. Camlibel, and P. Tesi, "Willems' Fundamental Lemma for State-Space Systems and Its Extension to Multiple Datasets," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 602–607, 2020.
- [10] D. Alpagó, F. Dörfler, and J. Lygeros, "An Extended Kalman Filter for Data-Enabled Predictive Control," *IEEE Control Systems Letters*, vol. 4, no. 4, pp. 994–999, 2020.
- [11] J. Berberich, J. Kohler, M. A. Muller, and F. Allgower, "Data-Driven Model Predictive Control with Stability and Robustness Guarantees," *IEEE Transactions on Automatic Control*, vol. 66, no. 4, pp. 1702–1717, 2021.
- [12] J. Coulson, J. Lygeros, and F. Dörfler, "Distributionally Robust Chance Constrained Data-enabled Predictive Control," no. Id, pp. 1–14, 2020. [Online]. Available: <http://arxiv.org/abs/2006.01702>
- [13] L. Huang, J. Zhen, J. Lygeros, and F. Dörfler, "Robust Data-Enabled Predictive Control: Tractable Formulations and Performance Guarantees," 2021. [Online]. Available: <http://arxiv.org/abs/2105.07199>
- [14] Y. Lian and C. N. Jones, "From System Level Synthesis to Robust Closed-loop Data-enabled Predictive Control," 2021. [Online]. Available: <http://arxiv.org/abs/2102.06553>
- [15] Y. Lian, J. Shi, M. P. Koch, and C. N. Jones, "Adaptive Robust Data-driven Building Control via Bi-level Reformulation: an Experimental Result," pp. 1–12, 2021.
- [16] C. Verhoeck, H. S. Abbas, R. Tóth, and S. Haesaert, "Data-Driven Predictive Control for Linear Parameter-Varying Systems," 2021. [Online]. Available: <http://arxiv.org/abs/2103.16160>
- [17] Y. Lian, R. Wang, and C. N. Jones, "Koopman based data-driven predictive control," vol. XX, no. Xx, pp. 1–11, 2021. [Online]. Available: <http://arxiv.org/abs/2102.05122>
- [18] P. C. N. Verheijen, G. R. Gonçalves Da Silva, and M. Lazar, "Data-driven predictive control with estimated prediction matrices and integral action," 2021.

- [19] H. Yang and S. Li, "A data-driven predictive controller design based on reduced Hankel matrix," *2015 10th Asian Control Conference: Emerging Control Techniques for a Sustainable World, ASCC 2015*, no. 2014, pp. 1–7, 2015.
- [20] G. Prando, "Model Order Selection in System Identification: New and Old Techniques," Ph.D. dissertation, 2013, Università degli Studi di Padova.
- [21] J. C. Willems, "The Behavioral Approach to Open and Interconnected Systems," *IEEE Control Systems*, vol. 27, no. 6, pp. 46–99, 2007.
- [22] I. Markovsky and P. Rapisarda, "Data-driven simulation and control," *International Journal of Control*, vol. 81, no. 12, pp. 1946–1959, 2008.
- [23] E. C. Kerrigan and J. M. Maciejowski, "Designing model predictive controllers with prioritised constraints and objectives," *IEEE International Symposium on Computer Aided Control System Design*, pp. 33–38, 2002.
- [24] R. Lowes, J. Rosenow, M. Qadrdan, and J. Wu, "Hot stuff: Research and policy principles for heat decarbonisation through smart electrification," *Energy Research and Social Science*, vol. 70, no. March, p. 101735, 2020.
- [25] E. Atam and L. Helsen, "Control-Oriented Thermal Modeling of Multizone Buildings: Methods and Issues: Intelligent Control of a Building System," *IEEE Control Systems*, vol. 36, no. 3, pp. 86–111, 2016.
- [26] A. Kathirgamanathan, M. De Rosa, E. Mangina, and D. P. Finn, "Data-driven predictive control for unlocking building energy flexibility: A review," *Renewable and Sustainable Energy Reviews*, vol. 135, no. January 2020, p. 110120, 2021.
- [27] DOE, "EnergyPlus — EnergyPlus," p. 1, 2017. [Online]. Available: <https://www.energyplus.net/>
- [28] E. O'Dwyer, L. De Tommasi, K. Kouramas, M. Cychowski, and G. Lightbody, "Modelling and disturbance estimation for model predictive control in building heating systems," *Energy and Buildings*, vol. 130, pp. 532–545, oct 2016.
- [29] TABULA, "TABULA webtool." [Online]. Available: <http://webtool.building-typology.eu/#bm>
- [30] D. Sturzenegger, D. Gyalistras, V. Semeraro, M. Morari, and R. S. Smith, "BRCM Matlab Toolbox: Model generation for model predictive building control," *Proceedings of the American Control Conference*, pp. 1063–1069, 2014.
- [31] Met Office (2006), "MIDAS: UK hourly weather observation data. NCAS British Atmospheric Data Centre." [Online]. Available: <https://catalogue.ceda.ac.uk/uuid/916ac4bbc46f7685ae9a5e10451bae7c>
- [32] G. Buttitta, W. J. Turner, O. Neu, and D. P. Finn, "Development of occupancy-integrated archetypes: Use of data mining clustering techniques to embed occupant behaviour profiles in archetypes," *Energy and Buildings*, vol. 198, pp. 84–99, 2019.
- [33] Octopus Energy, "Frequently asked questions about Octopus Tracker — Octopus Energy," 2020. [Online]. Available: <https://octopus.energy/tracker-faqs/>