

Title: Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings

Authors: Bünning, F., Huber, B., Heer, P., Aboudonia, A. and Lygeros, J.

Please cite as: Bünning, F., Huber, B., Heer, P., Aboudonia, A. and Lygeros, J., 2020. Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings. *Energy and Buildings*, 211, p.109792.

Accepted article: <https://doi.org/10.1016/j.enbuild.2020.109792>

Preprint: <https://doi.org/10.31224/osf.io/axsjd>

Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings

Felix Büning^{1,2,*}, Benjamin Huber^{1,2,*}, Philipp Heer¹, Ahmed Aboudonia², John Lygeros²

¹Urban Energy Systems Laboratory, Swiss Federal Laboratories for Materials Science and Technology, Empa, Dübendorf, Switzerland

²Automatic Control Laboratory, Department of Electrical Engineering and Information Technology, ETH Zürich, Switzerland

**These authors contributed equally to this work*

Abstract

Model Predictive Control for room temperature control in buildings is an effective approach to energy management in buildings. However, the development and maintenance of physical models may be a bottleneck for widespread real life application. Data Predictive Control is an attempt to address this problem by learning the behaviour of the building from historical data and thus reducing the modelling effort. Here, we present an application of a Data Predictive Control approach, based on Random Forests with affine functions and convex optimization, to control the room temperature in a real life apartment. When compared to a conventional hysteresis controller, the applied approach saves 24.9 % of cooling energy while reducing the integral of comfort constraint violations by 72.0 % in a six-day experiment.

Keywords:

Data Predictive Control, Model Predictive Control, building energy, Random Forests, Machine Learning

1. Introduction

Predictive control has the potential to reduce the energy consumption for heating and cooling in the building domain. Model Predictive Control (MPC) has been demonstrated many times successfully for temperature control in buildings (see for example (Oldewurtel et al., 2010; Široký et al., 2011; Oldewurtel et al., 2012; Huang et al., 2015) or the articles reviewed in (Hameed Shaikh et al., 2014)). However, there is to the best of the authors' knowledge, no widespread industrial application of MPC in residential buildings. Some authors ((Žáčková et al., 2014; Sturzenegger et al., 2016; Jain et al., 2018)) argue that the effort to develop and maintain first principle models of buildings might in many cases be prohibitive for real life deployment of MPC to building energy management.

The increasing availability of monitoring data in the building domain gives rise to the possibility of using data-driven modelling approaches for buildings. Here, the challenge lies in finding methods that show good prediction performance for temperatures in buildings and at the same time allow online real-time optimization over the control inputs, which usually requires at least

a convex model. Regression trees approximate a function $y = f(x)$ by splitting x into partitions and approximating y with constant values within these partitions (more detail in Section 2). In contrast to conventional Artificial Neural Networks for example, regression trees in combination with affine models of the control inputs are convex in these inputs and can therefore be used for building thermal control.

Jain et al. (2018) use a combination of regression trees and linear regression for peak power reduction in buildings to optimally trade-off peak power reduction against thermal comfort, under the name of Data Predictive Control (DPC). The regression trees are built with the help of historical building measurement data to model the input-output relation of the current state of the system and future disturbances as inputs to the future state of the system as output. The regression tree is combined with an affine model that relates the control inputs to the future states of the system. This yields a model that can be used for receding horizon MPC with convex optimization in the same way as a state-space building model. In (Smarra et al., 2018b) the method is extended by using random forests, which are ensembles

of regression trees, instead of single regression trees to reduce the variance in prediction. In (Smarra et al., 2018a) the approach is further developed by replacing the affine control input - state relation by a switched state-space model. This allows the analysis of system properties such as stability.

All these studies use simulation for validation, where all disturbances are perfectly known and undesired disturbances and uncertainties can be excluded. There are a limited amount of studies ((Costanzo et al., 2016; Macarulla et al., 2017; Ferreira et al., 2012; Finck et al., 2019)) that use data-driven modelling techniques for control in the building domain in real life application case studies. However, to the best of the authors' knowledge, there exists no real life demonstration of DPC for room temperature control in buildings to minimize energy consumption, as shown for MPC in (Oldewurtel et al., 2010; Široký et al., 2011; Oldewurtel et al., 2012; Huang et al., 2015), where control inputs are minimized while keeping room temperatures within comfort constraints.

Here, we apply a variation of the DPC approach introduced in (Smarra et al., 2018b) to a real life residential apartment. The apartment features two identical bedrooms, one of which is DPC-controlled and the other one controlled with a standard hysteresis controller. Our results in a six-day experiment show that the room controlled with DPC consumes significantly less cooling energy than the conventionally controlled room, as well as showing fewer comfort constraint violations.

The remaining article is structured as follows. In Section 2, the DPC approach by Smarra et al. (2018b) is explained and adapted to our setting. In Section 3, the residential apartment for the experiment is introduced and the particular DPC controller is specified. In Section 4, the performance of the DPC controller is compared to the hysteresis controller in terms of constraint violations and energy usage. Conclusions and future directions are provided in Section 5.

2. Methodology

2.1. Splitting of historical data

We assume given a historical data set (X, Y) of measurements of a building. $Y = \{x(k+1), \dots, x(k+N)\}$ consists of measurements of room temperatures x . $X = \{u(k), \dots, u(k+N-1), d(k-\delta), \dots, d(k+N-1), x(k-\delta), \dots, x(k)\}$ consists of measurements of data that potentially has an influence on Y . These are control inputs u (e.g. cooling energy), disturbances d (e.g. ambient temperature and solar radiation forecasts, but also time

of day variables) and autoregressive terms of the room temperature x . N denotes the prediction horizon for the predictive controller, δ denotes the number of autoregressive terms.

We further split X in two separate sets $X_u = \{u(k), \dots, u(k+N-1)\}$, which includes all measured control inputs (all variables that can be influenced by the controller), and $X_d = \{d(k-\delta), \dots, d(k+N-1), x(k-\delta), \dots, x(k)\}$, which includes all measured disturbances and autoregressive terms of x (all variables that cannot be influenced by the controller).

2.2. Training of random forests

We train N random forests to approximate the functions $X_d \rightarrow Y_j$, in which $j \in 1, \dots, N$ denotes the columns of Y (and the steps in the prediction horizon). Random forests are ensembles of regression trees - a form of decision tree to approximate continuous functions. A decision tree is built by splitting the input data X_d^p into partitions p , in each of which the output data Y_j^p is approximated with a constant value \bar{y}_j^p (which is the mean of all elements of Y_j^p). The splitting variable and splitting point is obtained with the help of a greedy heuristic (see (Smarra et al., 2018b)). A random forest of T trees is built by bootstrapping T subsets from (X_d, Y_j) and building a tree based on each subset. The function evaluation of the forest for one element (row) x_d of X_d is then done by evaluating each tree, giving rise to T different $\bar{y}_j(x_d)$, and taking the average of these, giving rise to $\hat{y}_j(x_d)$. T is independent of N and needs to be chosen heuristically. For more details on the training process please refer to the original source (Smarra et al., 2018b) and to general literature on random forests (e.g. (Louppe, 2014)).

2.3. Fitting of affine functions for control inputs

Each tree has a number of leaves l_i , each based on a partition X_d^p of X_d to approximate Y_j^p with a constant value \bar{y}_j^p (as discussed before). For each set X_d^p there exists a corresponding set of measured control inputs X_u^p that has not been used in the model until here. (Smarra et al., 2018b) fit an affine model to approximate Y_j^p with X_u^p (to relate the response of the room temperature to the control inputs) with least squares in each leaf l_i . For the prediction of the future temperature of a single room $x(k+j)$ (which corresponds to y_j) in one leaf, this gives

$$x(k+j) = \beta_{j,0} + \sum_{n=1}^j \beta_{j,n} u(k+n-1) + e \quad (1)$$

in which $\beta_{ji,0}$ and $\beta_{ji,n}$ denote the fitted coefficients and e the model error. The temperature is therefore an affine function of all control inputs from the time instant k , when the forecast is made, to the forecasted time instant $k+j$. As each leaf l_i is based on a different partition of X_d , $\beta_{ji,0}$ and $\beta_{ji,n}$ are also different for each leaf.

Preliminary results suggested that the proposed approach leads to weak prediction performance in the case of $j > 3$. This can be explained by the high dimensionality of the fitting process for bigger prediction steps in combination with a limited amount of samples of pairs (x_u, y_j) in each leaf l_i . We therefore model the predicted temperature in the j^{th} step as

$$x(k+j) = \beta_{ji,0} + \beta_{ji,1} \sum_{n=1}^j u(k+n-1) + e, \quad (2)$$

in which the dimension of β_{ji} is reduced to 2 for all j . This approach is less realistic than eq. (1), as $u(k)$ will have the same effect on $x(k+j)$ as $u(k+j-1)$ has, whereas one would expect the effect of earlier time steps to be weaker. However, preliminary experiments suggested that the advantage of better model fits outweighs the disadvantage of a less realistic model. We can rewrite eq. (2) as

$$x(k+j) = \beta_{ji}[1, u(k), \dots, u(k+j-1)]^T + e, \quad (3)$$

with β_{ji} being a row vector of dimension $j+1$.¹ In the case where random forests are used instead of regression trees, the average of all relevant β_{ji} is taken, giving rise to β_j .

2.4. MPC-like use of DPC model

The learned β_j can be used to perform conventional receding horizon MPC. The corresponding convex optimization scheme follows

$$\begin{aligned} & \underset{u, \epsilon}{\text{minimize}} && \sum_{j=0}^{N-1} u_{k+j}^T R u_{k+j} + \lambda \epsilon_{k+j+1} \\ & \text{subject to} && x_{k+j} = \beta_j [1, u(k), \dots, u(k+j-1)]^T \\ & && x_{\min} - \epsilon_{k+j} \leq x_{k+j} \leq x_{\max} + \epsilon_{k+j} \\ & && u \in U \\ & && \epsilon \geq 0 \\ & && j = 1, \dots, N \end{aligned} \quad (4)$$

¹The dimensions of β_{ji} and $u(k), \dots, u(k+j-1)$ of course change if multiple rooms are considered because in that case also $x(k+j)$ becomes a vector.



Figure 1: NEST building with controlled residential unit UMAR marked in white © Zooney Braun, Stuttgart

in which R and λ denote costs for control inputs and constraint violations, ϵ is a slack variable for the comfort constraint, x_{\min} and x_{\max} denote the comfort constraints and U defines the set of allowed control inputs. The algorithm for the online phase of the DPC controller is shown in 1.

Algorithm 1

- 1: **procedure** DPC CONTROLLER (ONLINE)
 - 2: **for** k in runtime **do**
 - 3: **for** j in $1, \dots, N$ **do**
 - 4: find all leaves l_i in forest j based on current state and predicted disturbances $x_d(k)$
 - 5: construct β_j by averaging over all β_{ji} in all found l_i
 - 6: solve optimization problem (4)
 - 7: apply $u(k)$ to the system
-

3. Case study

3.1. General description

The controlled apartment is a residential unit in the NEST demonstrator building at Empa in Switzerland (Richner et al., 2018) and is marked in Figure 1. A schematic of the layout is shown in Figure 2. On the left and right side, the apartment features two nearly identical bedrooms that both have a floor area of 17.6 m² and a large window front that faces east-southeast. The right room is controlled by the DPC controller presented in this study, the left room is controlled by a hysteresis controller. The whole apartment is equipped with combined heating and cooling ceiling panels that are operated with water. In the cooling case (present in




Figure 2: Layout of the controlled building unit

this experiment) the water has a constant supply temperature of 18 °C. The amount of transferred heat is controlled by an individual open/closed valve in each room. The hysteresis-controlled room was occupied during the experiment, the DPC-controlled room was unoccupied. (This was switched in the validation experiments in Appendix A.)

3.2. Controller set-up

The applied DPC algorithm uses as inputs autoregressive terms and the forecast ambient temperature and solar radiation from the Federal Office of Meteorology and Climatology (MeteoSwiss), the time of the day encoded in a sinus and cosinus function and autoregressive terms of the room temperature and neighbouring rooms with $\delta = 16$.² We use a forest of $T=200$ trees with a minimum amount of 200 historical samples in each leaf. The forest is trained on 10 months of historical measurement data, sampled in a 10-minute interval. Instead of absolute room temperatures x , we use the approach to model and predict temperature changes Δx . This does not fundamentally change the algorithm but led to better prediction performance in preliminary experiments. Moreover, we use sample weighting in the historical measurement data set to train the forests in order to adjust the importance of individual samples based on their location in the input space. The prediction horizon is $N=6$ with a sampling time of one hour, the cost per squared unit of cooling is $R=1$ and the cost per unit of comfort constraint violation is $\lambda=100$. The comfort constraints are 22-23 °C during the night (8.00 pm to 8.00 am) and 22-25 °C during the day (8.00 am to 8.00 pm).³

²This number might seem excessive, however, additional inputs do not affect the fitting process of regression trees negatively, as unimportant inputs do not get chosen as splitting/partitioning variables.

³This is done to allow a better assessment of the predictive capabilities of the controller and to simulate a situation where the occupants are out-of house (e.g. at work) during the day.

Measurements of the autoregressive terms and weather forecasts are queried via a Python-SQL-client from a database, the DPC algorithm is implemented in Python 3 using Scikit-learn (Pedregosa et al., 2011), communication with the actuators in the building is established via a Python-OPC UA client.

The hysteresis controller works in the following way: It switches on when the upper DPC constraint is reached, it switches off when the temperature is 1 °C below the upper constraint. We do not force the hysteresis controller to cool down until the lower DPC constraint as this would result in a much higher energy consumption for the hysteresis controller and non-comparable average room temperatures. The connection to the database and actuators is established the same way as for the DPC controller.

4. Results and Discussion

Figure 3 shows the results for a six-day experiment from 2019-06-20 2.00 pm until 2019-06-26 2.00 pm in the apartment with a cooling demand. The bottom two graphs show the ambient conditions measured at the roof-top of the building. It can be seen that ambient temperatures between 13.7 and 39.5 °C and solar radiations between 0 and 1310 W/m² (which is unrealistic due to the use of a non-calibrated sensor, but irrelevant as the used approach is data-driven) are present during the time of the experiment. The conditions allow a good estimation of the controller performance as the experiment covers the range from mild cooling on 06-22 with a maximum daily temperature of 22.7 °C and heavy cooling (to the cooling system's limit) at the end of the experiment with a maximum daily temperature of 39.5 °C.

The top graph of Figure 3 shows the trajectory of the room temperatures (with a quantisation of 0.15 °C) for the DPC-controlled room in blue, and the hysteresis-controlled room in dashed orange. The DPC comfort constraints are plotted in grey. First, it should be noted that the rooms indeed show very similar dynamics. For example, in the middle of day 06-24 (marked by the second grey period in the figure), where cooling is switched off in both rooms and the constraint jumps to 25 °C, the temperatures follow the same upwards trend. This observation is confirmed by the validation experiments in Appendix A.

It can be seen that the DPC controller shows much less constraint violations than the hysteresis controller. More specifically, the integral of the constraint violations is 72.0 % smaller in the DPC case. This is to be expected as the hysteresis controller has no predictive

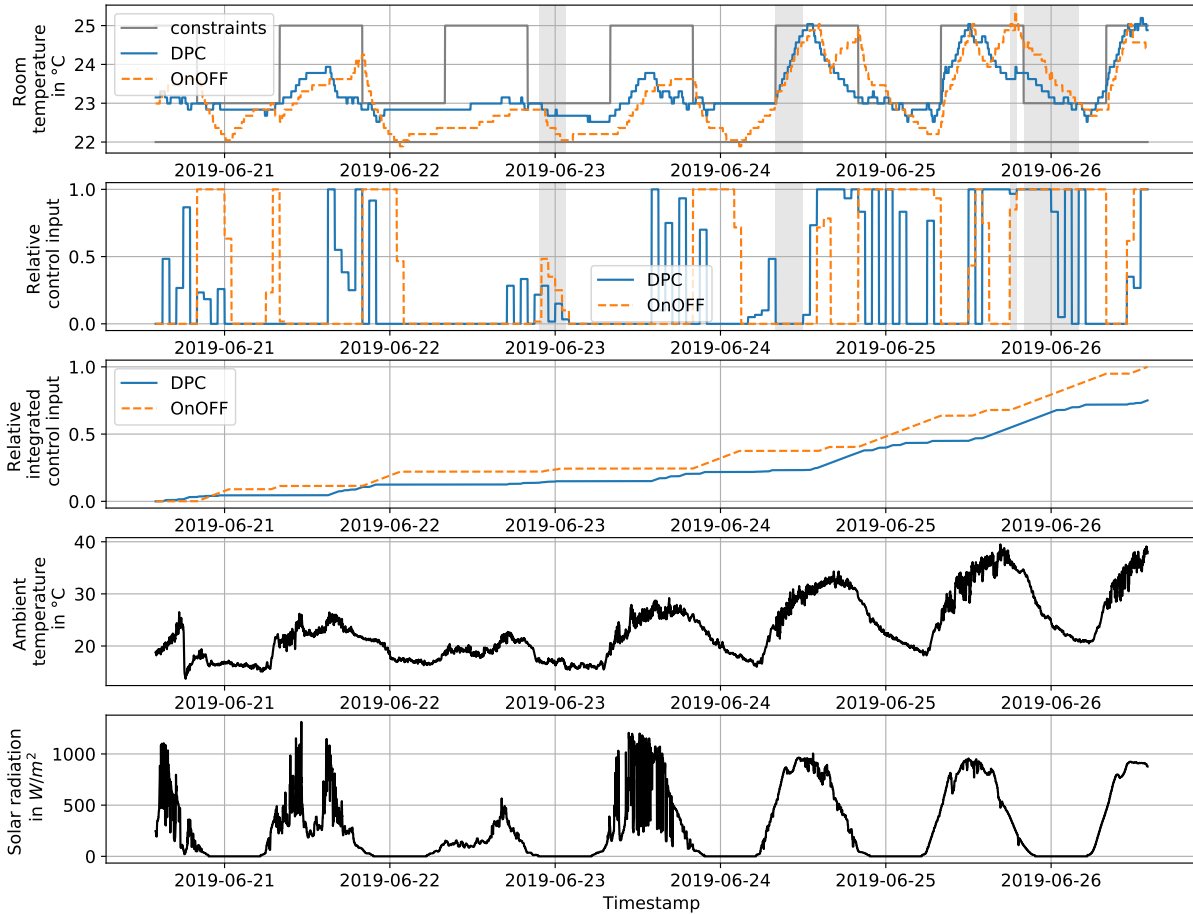


Figure 3: DPC cooling experiment results

capabilities and only starts cooling once it hits a constraint. But it also shows that the prediction accuracy of the DPC-model is accurate enough to start cooling early during the day in order to not violate the upper comfort constraint during the night: The temperature trajectory in the DPC controlled room exactly meets the lowered temperature constraint at 8.00 pm on all days except 06-25. On the evening of 06-25 the constraint is not met because the cooling system operated at it's maximum capacity with a maximum ambient temperature of 39.5 °C on that particular day and an ambient temperature of still 34.6 °C at 8.00 pm (fourth grey period in the figure). This issue could possibly addressed with a longer prediction horizon to allow pre-cooling. During night time, the temperature in the DPC-controlled room stays

close to the upper constraint most of the time. This is desirable for a predictive controller in the cooling case because it minimizes the control input.

The second graph of Figure 3 shows the control input (relative between minimum cooling capacity of 0 and maximum of -0.6 kWh/h) during the experiment. It can be seen that the hysteresis controller uses much more cooling energy than the DPC-controller. This is partly due to always cooling until the room temperature is one degree below the upper comfort constraint. However, the average temperature in the DPC-controlled room is only 0.1 °C above the hysteresis-controlled room, which is in the range of the quantisation error of the temperature sensor. This shows that the cooling energy savings are not due to a higher average room temperature. Over

the course of the whole experiment, the DPC-controller saved 24.9 % cooling energy compared to the hysteresis controlled room, as can be seen in the middle graph that shows the relative integral of the control input.

Both top graphs show two more areas with grey background. In the first period between 06-22 and 06-23 the controllers lost connection to the actuators multiple times, which caused the valves in the cooling system to open and close multiple times - thus the part load in the hysteresis controlled room. In the third period on 06-26, the windows were opened in both rooms for a period of 10 minutes. This automatically switches off cooling and together with an ambient temperature higher than the room temperature resulted in a slight increase in room temperature.

5. Conclusion

Data Predictive Control in buildings constitutes an attempt to lower the effort for model development and maintenance for predictive thermal building control. Despite good simulation results there is no practical experience in real life cases for room temperature control between comfort constraints and minimization of cooling or heating input with such methods so far. Here, we apply a DPC approach based on Random Forests with affine functions and convex optimization to temperature control in a real life residential apartment and compare the controller to a conventional hysteresis controller. The temperature in the room controlled by DPC shows 72.0 % less integrated comfort constraint violations than in the room controlled with hysteresis-control. Moreover, the DPC controller saves 24.9 % of cooling energy. The experiment constitutes a promising first real life application of the algorithm. A comparison to other control approaches and the investigation of different prediction horizons and sampling rates as well as experiments regarding heating are part of ongoing and future research.

Acknowledgements

We would like to thank Kristina Orehounig, Viktor Dorer and Roy S. Smith for their valuable help and support. We are also grateful to Annika Eichler, Benjamin Flamm, Andrea Iannelli, Mohammad Khosravi, Jeffrey Lungthok, Anil Parsi, Nicolas Schmid, Bratislav Svetozarevic, Christoph Waibel, Joseph Warrington and Suli Zou for fruitful discussions. Moreover, we would like to thank Reto Fricker, Ralf Knechtle and Sascha

Stoller for sharing their expertise and help with the implementation, and Spyros Georgakos for showing patience and engagement with our experiments.

This research project is financially supported by the Swiss Innovation Agency Innosuisse and is part of the Swiss Competence Center for Energy Research SCCER FEEB&D.

Conflict of interest

None.

Appendix A. Validation experiments

We have conducted two experiments from 2019-05-31 0.00 am to 2019-06-06 0.00 am (validation experiment 1) and 2019-06-15 0.00 pm to 2019-06-20 0.00 pm (validation experiment 2) before the experiment described in the main body of this article. The results, shown in Figure A.4 and Figure A.5 respectively, were not used as the main result for this study because both DPC and hysteresis controller were not yet operating totally correct: The DPC controller sometimes saw the lowered upper constraint at 8.00 pm too late because the optimization was performed at 1.59 pm instead of 2.00 pm, resulting in increased constraint violations. The hysteresis controller sometimes started cooling before hitting the upper constraint.

However, these results are suitable to confirm 1. that the thermal dynamics of both rooms are indeed very similar: validation experiment 1 and 2 were performed with the DPC controller operating in the left room and the hysteresis controller operating in the right room in contrast to the main experiment; and 2. that the order of magnitude of energy savings is realistic and the presented main result is not a best-case example. In validation experiment 1 the energy saving amounted to 31.9 % and in validation experiment 2 to 29.7 %, both compared to the respective hysteresis controller. The difference in average room temperature was less than 0.05 °C in both experiments. The integrated constraint violation was reduced by 74.8 % and 79.7 % respectively.

References

- Costanzo, G., Iacovella, S., Ruelens, F., Leurs, T., Claessens, B., 2016. Experimental analysis of data-driven control for a building heating system. *Sustainable Energy, Grids and Networks* 6, 81–90. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2352467716000138>, doi:10.1016/j.segan.2016.02.002.

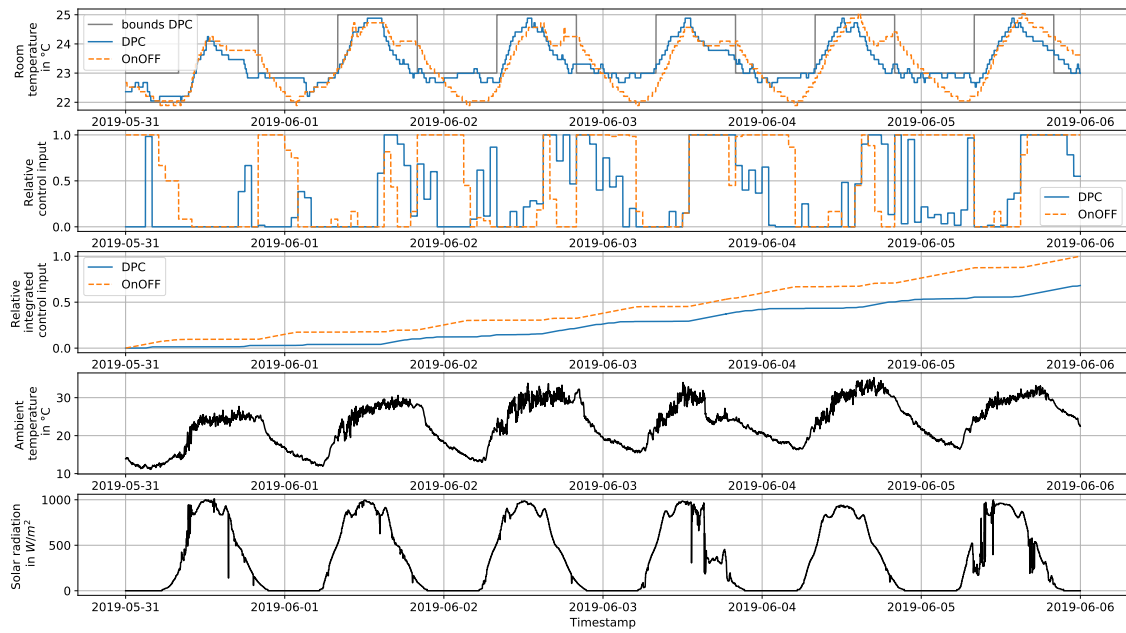


Figure A.4: DPC cooling validation experiment 1

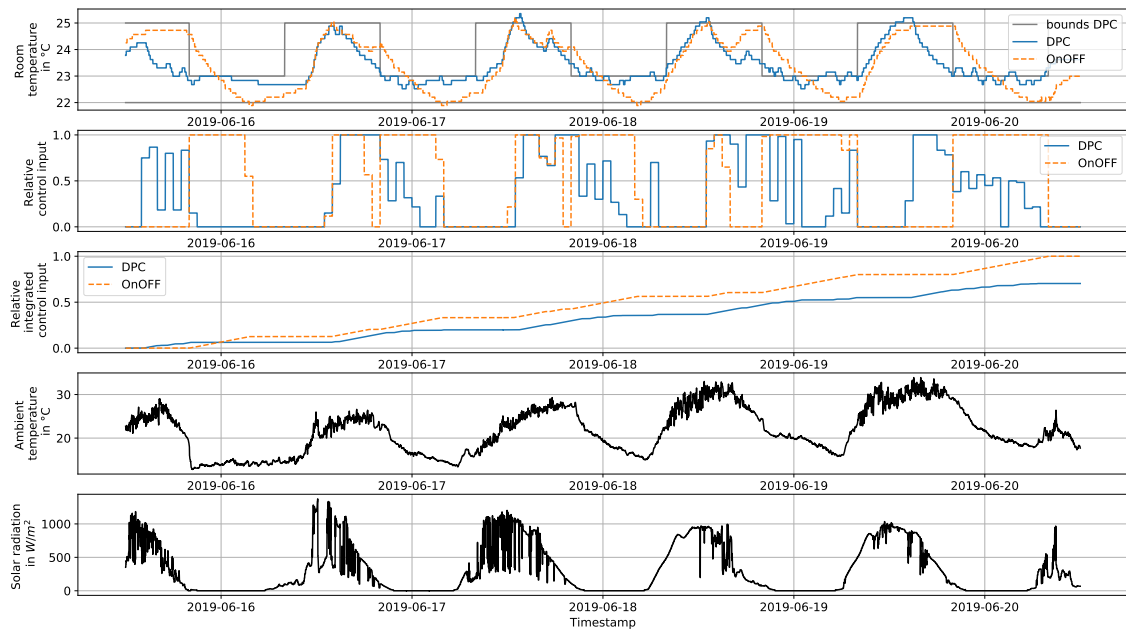


Figure A.5: DPC cooling validation experiment 2

- Ferreira, P., Ruano, A., Silva, S., Conceição, E., 2012. Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and Buildings* 55, 238–251. URL: <https://linkinghub.elsevier.com/retrieve/pii/S037877881200388X>, doi:10.1016/j.enbuild.2012.08.002.
- Finck, C., Li, R., Zeiler, W., 2019. Economic model predictive control for demand flexibility of a residential building. *Energy* 176, 365–379. doi:10.1016/j.energy.2019.03.171.
- Hameed Shaikh, P., Bin Mohd Nor, N., Nallagownden, P., Elamvazuthi, I., Ibrahim, T., 2014. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews* 34, 409–429. URL: <http://dx.doi.org/10.1016/j.rser.2014.03.027>.
- Huang, H., Chen, L., Hu, E., 2015. A new model predictive control scheme for energy and cost savings in commercial buildings: An airport terminal building case study. *Building and Environment* 89, 203–216. URL: <http://dx.doi.org/10.1016/j.buildenv.2015.01.037>.
- Jain, A., Smarra, F., Behl, M., Mangharam, R., 2018. Data-Driven Model Predictive Control with Regression Trees—An Application to Building Energy Management. *ACM Transactions on Cyber-Physical Systems* 2, 1–21. URL: <http://dl.acm.org/citation.cfm?doid=3174275.3127023>, doi:10.1145/3127023.
- Loupe, G., 2014. Understanding Random Forests: From Theory to Practice URL: <http://arxiv.org/abs/1407.7502>, arXiv:1407.7502.
- Macarulla, M., Casals, M., Forcada, N., Gangoles, M., 2017. Implementation of predictive control in a commercial building energy management system using neural networks. *Energy and Buildings* 151, 511–519. URL: <https://www.sciencedirect.com/science/article/pii/S0378778817300907>, doi:10.1016/J.ENBUILD.2017.06.027.
- Oldewurtel, F., Parisio, A., Jones, C., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Wirth, K., 2010. Energy Efficient Building Climate Control using Stochastic Model Predictive Control and Weather Predictions, in: *Proceedings of the 2010 American Control Conference*, Ieee Service Center, 445 Hoes Lane, Po Box 1331, Piscataway, Nj 08855-1331 Usa. pp. 5100–5105. URL: <https://infoscience.epfl.ch/record/169733>.
- Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., Morari, M., 2012. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings* 45, 15–27. URL: <https://www.sciencedirect.com/science/article/pii/S0378778811004105>, doi:10.1016/J.ENBUILD.2011.09.022.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, É., 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12, 2825–2830. URL: <http://www.jmlr.org/papers/v12/pedregosa11a.html>.
- Richner, P., Heer, P., Largo, R., Marchesi, E., Zimmermann, M., Zimmermann, M., 2018. NEST – una plataforma para acelerar la innovación en edificios. *Informes de la Construcción* 69, 222. URL: <http://informesdelaconstruccion.revistas.csic.es/index.php/informesdelaconstruccion/article/view/5879>, doi:10.3989/id.55380.
- Široký, J., Oldewurtel, F., Cigler, J., Prívará, S., 2011. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy* 88, 3079–3087. URL: <https://www.sciencedirect.com/science/article/pii/S0306261911001668>, doi:10.1016/J.APENERGY.2011.03.009.
- Smarra, F., Jain, A., Mangharam, R., D’Innocenzo, A., 2018a. Data-driven Switched Affine Modeling for Model Predictive Control. *IFAC-PapersOnLine* 51, 199–204. URL: <https://www.sciencedirect.com/science/article/pii/S2405896318311509>, doi:10.1016/J.IFACOL.2018.08.034.
- Smarra, F., Jain, A., de Rubeis, T., Ambrosini, D., D’Innocenzo, A., Mangharam, R., 2018b. Data-driven model predictive control using random forests for building energy optimization and climate control. *Applied Energy* 226, 1252–1272. doi:10.1016/j.apenergy.2018.02.126.
- Sturzenegger, D., Gyalistras, D., Morari, M., Smith, R.S., 2016. Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost–Benefit Analysis. *IEEE Transactions on Control Systems Technology* 24, 1–12. URL: <http://ieeexplore.ieee.org/document/7087366/>, doi:10.1109/TCST.2015.2415411.
- Žáčková, E., Vána, Z., Cigler, J., 2014. Towards the real-life implementation of MPC for an office building: Identification issues. *Applied Energy* 135, 53–62. URL: <https://www.sciencedirect.com/science/article/pii/S0306261914008071>, doi:10.1016/j.apenergy.2014.08.004.