Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Hardware-in-the-loop test of learning-based controllers for grid-supportive building heating operation

Lilli Frison * **Sweetin Paul** * **Torsten Koller** ** **David Fischer** *
Gianluca Frison *** **Joschka Boedecker** ** **Peter Engelmann** *

* *Fraunhofer Institute for Solar Energy Systems ISE, Freiburg, Germany (e-mail: lilli.frison@ise.fraunhofer.de).*
** *Neurorobotics Lab, Dept. of Computer Science, University of Freiburg, Germany*
*** *Dept. of Microsystems Engineering, University of Freiburg, Germany*

Abstract: While MPC is the state-of-the-art approach for building heating control with proven cost savings and improvement in energy flexibility, in practice, buildings are operated by simple rules-based controllers which are not able to accomplish an energy efficient and flexible operation. This paper explores the suitability of deep neural networks for approximating optimal economic MPC strategies for this task. In particular, we develop a convolutional neural network controller and test it in a closed-loop simulation against MPC and an improved predictive rule-based controller. The learned controller is easy to implement and fast to process on standard building control hardware. The feasibility, performance and robustness of the learned controller is validated in a realistic hardware-in-the-loop test setup for the demand-responsive operation of a heat pump combined with a storage tank.

Keywords: Learning-based Control; Neural Network Control; Economic MPC; Energy Storage Operation and Planning; Building Automation; Smart Grids

1. INTRODUCTION

In recent years, many energy efficient model predictive control (MPC) approaches have been reported for control of building heating systems, cf. Oldewurtel et al. (2012). Cost saving results of MPC against rule-based controllers are shown in Fischer et al. (2017). The ability to optimize different objective criteria while satisfying system constraints becomes necessary when a demand responsive building operation is required, cf. Halvgaard et al. (2012); Péan et al. (2018). Particularly useful for demand-side power management in buildings are heat pumps and other power-to-heat systems that transform electricity, and possibly additional "cheap" environmental energy, into heat. To enable load scheduling, thermal storage such as water tanks or the building mass are used. Adapting the electricity consumption for heating to the requirements of the power grid can be incentivised by time-variable electricity prices with high cost during times of peak demand or low renewable energy production.

The basic principle of MPC is to solve an optimization problem at each sampling interval using new measurements and forecasts of the grid signal, building load and environmental conditions. Despite intensive research efforts, the practical adoption of MPC within building controllers is still in its early stages due to the cost and time related to creating the building and heating system model and solving the optimization problem on site Cigler et al. (2013); Zong et al. (2017). Setting up an MPC controller on site is challenging because it requires expert knowledge of the building operators for design and tuning. The nonlinear and hybrid nature of the building energy system increases the solution complexity, cf., e.g., Bürger et al. (2018); Fischer et al. (2017). Most technicians and commission engineers are not well-versed on numerical optimization techniques to set up complex control systems. Unlike the industrial applications of MPC, most buildings are not operated with on-site engineers who are constantly monitoring and supervising the functioning of the employed control system.

Researchers have identified two main approaches to get MPC-like performance without its drawbacks by using machine learning tools: by improving modeling and system identification and by reducing the computational requirements to solve the MPC optimization problem. Jain et al. (2017) as well as Balint and Kazmi (2019) use neural networks and regression trees to estimate the internal building model. Afram et al. (2017) apply feed-forward neural networks for model identification of the heating and building system. In these works, the control actions are computed by solving an optimization problem. Learning algorithms for imitating MPC control have been explored by Drgoňa et al. (2018). In that paper, various machine learning methods such as regression models, support vector machines and a feed-forward time-delay neural network are used to

⋆ The authors acknowledge the financial support by the German Federal Ministry of Economics and Energy in the project "flexControl" (project code: 03ET1359A).

directly predict the control actions. These research efforts show that the potential performance gain compared to traditional rule-based controllers of the learned controllers based on deep neural networks has been recognized in theoretical studies. However, none of the studies consider additionally the practical feasibility and robustness of the learned controllers by implementation in the real plant.

In this work, we develop a controller based on deep neural networks that is able to infer rules from optimal MPC strategies for a grid-supportive heat pump operation with thermal storage. It is easy to implement and fast to process on building controller hardware without the need to solve an optimization problem. The learned controller is compared in closed loop simulation to MPC, a non-predictive and a predictive rule-based controller. In terms of grid flexibility, the simulation results show that the learned controllers are approximately 10% better than the predictive rule-based controller and 10% worse than MPC. Also the operation cost show a similar performance however the exact extend depends on the considered price signal. Our second main contribution is the implementation of the learned controller into a hardware-in-the-loop (HIL) test bed consisting of a heat pump and storage system to examine its suitability for closed-loop control in real buildings.

2. PROBLEM FORMULATION

2.1 HIL test setup

Hardware-in-the-Loop (HIL) testing combines real plants and system components with emulated modules that can contain numerical simulations. The goal is to evaluate the complete interaction of thermal, hydraulic, electrical and communication interfaces under realistic settings emulated in a lab environment. Compared to a purely numerical simulation, HIL provides a more realistic behavior of the plant and avoids modeling inaccuracies. The drawback is that experiments are more costly and time-consuming to conduct. The employed test bed is shown in Figure 1. The software and hardware communication infrastructure as well as the set-up of the MPC experiment were previously described in Frison et al. (2019).

Fig. 1. Testbed with heat pump (right), storage tank, circulation pump (middle) and heat meter (left)

Figure 2 depicts the corresponding heating system scheme where the heat pump is connected in parallel with a storage tank. The water for space heating is injected into the top of the storage tank while the return water for the heat pump is taken from the bottom. The ground water source and the building are emulated in the lab. The heat pump employed in the test bed allows to modulate the thermal power output in the range of approximately 7 to 25 kW. The water tank has a volume of 1m^3, which gives a maximum storage capacity of approximately 23kWh (assuming a low-temperature heat supply of 35°C and a large storage over-heating temperature of 20K). How long the system can maintain the energy storage capacity during a demand response event and provide up- or down-regulation to the grid depends on the heating load profile.

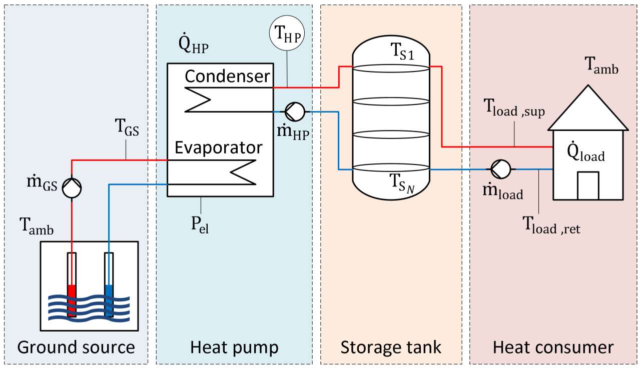

Fig. 2. Simplified test bed model

The low-level heat pump controller is in charge of keeping the heat pump supply temperature to the storage tank close to the control variable T_{HP} by selecting an appropriate compressor frequency. The high-level MPC controller determines T_{HP} as well as the heat pump operation times in order to facilitate a grid-supportive operation.

2.2 Model description

In the following, the heat pump and storage tank models suitable for optimization are described.

Heat pump model The heating power

$$\dot{Q}_{\text{HP}} = \dot{m}_{\text{HP}} c_p (T_{\text{HP}} - T_{S_N})$$

is calculated from the difference between return temperature to the heat pump T_{S_N} and the supply temperature from the heat pump to the tank T_{HP}. According to the test bed settings, the mass flow is set constant to $\dot{m}_{\text{HP}} = 0.583$. $c_p \approx 4181\,\text{J/kg K}$ is the specific heat capacity of water. The used electricity P_{el} is expressed as the ratio of heat produced $\dot{Q}_{\text{HP}}$ and the heat pump's coefficient of performance (COP). The COP is expressed by a standard 2nd-order polynomial fit with manufacturer data with non-linear behavior over compressor speed f_{comp}, sink and source temperatures T_{HP} and T_{GS}. Experience with heat pump tests in the lab has shown that it reacts very slowly to changes. Consequently, the desired heat pump temperature is reached only after some minutes while the above model assumes immediate reaction. The heat pump dynamics can be satisfactorily modeled by a simple first-order system with time constant of 5 minutes.

Storage model The water tank is modeled by a stratified storage model based on nodal energy and mass flow balances for the temperature evolution of each layer (Eicker, 2003). According to the sensors in the test bed, the number of layers equals 4. The temperature evolution of each layer i of mass m_i, transmission losses to the exterior $\dot{Q}_{\mathrm{loss},i}$, heat conduction between layers $\dot{Q}_{\mathrm{cond},i}$ and mixing introduced during charge ($\dot{Q}_{\mathrm{HP},i} > 0$) and discharge ($\dot{Q}_{\mathrm{load},i} > 0$) cycles is formulated as follows:

$$m_i c_p \frac{dT_{S_i}}{dt} = -\dot{Q}_{\mathrm{loss},i} + \dot{Q}_{\mathrm{cond},i} + \dot{Q}_{\mathrm{HP},i} - \dot{Q}_{\mathrm{load},i} + \delta_i^+ \dot{m}_i c_p (T_{S_{i-1}} - T_{S_i}) - \delta_i^- \dot{m}_{i+1} c_p (T_{S_i} - T_{S_{i+1}}).$$

The effective mass flow $\dot{m}_i = \dot{m}_{\mathrm{HP}} - \dot{m}_{\mathrm{load}}$ is positive if energy enters from the upper layer $i-1$, i.e. $\dot{m}_{\mathrm{HP}} > \dot{m}_{\mathrm{load}}$, in which case the parameter $\delta_i^+ = 1$ (otherwise $\delta_i^+ = 0$). A negative effective mass flow from the lower layer $i+1$, i.e. dominance of the load mass flow and thus cooling of layer i, is similarly taken into account by parameter δ_i^-.

Building heating demand For a given heat demand profile $\dot{Q}_{\mathrm{load}}$ the mass flow at the storage is given by

$$\dot{m}_{\mathrm{load}} = \dot{Q}_{\mathrm{load}} / (c_p (T_{S_1} - T_{\mathrm{ret,set}})).$$

The supply and return temperature of the heating distribution system are denoted by $T_{\mathrm{sup,set}}$ and $T_{\mathrm{ret,set}}$, respectively, and are calculated by using a heating curve depending on the ambient temperature and the characteristics of the building heat emission system. In order to satisfy the heating demand, the temperature of the upper layer of the storage tank must be sufficiently high, i.e. $T_{S_1} \geq T_{\mathrm{sup,set}}$ whenever $\dot{Q}_{\mathrm{load}} > 0$.

3. MPC FORMULATION

The economic nonlinear MPC problem that minimizes the cost of the used energy while ensuring that the provided water temperature is sufficiently high to be able to satisfy the building heating demand is as follows:

$$\min_{x,u,p,s} \int_0^{t_f} c_{\mathrm{el}}(t) P_{\mathrm{el}}(t) + c\, \dot{Q}_{\mathrm{load}}(t) s(t)^2 \, \mathrm{d}t \tag{1a}$$

$$\text{s.t.} \quad \dot{x}(t) = f(x(t), u(t), p(t)) \tag{1b}$$

$$x(0) = x_0 \tag{1c}$$

$$T_{S_1}(t) \geq T_{\mathrm{sup,set}}(t) - s(t) \tag{1d}$$

$$s(t) \geq 0 \tag{1e}$$

$$\dot{Q}_{\mathrm{HP,max}}(t) \geq \dot{Q}_{\mathrm{HP}}(t) \geq \dot{Q}_{\mathrm{HP,min}}(t) \tag{1f}$$

The state vector

$$x = (T_{S_1}, \ldots, T_{S_N})$$

holds the storage layer temperatures. The controllable input is the heat pump supply temperature

$$u = T_{\mathrm{HP}}.$$

The time dependent system parameters are the 24h-forecasts of

$$p = \left[\dot{Q}_{\mathrm{load}}, c_{\mathrm{el}}, T_{\mathrm{amb}}\right],$$

i.e. the non-negative demand for space heating, the grid signal and the ambient temperature, respectively. The supply and return set point temperature of the heat distribution system $T_{\mathrm{sup,set}}$ and $T_{\mathrm{ret,set}}$ can be directly calculated from the ambient temperature by the heating curve. In order to satisfy the user comfort demands, the upper storage layer temperature must be sufficiently high. This is realized by a soft constraint formulation (1d) with weighting factor c and slack variable $s(t)$ that is only active whenever the heating demand is positive. The thermal capacity of the heat pump depends on T_{HP} and is included as time variant lower and upper bounds $\dot{Q}_{\mathrm{HP,min}}(t)$ and $\dot{Q}_{\mathrm{HP,max}}(t)$ in Constraint (1f).

3.1 Relaxation of hybrid formulation

State-dependent discontinuities arise from the temperature dependent lower bound $\dot{Q}_{\mathrm{HP,min}}$, which equals 0 if the heat pump is off and a minimum output power larger than 6kW otherwise, and the upward or downward direction of the mass flows between different layers in the storage model. In order to apply a gradient-based optimization methods, the discontinuous terms are relaxed by using approximate heuristic rules, e.g. a post-processing step to meet the lower heat pump bound and the simplifying assumption that the building load mass flow is smaller than that of the heat pump.

3.2 Solution of the MPC problem

The optimal control problem is transformed into a nonlinear program using direct collocation (Rawlings et al., 2017) on a time grid of step size $h = 450$s, which also equals the sampling length for MPC and the high-level control in the test bed. This duration prevents an over-frequent cycling of the heat pump and is sufficiently small for operating the real plant. The MPC algorithm as well as the system model are implemented in Python, using the algorithmic differentiation framework CasADi (Andersson et al., 2018) and the nonlinear programming solver IPOPT (Wächter and Biegler, 2006). The prediction horizon is 24h, resulting in $n_{\mathrm{k}} = 192$ time steps. The solution time for a daily closed-loop simulation takes approximately 6-7 min on a standard PC.

4. DEEP LEARNING CONTROLLER

Neural networks (NN) are powerful function approximators that find a function $f_W : \mathbb{R}^{n_\xi} \to \mathbb{R}^{n_{\mathrm{u}}}$ that approximates the training data well and can generalize to unknown data from a test set. The training data $\{(\xi^{(1)}, u^{(1)}), \ldots, (\xi^{(m)}, u^{(m)})\}$ is composed of the input vector $\xi^{(j)} \in \mathbb{R}^{n_\xi}$ and the response variable $u^{(j)} \in \mathbb{R}^{n_{\mathrm{u}}}$.

4.1 Training process based on MPC data

Similar to the MPC closed-loop iteration, each input training data point $\xi^{(j)} \in \mathbb{R}^{n_\xi}$ consists of the state vector at the current time instance $(T_{S_1}, \ldots, T_{S_N})$ and the parameter forecasts in the form of time series of size $3 \times n_{\mathrm{k}}$ composed of the forecasts of the heating load, grid signal and ambient temperature. Hence, $n_\xi = 3n_{\mathrm{k}} + N$. The response variable $u^{(j)} \in \mathbb{R}^{n_{\mathrm{u}}}$ represents the heat pump temperature T_{HP} with $n_{\mathrm{u}} = 1$. The total number of data points m equals the number of MPC instances.

The resulting learned controller is a model-free control approach that can be learned from a supervisory MPC control action. The learned controller makes predictions based on the sensor information without solving an optimization problem. Simulated MPC data is used as expert data because collection of enough real world data for

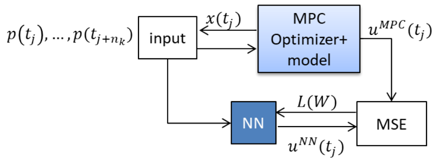

Fig. 3. Training process based on supervisory MPC

training would take a long time. The training process is illustrated in Figure 3. An output from the network is a linear combination of activation functions constructed from the inputs, which are passed through a non-linear transformation to produce the output prediction. Learning or training involves tweaking the weights W of the network such that the mean squared error (MSE) loss $L(W)$ between network output $f_W(\xi^{(j)}) = u^{NN}(t_j)$ and training data output $u^{MPC}(t_j)$ is minimized.

4.2 Network architecture

There exist different neural network architectures suitable for different tasks. Convolutional neural networks (CNN) are used in image processing to extract patterns from images. Since time series data can be considered as a 1D pattern in time, a modified CNN should be able to identify the output based on different input trajectories. Such patterns which the CNN is being trained to find could represent for instance the increase of heat pump power performed by MPC when it encounters low price signals and tries to store heat and the reduction of heat pump power when prices increase so that heat from the storage tank can be used. CNNs are comparable to the simpler feedforward networks but have at least one convolutional layer among the fully connected layers. The CNN architecture used is represented in Figure 4. An

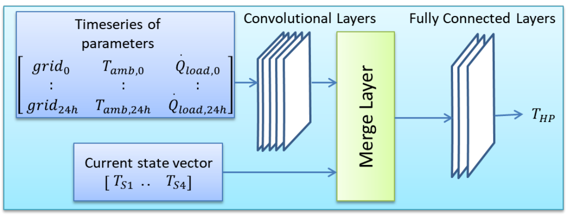

Fig. 4. CNN architecture with input features and output

input to the learning based controller is composed of two parts: (i) the forecast of grid price, ambient temperature and heat demand, as a time series matrix and (ii) the current state of the storage tank. The time series matrix is processed by the convolutional layers, which identify the patterns in the time series. These patterns combined with the storage temperature vector are then processed through a full connected layer, which predicts how much heat should be generated by the heat pump.

4.3 Weighted training

In order to increase the performance of the model, certain data points are given more emphasis during training. The intuition is that this will push the network towards identifying the patterns in the emphasized data points more effectively. During training this is accomplished by modifying the loss function with a new weight parameter $\alpha \in \mathbb{R}^m$ that represents the weights of each data point. In order to, e.g., give more weight to data points with peaks in the heating load, the weight for input points with large heat demand are increased. Similarly, more weight is given to data points with peaks and valleys in the grid price signal.

4.4 Constraint handling

The main challenge of operating the learned controller based on NN is the handling of constraints. While simple constraints on the outputs such as upper and lower bounds on the heat pump control can be easily enforced by choosing an appropriate output layer of the network, guaranteeing the satisfaction of state constraints of the system is an open problem. The difficulty here is that we can influence the internal state of the system only indirectly through the controls. Even if the expert data, i.e. the MPC controls, never leads to violations of the state constraints, the deviations of the sub-optimal NN controller could still lead to constraint violations. Because of these difficulties, we implement state constraints by a post-processing as described in Section 6.1.

5. PREDICTIVE RULE-BASED CONTROLLER

The current practice in building energy systems are rule-based controllers (RBC). They control the heat pump temperature in order to maintain the set-point temperature by heuristic rules. The employed RBC provides the required heat by increasing the heat pump temperature whenever there is heat demand:

(1) $\dot{Q}_{\text{HP}} = \min \dot{Q}_{\text{HP}}$ such that $T_{\text{HP}} \geq T_{\text{sup,set}}$ whenever $\dot{Q}_{\text{load}} > 0$, otherwise $\dot{Q}_{\text{HP}} = 0$.

An additional predictive rule-based controller (PRBC) is considered, which makes use of the forecast for the next 12 hours of the disturbance parameters. It employs a basic cost saving mechanism by operating the heat pump at full power when grid prices are negative. It includes the following additional rules:

(2) $\dot{Q}_{HP} = \max \dot{Q}_{HP}$ if $c_{\text{el}} \leq 0$
(3) $\dot{Q}_{HP} = \text{mean}\, \dot{Q}_{HP}$ if $\dot{Q}_{\text{load},0}, \dot{Q}_{\text{load},1} = 0$ and $\text{mean}\{\dot{Q}_{\text{load},2}, \ldots, \dot{Q}_{\text{load},6}\} > 10\text{kW}$.

Rule (3) is to pre-heat the storage tank if the heat pump was turned off and there is high demand in the horizon. This ensures that little comfort violation occurs.

6. RESULTS

6.1 Closed-loop simulation

The performances of MPC, the learned controller and the predictive and non-predictive rule-based controllers are compared in closed-loop simulation. At each simulation step, given the current storage state and parameter forecasts, the controller computes one control step $u(t)$, cf. Figure 5. A post-processing step is applied in order to make sure that the predicted temperature is within the heat pump bounds and not lower than the required building supply set-point.

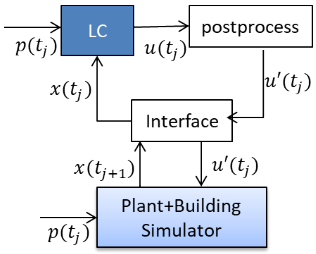

Fig. 5. Closed-loop simulation scheme

6.2 Data sets

For training the neural network, we use a 60 day data set of a standard office building and an artificial price signal for the electricity cost. The first day is excluded. For validation, two different test sets with different office buildings and grid signals are used:

(i) the first day of the training data set
(ii) one week of a new data set

Test set (i) is also used in the HIL experiment. The new data set (ii) contains different load, ambient temperature and grid trajectory based on the variable electricity price of the year 2012 [1]. The grid signals are plotted in each diagram on the right y-axis. The results are evaluated by the following indicators:

(1) cost of the used electricity is to be minimized,
(2) grid flexibility compared to the average grid signal,
(3) comfort deviation remains in an acceptable range, i.e., total deviation $\leq$ 1Kh/day and maximum deviation $\leq$ 2K (according to Klein et al. (2017)).

The grid flexibility (2) is evaluated by the grid support coefficient GSC (Klein et al. (2017)). It describes the proportion of time that the employed strategy uses electricity at a grid signal above (grid-adverse) or below (grid-supportive) average. In terms of a price-based grid signal, a grid-supportive strategy implies a cost saving compared to the average price. Due to the comparison to the mean grid signal, this term is suitable for a comparison between test sets of different grid signals and duration.

6.3 Simulation results

Test data set (i) Table 1 summarizes the results in terms of cost and comfort deviation. The trajectories resulting from the closed-loop simulation are given in Figures 6 and 7. The learned controller is better than both rule-based controllers in terms of lower cost and better grid support, but does not reach the performance of MPC. Both rule-based controllers follow the heating load curve. The predictive version additionally exploits the fact that negative prices occur frequently in the grid signal and shows a good performance.

Test data set (ii) The summarized results in Table 2 and the charging behavior of the storage tank in Figure 8 confirm the previous results. Even for this completely different data set, the learned controller is able to imitate the MPC controller in most cases. In terms of grid flexibility and cost, the simulation results show that the learned controller is approximately 10% better than the predictive rule-based controller and 10% worse than MPC. Compared to the previous test, negative grid prices do not occur frequently and the predictive RBC has a worse performance.

[1] European Energy Exchange AG. EEX transparency in energy markets, available from `www.transparency.eex.com/de`.

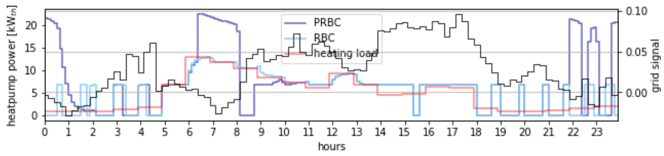

Fig. 6. Simulation of test case (i) showing heat pump behavior resulting from rule-based and predictive redule-based controller and building heating load.

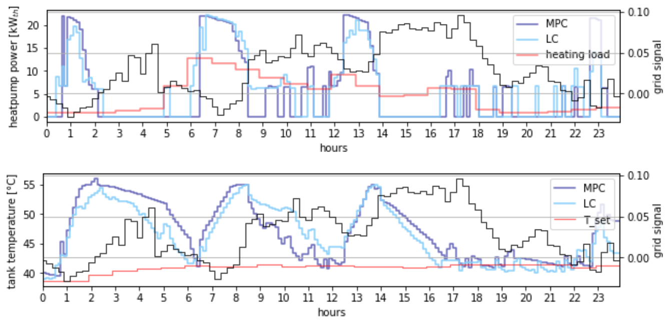

Fig. 7. Simulation of test case (i) showing heat pump behavior resulting from MPC and learned controller (top) and charging of storage tank (bottom).

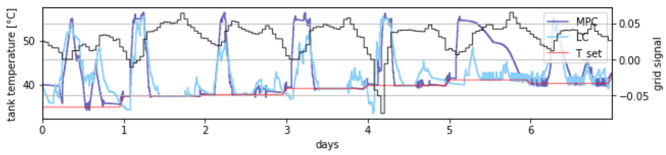

Fig. 8. Charging of storage tank for one week simulation (test case (ii)) of MPC and the learned controller.

6.4 Results of HIL test

The results of the HIL experiment for the control of the real plant by MPC and by the learned controller are plotted in Figure 9. Similar to the closed-loop simulation in

Table 1. Closed-loop performance in simulation for test case (i) (HIL test day).

	RBC	PRBC	LC	MPC
Cost	0.75	0.5	0.38	0.19
Energy [kWh]	129.0	171.5	138.8	128.6
Max. comfort dev. [K]	0.73	0.7	0.5	0.1
Grid support coefficient	1.11	0.44	0.40	0.19

Table 2. Closed-loop performance in simulation for test case (ii) (one week data set).

	RBC	PRBC	LC	MPC
Cost	4.79	4.69	4.18	3.50
Energy [kWh]	780.6	814.3	828.1	755.0
Max. comfort dev. [K]	1.4	0.0	1.2	0.8
Grid support coefficient	1.19	1.06	0.86	0.77

Figure 7, the learned controller imitates MPC in the most important regions and can exploit the characteristics of the grid signal. The summarized results concerning cost, flexibility and comfort deviation in Table 3 show that the learned controller operates the system equally well as MPC. This is due to the fact that in the numerical simulation experiments the same model is used for the MPC controller and for system simulation while in the HIL experiment there is a model-plant mismatch due to unmodeled plant behavior (such as the internal heat pump controller). This affects MPC performance.

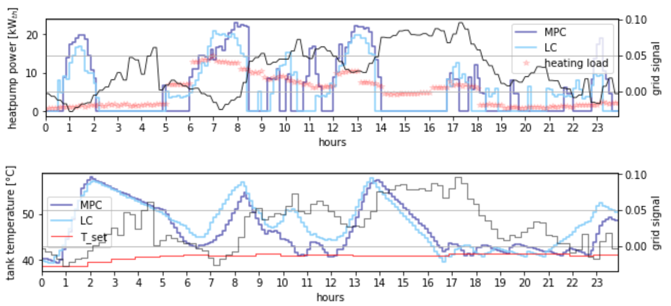

Fig. 9. One day HIL experiment of MPC and learned controller showing heating load and heat pump behavior resulting from MPC and learned controller (top) and charging of storage tank (bottom).

Table 3. HIL results

	MPC	LC
Cost	0.56	0.55
Energy [kWh]	138.2	128.5
Max. comfort dev. [K]	0.4	1.4
Grid support coefficient	0.51	0.46

7. CONCLUSION

In this paper, we analyzed how well a model-free control approach can be learned from a supervisory MPC control action such that the learned controller gives predictions based only on the sensor information. The evaluation of the learned controllers was done in a closed-loop simulation with different test sets and by implementation in a HIL test bed. In terms of grid flexibility, the simulation results show that the learned controllers are approximately 10% better than the predictive rule-based controller and 10% worse than MPC. Also the operation cost show a similar performance however the exact extend depends on the considered test set. In the practical HIL experiment, the learning-based controller operates the system equally well in terms of comfort violation, cost and grid flexibility as the MPC controller does.

REFERENCES

Afram, A., Janabi-Sharifi, F., Fung, A.S., and Raahemifar, K. (2017). Artificial neural network based model predictive control and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy and Buildings*, 141, 96–113.

Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., and Diehl, M. (2018). CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*.

Balint, A. and Kazmi, H. (2019). Determinants of energy flexibility in residential hot water systems. *Energy and Buildings*, 188-189, 286 – 296.

Bürger, A., Zeile, C., Altmann-Dieses, A., Sager, S., and Diehl, M. (2018). An algorithm for mixed-integer optimal control of solar thermal climate systems with mpc-capable runtime. In *2018 European Control Conference (ECC)*, 1379–1385.

Cigler, J., Gyalistras, D., Široky, J., Tiet, V., and Ferkl, L. (2013). Beyond theory: the challenge of implementing model predictive control in buildings. In *Proceedings of 11th Rehva world congress, Clima*, volume 250.

Drgoňa, J., Picard, D., Kvasnica, M., and Helsen, L. (2018). Approximate model predictive building control via machine learning. *Applied Energy*, 218, 199 – 216.

Eicker, U. (2003). *Solar technologies for buildings*, volume 341. Wiley Online Library.

Fischer, D., Bernhardt, J., Madani, H., and Wittwer, C. (2017). Comparison of control approaches for variable speed air source heat pumps considering time variable electricity prices and PV. *Applied Energy*, 204, 93–105.

Frison, L., Kleinstück, M., and Engelmann, P. (2019). Model-predictive control for testing energy flexible heat pump operation within a hardware-in-the-loop setting. *Journal of Physics: Conference Series*, 1343, 012068.

Halvgaard, R., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2012). Economic model predictive control for building climate control in a smart grid. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, 1–6.

Jain, A., Smarra, F., and Mangharam, R. (2017). Data predictive control using regression trees and ensemble learning. In *Proceedings of the 2017 Conference on Decision and Control*. IEEE.

Klein, K., Herkel, S., Henning, H.M., and Felsmann, C. (2017). Load shifting using the heating and cooling system of an office building: Quantitative potential evaluation for different flexibility and storage options. *Applied Energy*, 203, 917 – 937.

Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Morari, M. (2012). Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45, 15 – 27.

Péan, T., Salom, J., and Costa-Castelló, R. (2018). Review of control strategies for improving the energy flexibility provided by heat pump systems in buildings. *Journal of Process Control*.

Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI.

Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1), 25–57.

Zong, Y., Böning, G.M., Santos, R.M., You, S., Hu, J., and Han, X. (2017). Challenges of implementing economic model predictive control strategy for buildings interacting with smart energy systems. *Applied Thermal Engineering*, 114, 1476 – 1486.