PROCEEDINGS OF ECOS 2023 - THE 36TH INTERNATIONAL CONFERENCE ON
EFFICIENCY, COST, OPTIMIZATION, SIMULATION AND ENVIRONMENTAL IMPACT OF ENERGY SYSTEMS
25-30 JUNE 2023, LAS PALMAS DE GRAN CANARIA, SPAIN

Machine-learning-based approximation of the hierarchical model predictive control of multi-use PV-battery systems in non-residential buildings

Laura Maier[a], Sönke Quast[a], Dominik Hering[a] and Dirk Müller[a]

[a] RWTH Aachen University, E.ON Energy Research Center, Institute for Energy Efficient Buildings and Indoor Climate, Aachen, Germany, laura.maier@eonerc.rwth-aachen.de, CA

Abstract:

Model predictive control (MPC) has proven to be a promising method to exploit energy saving potentials in building energy systems. However, they are not widespread in practice due to high hard- and software requirements, high computational effort, and missing trust and know-how among practitioners. Approximate MPC can address these challenges by replacing the hard- and software-intensive optimization program by black box models. Machine learning models such as Artificial Neural Networks or tree-based algorithms have been widely investigated by the scientific community. However, a comparison of Artificial Neural Networks with advanced tree-based models like Random Forest and Gradient Boosting is still missing. In addition, the relation between the models' training and the resulting control performance has not yet been assessed. We close these gaps by investigating the optimal control based on an MPC of a PV-battery system in a non-residential building. The MPC optimizes the battery's power based on a preceding peak load optimization. The MPC is imitated by three machine learning models, namely, an Artificial Neural Network, a Random Forest, and Gradient Boosting, whose performance is subsequently evaluated open- and closed loop. While Gradient Boosting results in the highest open-loop performance with an R^2 of 0.83, it deviates more significantly from the optimal control trajectory than, e.g., the Artificial Neural Network. Nonetheless, Gradient Boosting even outperforms the teacher MPC when considering the system's annuity. This is explained by its ability to push beyond the peak load constraints which are set within the optimization. A rule-based backup controller is, therefore, included for all approximator-based controllers. Based on this, the approximators result in a peak load reduction between 5 % and 7 % compared to the benchmark and a change in annuity between -1 % and 4 % compared to the MPC. To summarize, all approximators can retain most of the MPC's advantages but do not surpass its overall performance.

Keywords:

Approximate Optimal Control, Model Predictive Control, Machine Learning.

1. Introduction

In 2022, the United Nations Environment Program published its global status report for the buildings and construction sector, highlighting the need for immediate action to cut emissions [1]. Following a slowdown caused by the global pandemic, the CO_2 emissions from building operations in 2021 even exceeded their all-time maximum of 2019 by 2 %. When taking into account the emissions of building material production, buildings made up around 37 % of the global CO_2 emissions in 2021 [1]. In addition to investing in refurbishment strategies, optimizing building energy system operation can contribute to the goal of an emission-free building sector by 2050 [1, 2].

Model predictive control (MPC) as a representative of advanced control methods has proven to optimize building operation tremendously. In a review study, Drgoňa et al. [2] find that the realized energy savings range between 15 % and 50 % based on selected case studies [3–6]. MPC is also a valuable method for providing grid flexibility services by, e.g., exploiting price incentives or providing demand response services [2]. The latter is crucial for future grid operation as the building sector will be electrified and, hence, interact with the grid more intensely. However, despite its significant potential MPC is not widespread in practical applications. The reasons lie in the high requirements of hard- and software and data infrastructure, missing know-how among the operating personnel, time-consuming modeling and maintenance, and network and privacy security concerns [2, 7–10]. This is why conventional rule-based and PID controllers are still state-of-the-art in nowadays buildings. A method that bridges the gap between the performance of MPC applications and the simplicity of rule-based controllers is approximate MPC. The idea is to imitate the MPC's output using black box models. The MPC serves as a teacher that generates training data. The implicit MPC-based controller is finally replaced by an explicit controller [11]. The black-box-based explicit controller requires less advanced hard- and software as well as data infrastructure. The used black box models in building energy systems range from

https://doi.org/10.52202/069564-0287

simple linear regression models [9, 12], over decision trees [11, 13–16] to sophisticated machine learning models like Artificial Neural Networks (ANNs) [11, 17–20]. These so-called approximator-based controllers have proven to retain most of the MPC's performance while significantly reducing the required computational effort and processing time [11]. In this context, decision-tree-based approximators are favorable as they resemble the rule-based controller-like "if-condition-then-action" structure and, hence, promote comprehensibility and address the challenge of missing know-how. However, decision trees tend to overfit data and are very sensitive to input data [21]. Consequently, advanced methods like ANNs are often selected as approximators [11, 19, 20]. High-performing ensemble methods like Random Forest (RF) and Gradient Boosting (GB) can address the decision trees' disadvantages of overfitting. However, to the authors' best knowledge, studies investigating high-performing, tree-based machine learning algorithms like RF and GB in the context of approximate MPC are still missing. Furthermore, even though approximate MPC has been successfully applied to building energy systems, there exists no use case focusing on a grid-connected PV-battery system as a part of a building energy system.

In contrast to that, the scientific community has come up with related methods targeting purely grid-related challenges. Smart grids will play a key role in the energy transition to guarantee safe grid operation. The underlying concept is referred to as the optimal power flow problem, which can focus on AC- and DC-based applications. Here, the main challenge lies in the real-time solution of a highly complex optimization problem that optimizes power flows in a grid, e.g., the optimal power that a set of generators have to produce [22–25]. [23] and [24], e.g., both focus on an AC optimal power flow problem. For example, Zamzam et al. [23] learn the mapping of system loading and optimal generation values using an ANN serving as an input for the subsequent power flow solver. Doing so, they speed up the calculation process by a factor between 8 and 15 while still achieving near-optimal results compared to directly solving the optimal power flow problem. Furthermore, the authors of [24] approximate the cost function and give a feasibility prediction for the AC optimal power flow problem. Again, the accuracy is high and the computational effort is reduced significantly. As approximators, they compare the performance of linear regression, piecewise regression, Gaussian Processes, and ANNs. Apart from that, [22] use graph NNs and [25] test deep NNs, and obtain similar trends for an AC and DC optimal power flow problem, respectively. The aforementioned grid-focused studies all apply their method to simplified use cases and learn the output of static optimal power flow problems. De Jongh et al. [26] are the first to use ANNs to learn a dynamic MPC-based problem of a smart distribution grid. The distribution grid covers 15 nodes and 13 households and considers electric vehicles and heat pumps. The MPC optimizes the power schedule for all flexible devices. The ANN is trained based on a full-year simulation with three months of testing. This time, the closed-loop performance is evaluated and, again, near-optimal results are achieved with a 55 times quicker processing time.

In addition to these aspects, the training process for the approximators differs significantly among the studies. While the grid-focused studies generate randomized training samples using distribution assumptions for their inputs [22, 24, 25], the building-focused ones tend to use closed-loop operation results [11, 12, 14]. The training performance is usually evaluated open-loop, i.e., the true and predicted outputs are compared without system interaction. However, it is unclear how the open-loop performance affects the closed-loop one, which we identify as an additional research gap.

1.1. Contributions and structure of this study

To summarize, the state of research highlights that approximate MPC applications are promising for building energy systems. However, we detect a gap in studies focusing on building energy systems while considering their interaction with the grid. In addition, studies are missing that compare sophisticated machine learning models like ANNs to advanced tree-based ones like RF and GB. Furthermore, the scientific literature has not yet presented a detailed investigation of the relation between the open- and closed-loop performance.

This study closes these gaps through the following aspects:

- We present a PV-battery system of a non-residential building in Berlin, Germany (see subsection 2.1.). The battery's operation is optimized based on an MPC, used as teacher MPC, that has been presented in previous work [27] (see subsection 2.2.).
- The teacher MPC is taken as a training basis for three machine learning models, namely, an ANN, an RF, and GB. They are trained based on a full-year simulation of the MPC-controlled system (see subsection 2.4.).
- section 3. presents the open- and closed-loop performance of these approximators, which are used as substitutes for the MPC. For this, we define open- and closed-loop metrics (see subsection 2.5.).
- We finally discuss the relation between the open- and closed-loop performance as well as the comparison of the machine learning models in section 4., conclude our findings (see section 5.), and give an outlook into future work (see section 6.).

https://doi.org/10.52202/069564-0287

By investigating these aspects, we try to answer the two following research questions in the context of approximate MPC:

1. Can advanced tree-based machine learning models outperform the most prominent Artificial Neural Network?
2. Is there a correlation between open- and closed-loop performance?

2. Methodological approach

2.1. Use case: PV-battery system of a non-residential building

The use case is the electrical system of a laboratory and office building in Berlin, Germany. Figure 1 illustrates the central components and their interaction. The energy supply for heating, ventilation, and air conditioning (HVAC) is solely based on electricity. A battery energy storage system aims at reducing potentially arising peak loads and maximizing the electricity generation of the PV power plant. In addition to the demands caused by the HVAC system, the electrical load of the building contains the electricity consumption of the building's, i.e., its tenants', equipment (servers, laboratory equipment, etc.) and lights. The overall system can exchange electricity with the grid.

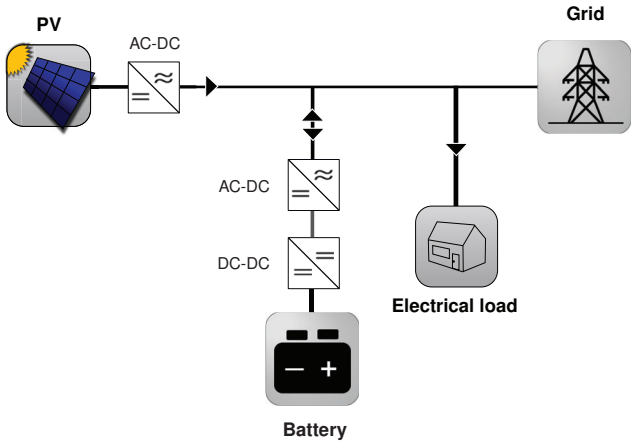

Figure 1: Use case: electrical system setup of the non-residential building including a battery energy storage system, a PV power plant, the building load, and the grid.

The building's electrical loads have been simulated using Modelica as a modeling language. The underlying toolchain to obtain realistic electrical profiles has been presented in [28]. The simulated electrical load is illustrated in Figure 2. The resulting maximum electrical load of the building is 955 kW, the battery's capacity is 2500 kWh and its maximum discharging and charging power is 1250 kW. The PV power plant has a peak power of 500 kWp.

2.2. Model predictive control and data base

In Figure 3, the MPC and AMPC toolchain is illustrated. The target system, which is described in subsection 2.1., is simulated in the modeling language Modelica using the functional mock-up interface. The battery's charging and discharging powers are the manipulated variables that are transferred using the open-source Python package *fmpy*. For every iteration, these manipulated variables are optimized by an mixed-integer linear program (MILP). The MILP is formulated in the Python-based optimization modeling language *Pyomo* [29]. Based on perfect forecast of the disturbances, the battery's state of charge and other state variables, the optimization is solved for two different prediction horizons, resulting in an hierarchical structure. The upper optimization layer is a full-year optimization that determines the optimal electrical peak. This electrical peak is subsequently transferred as a constraint to the second optimization stage that follows the receding horizon scheme. Here, a prediction horizon of 16 h and a timestep of 900 s are applied and the control loop is repeated on an hourly basis. The lower layer optimization optimizes the battery's power based on the PV generation, the building load, as well as economic boundary conditions, such as the electricity price. Since the MPC-based battery operation deals with different sources of revenue, we refer to it as multi-use PV-battery system. Further details on the MILP and the hierarchical MPC are presented in [27].

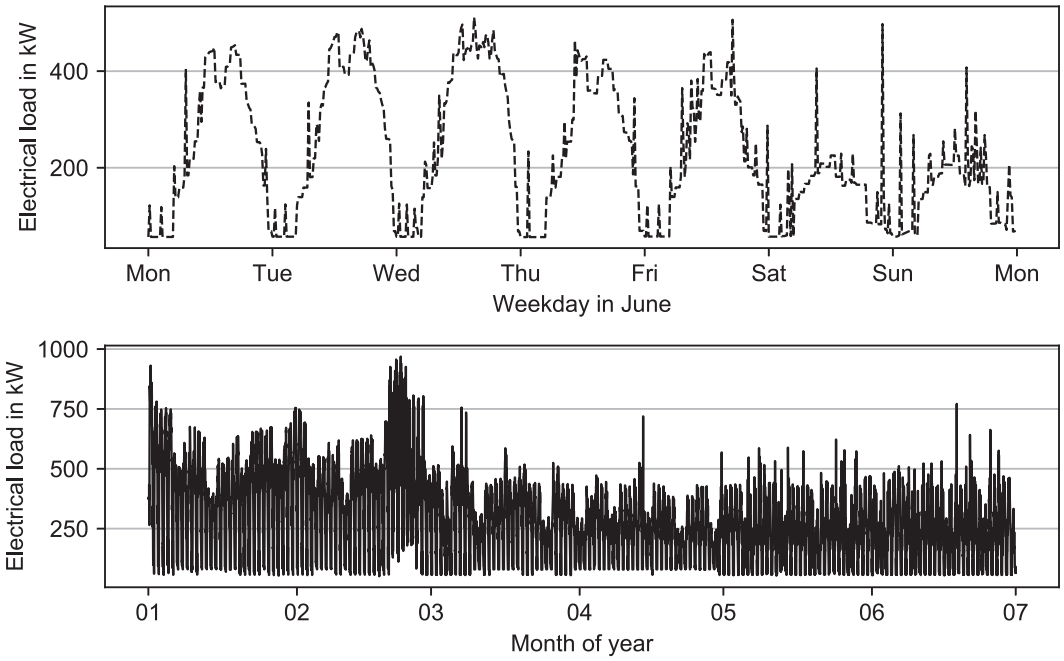

Figure 2: Simulated electrical building load for an exemplary week (top) and the first 6 months (bottom).

The perfectly predicted disturbances comprise:

- **Weather forecasts:** Ambient temperature, global irradiation, and wind speed to compute PV power output based on standardized test reference years [30]
- **Real time electricity prices:** Time-varying EEX electricity prices of 2019 are taken as a basis [31].
- **Electrical building load:** The electricity consumption of the building is simulated in the modeling language Modelica using typical user profiles [28, 32].

This study's aim is to replace the implicit optimization-based controller by an explicit black-box-model-based one. Therefore, we use the open-source machine learning package scikit-learn [33].

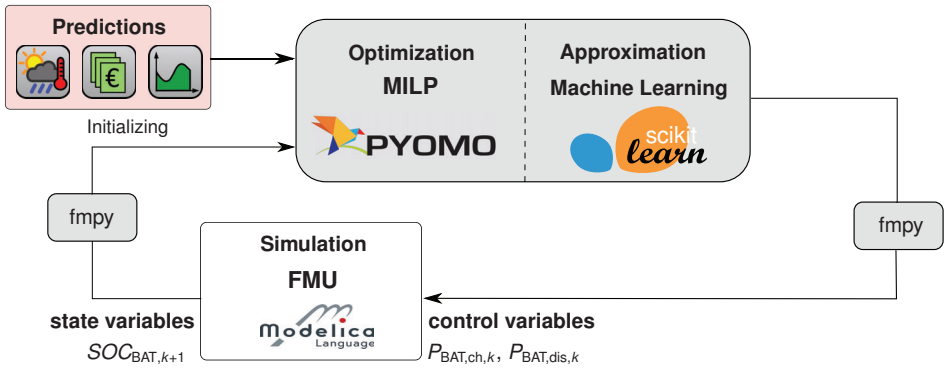

Figure 3: MPC and approximate MPC framework and control loop that calculates the battery's charging $P_k^{BAT,ch}$ and discharging power$P_k^{BAT,ch}$ to the simulation model for each timestep k. Based on measurements like the battery's state of charge SOC_k^{BAT}, the next iteration starts.

https://doi.org/10.52202/069564-0287

2.3. Feature construction and selection

The general approach for approximate MPC is to provide the MPC's input data for the approximator so that it can learn the MPC's output. However, it is recommended to include additional features that support the imitation learning process. For the present study, 14 features have been selected, which are derived based on sensitivity analyses and listed in Table 1. The features comprise measurement data (simulated), set points of a former optimization, synthetic data, and disturbance predictions. The measurements of the previous $k-1$ timestep are used to predict the next timestep's k output. Among these measurements are the simulated battery charging or discharging power, its SOC, the cells' temperatures, the voltage as well as the power from or to the grid. Furthermore, the top layer optimization's computed optimal peak load of timestep k as well as future timesteps serves as input to the approximators. In addition to the measurements, disturbance predictions for 64 timesteps are included, standing for a time period of 16 h, i.e., the MPC's prediction horizon.

Table 1: Overview over features selected to predict the MPC's output.

Feature	timestep	Type
Power of battery (charge or discharge)	k-1	Measurement
State of charge of battery	k-1	Measurement
Temperature of battery cells	k-1	Measurement
Voltage in battery stack	k-1	Measurement
Power to/from grid	k-1	Measurement
Maximum allowed grid demand	k,...,k+63	Optimization set point prediction
Building's electrical load	k,...,k+63	Disturbance prediction
PV electricity generation	k,...,k+63	Disturbance prediction
Electricity price	k,...,k+63	Disturbance prediction
Ratio of building load to maximum allowed grid demand	k,...,k+63	Synthetic (measurement and disturbance prediction)
Ratio of PV generation to building load	k,...,k+63	Synthetic (measurement and optimization set point prediction)
Hour of day	k	Synthetic
Day of week	k	Synthetic
Month of year	k	Synthetic

2.4. Imitation learning process

In approximate MPC applications, the optimization is replaced by an approximated model that imitates the so-called teacher MPC. Machine learning models are the most common approximators in this context. This study deals with a continuous manipulated variable, namely the charging $P_{\mathrm{BAT,ch}}$ and discharging $P_{\mathrm{BAT,dch}}$ power of the battery. We simplify the manipulated variable to P_{BAT}. Due to the continuous nature of the variable, the approximator needs to perform a regression task. The objective of this regression task is to find a regression function $f_\theta : \mathbb{R}^{n_\varepsilon} \rightarrow \mathbb{R}$ that minimizes the squared error between the true manipulated variable P_{BAT} being the MPC's output and the predicted one $\hat{P}_{\mathrm{BAT},k}$ by tuning the parameters θ:

$$\min_{\theta} \sum_{k=1}^{n} (\hat{P}_{\mathrm{BAT},k} - P_{\mathrm{BAT},k})^2, k \in \mathbb{N}, \theta \in \mathbb{R}^{n_\theta} \tag{1}$$

The training data is discretized for each timestep k over the number of relevant samples n.

The regression parameter fitting process is carried out using the open-source Python framework *AddMo* (Automated data-driven modeling) [34]. *AddMo* comprises all relevant steps needed to obtain a well-trained machine learning model. The steps cover the data tuning including data preprocessing, period selection, and feature construction. Subsequently, the tool enables automatic feature selection and hyperparameter tuning. The framework's basis is the open-source machine learning library *scitkit-learn* [33] and covers ANNs, GB, Lasso, RF, and Support Vector Regression as potential model choices. The interested reader is referred to [34] for a more detailed framework presentation.

This study selects ANNs, GB, and RF as approximators as motivated in section 1.

2.5. Open- and closed-loop evaluation scheme

The evaluation includes both the open- and the closed-loop performance of the MPC- and approximator-based controllers. The open-loop performance aims at assessing the approximator's ability to predict the MPC's output, i.e., the charging and discharging power of the battery. As the training process is categorized as supervised machine learning for regression tasks, the coefficient of determination R^2 and the mean absolute

error *MAE* are selected as key performance indicators. Specifically, for each timestep k, the true and predicted output are compared and the respective error metrics are determined. As a simplified initial comparison, these statistical metrics are also computed for the closed-loop operation even though the focus should lie on system- and control-specific metrics and should, thus, be oriented towards the MPC's objectives. For the MPC, the economic evaluation is based on the PV-battery system's annuity. The annuity C_{ANN} simultaneously considers the investment costs (CAPEX, C_{CAPEX}) as well as the operating costs (OPEX, C_{OPEX}). In the context of this study, the following calculation scheme is applied, which is based on the German pricing and regulatory system:

$$C_{\mathrm{ANN}} = f_{\mathrm{AF}} C_{\mathrm{CAPEX}} + C_{\mathrm{OPEX,energy}} + C_{\mathrm{OPEX,power}} \tag{2}$$

$$C_{\mathrm{CAPEX}} = c_{0,\mathrm{PV}} P_{\mathrm{PV,peak}} + c_{0,\mathrm{BAT}} E_{\mathrm{BAT,cap}} \tag{3}$$

$$\begin{aligned} C_{\mathrm{OPEX,energy}} = {} & E_{\mathrm{grid,load}}(c_{k,\mathrm{EEX}} + c_{\mathrm{Tax}} + c_{\mathrm{EEG}}) + E_{\mathrm{grid,BAT}}(c_{k,\mathrm{EEX}} + c_{\mathrm{Tax}} + 0.6 c_{\mathrm{EEG}} \eta_{\mathrm{BAT}}^2) \\ & + 0.4(E_{\mathrm{PV,load}} + E_{\mathrm{BAT,load}}) c_{\mathrm{EEG}} + c_{\mathrm{Network}}(E_{\mathrm{grid,load}} + E_{\mathrm{grid,BAT}}) - E_{\mathrm{BAT,grid}} c_{k,\mathrm{EEX}} - E_{\mathrm{PV,grid}} c_{\mathrm{PV,feedin}} \end{aligned} \tag{4}$$

$$C_{\mathrm{OPEX,power}} = P_{\mathrm{peak}} c_{\mathrm{Peak}} \tag{5}$$

The associated costs of Equation 2 to Equation 5 are listed in Table 2. In the context of this paper, we define costs as positive quantities. In Equation 2, C_{CAPEX} includes the specific investment costs of the battery $c_{0,\mathrm{BAT}}$, the PV power plant $c_{0,\mathrm{PV}}$, including the required DC-DC and DC-AC inverters, respectively (see Equation 3 and Table 2). The battery's investment costs depend on its expected lifetime that is affected by aging. Calendrical and cyclical aging is considered in the simulation model. More specifically, keeping the battery on high SOCs or triggering many cycles leads to a degradation of its capacity and performance. For more details, the interested reader is referred to [27]. The CAPEX are multiplied with the annuity factor f_{AF} that depends on the interest rate i and the observation period T. We assume that all components except for the battery have a lifetime of T. If the battery's lifetime is shorter than the observation period, we consider a price degression d. The operation-related costs are further divided into energy demand- $C_{\mathrm{OPEX,energy}}$ and power-related operating costs $C_{\mathrm{OPEX,power}}$ (see Equation 2). For $C_{\mathrm{OPEX,energy}}$, the energy flows of the PV-battery system, the building, and the grid must be distinguished since different pricing schemes apply. Equation 4 denotes energy flows from source to sink. For example, $E_{\mathrm{grid,load}}$ is the building's consumed electrical energy covered by the grid. In addition to taxes c_{Tax} and network charges c_{Network}, the German pricing scheme includes a charge to support renewable energy sources c_{EEG}, whose quantity depends on the energy source. Therefore, Equation 4 differentiates between flows of the PV power plant, the grid, and the battery. Furthermore, the operation-related costs depend on the EEX market prices of each timestep k. The EEX prices also serve as source of revenue if electrical power is fed back into the grid. As an additional revenue, fed-in electricity from the PV power plant is rewarded with the feed-in price $c_{\mathrm{PV,feedin}}$.

Apart from the purely economic evaluation, this study also focuses on the system's peak load P_{peak} . The system's peak load is taken as an additional metric because the teacher MPC's aims to determine the system's optimal maximum peak load and control the battery's power accordingly. The associated costs are calculated based on a peak power price c_{Peak} (see Equation 5).

Table 2: Assumptions of economic boundary conditions.

Type	Mathematical description	Quantity
Initial invest in battery including DC-DC inverter	$c_{0,\mathrm{BAT}}$	725 €/kWh
Initial invest in PV power plant including DC-AC inverter	$c_{0,\mathrm{PV}}$	1170 €/kWp
Lifetime of battery	T_{BAT}	Simulated
Observation period	T	20 a
Relative price degression	d	6 %/a
Interest rate	i	1.3 %
Real-time pricing	$c_{k,\mathrm{EEX}}$	EEX 2019
PV feed-in pricing	$c_{\mathrm{PV,feedin}}$	6.62 ct /kWh
Tax-related charges	c_{Tax}	3.17 ct /kWh
Network charges	c_{Network}	1.65 ct /kWh
Peak power charges	c_{Peak}	53.53 €/kW

https://doi.org/10.52202/069564-0287

3. Results

The following section presents both the open- and the closed-loop results. The open-loop results stem from comparing the approximator's and the MILP's output regarding the battery's charging and discharging power. Consequently, the control loop is not closed. In contrast, the second part of this section focuses on closed-loop simulation results, for which the approximators are used to control the PV-battery system.

3.1. Open-loop performance

The open-loop analysis compares the predicted output $\hat{P}_{\mathrm{BAT}}$ with the true MPC output P_{BAT} based on a time series comparison. Open-loop training serves as an indicator of how well the machine learning models imitate the controller's output. However, since there is no interaction with the system, we cannot conclude on the closed-loop performance. Figure 4 illustrates the coefficient of determination R^2 as well as the mean absolute error *MAE* for the testing period for both open- and closed-loop operation. For this section, we concentrate on the open-loop results, i.e., the red bars and line graphs. The models are trained based on six months of training data and tested on six months of unseen data. The first half of the year serves as training data, while the second half serves as a testing period. The overall open-loop accuracy is high for all three models. GB results yields the highest R^2 of 0.83 and the lowest *MAE*, while ANN results in the lowest R^2. The ANN and RF yield a similar performance when taking the *MAE* as a basis. However, the ANN's R^2 is slightly lower than the RF's one. Consequently, the GB is favorable from an open-loop performance perspective. Nonetheless, the variations in performance are slight.

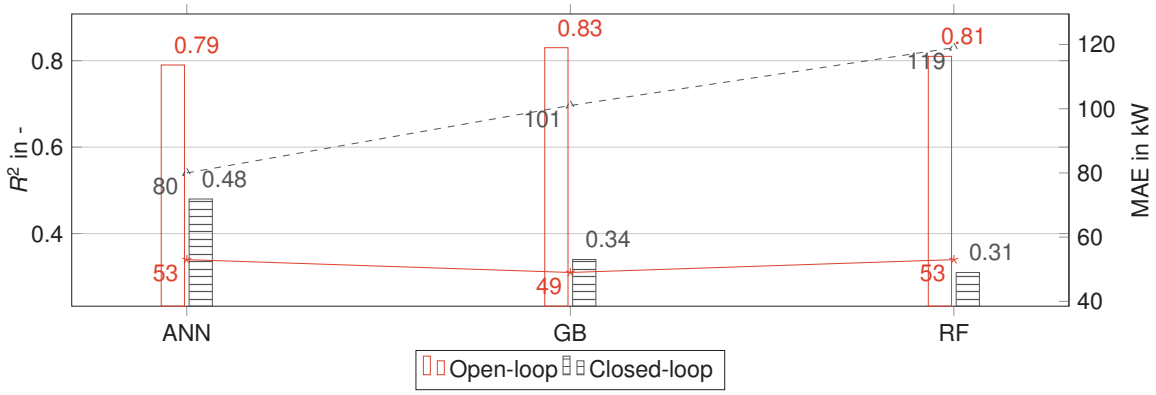

Figure 4: Time-series-based open- and closed-loop prediction results for 6 months training and testing.

3.2. Closed-loop performance

In addition to the open-loop performance not considering any interaction with the simulation model, this section investigates the closed-loop performance. At first, we focus on a purely time-series-based comparison between the open- and closed-loop performance. For this evaluation, Figure 4 also shows the statistical metrics for the closed-loop operation in addition to the open-loop comparison. The data is based on a closed-loop simulation of 6 months, from January to the end of June. The metrics derive from comparing the controllers' outputs, i.e., the battery's set charging and discharging power. The results indicate a different picture compared to the open-loop performance. The ANN clearly outperforms the tree-based models with an R^2 of 0.48 compared to 0.34 for GB and 0.31 for the RF, respectively. The same trend applies to the *MAE*. Based on this, we cannot see any correlation between the open- and the closed-loop performance of the approximators. However, the comparison above is based on time series only and does not consider any MPC performance metrics.

For this reason, we analyze the actual closed-loop performance taking MPC relevant objectives into consideration (see subsection 2.5.). The bar chart on the left of Figure 5 illustrates the resulting annuity of the PV-battery system for the MILP and the approximators, i.e., the ANN-, GB-, and RF-based controller. We like to highlight at this point that the annuities are negative because we mainly consider costs and the rewards through electricity feed-in are small for our use case. The resulting annuity ranges between $-425\,\mathrm{k€}$ for the GB-controlled and $-440\,\mathrm{k€}$ for the ANN-controlled system. Among the approximators, the GB performs best, followed by the RF. The ANN results in the lowest annuity. The GB-controlled system even outperforms the MILP regarding the annuity.

Figure 6 further illustrates the resulting operation for an exemplary week in June. The top plot shows the manipulated variables, namely, the battery's set power for the MILP- and approximator-controlled system. In the middle, the resulting grid load is depicted. As a reference, the optimized peak load is marked, too. A

https://doi.org/10.52202/069564-0287

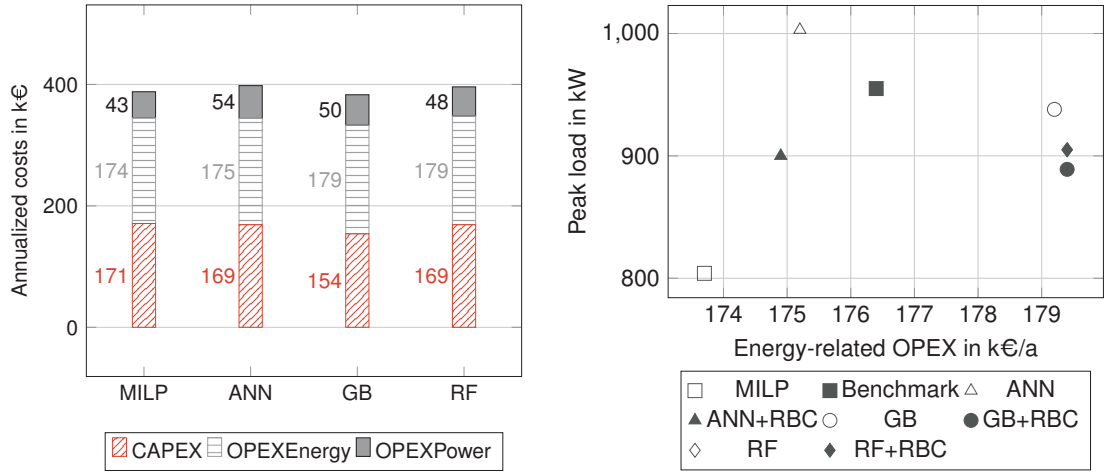

Figure 5: Closed-loop simulation results: The figure illustrates the holistic annuity of the PV-battery system without backup controllers (left) and the comparison of the resulting electrical peak load and demand-related operating costs with (+RBC) and without backup controllers (right).

positive battery power denotes battery charging, while a negative one signifies discharging. Furthermore, positive grid loads denote an energy flow from the grid to the system and vice versa. The bottom part of Figure 6 shows the variable electricity tariff. We observe that the approximators sufficiently imitate the teacher MPC's output. However, the resulting grid load shown in the middle part highlights that some approximators exceed the optimized grid limit. The tendency to overshoot the grid limit applies the most for the ANN-based controller. On the illustrated Saturday afternoon, the ANN lets the battery charge due to low electricity prices. This trend is partially observed for the residual controllers but to a smaller extent. The high charging power results in a grid load exceeding the optimized peak load. This behavior is also present on Thursday and Friday afternoon. Again, the ANN triggers a charging period due to low electricity prices even though the teacher MPC's output is 0 kW. Here, we detect the approximator's correlation between the in- and output without considering the constraints set within the MILP.

Figure 6 highlights the need for an additional backup controller that reinforces the optimized peak. To support the approximators in reducing the peak power, an additional rule-based backup controller is considered. The controller is non-predictive, meaning it relies on present or past measurements of the system. Based on the current grid load and the approximator's output, the rule-based controller limits the battery's set charging or discharging power. The result is depicted on the right-hand side of Figure 5. Here, we illustrate the realized electrical peak load and the energy-related OPEX for two sets of an approximator-controlled, the MILP-controlled system, and a benchmark model. The two sets of approximators consider the machine learning models with and without additional rule-based backup controllers ("X+RBC"). The MILP, i.e., the MPC, runs a pre-optimization for a whole year to determine the optimal peak load. This peak load is 803 kW for the given scenario and a constraint for the MILP. Hence, the MILP does not surpass the optimal peak, so the maximum peak for the 6-months simulation is 803 kW, too. In addition, the MPC also yields the lowest energy-related operating costs $C_{\mathrm{OPEX,energy}}$ of approximately 174 k€. Apart from that, we determine a benchmark model that does not include a battery. We solely include the benchmark model in the right-hand side plot of Figure 5 because the benchmark does not involve an investment in the PV-battery system. Consequently, a comparison based on the annuity is not expedient. The benchmark control does not shift any load and consequently solely depends on the variable electricity price. It yields a peak load of 955 kW and operating costs of 176 k€. Thus, the MPC achieves a peak load reduction of 16 % and a decrease in energy-related operating costs of 1 %. In general, the energy-related operating costs do not vary significantly among the controllers. However, it is noteworthy that high peak loads lead to higher power-related OPEX $C_{\mathrm{OPEX,power}}$. Combining both energy- and power-related OPEX, the MPC results in OPEX savings of 5 %.

In contrast, we observe a great difference in the peak loads. Without integrating a rule-based controller, only the GB- and the RF-based approximators realize a smaller peak load than the benchmark model of 938 kW and 905 kW, respectively. For the RF, the integration of a backup controller has only a negligible effect on the peak load of 0.3 kW. For the GB and the ANN, the effect is more significant. While the peak load is reduced from 938 kW to 889 kW in the case of GB, it decreases from 1003 kW to 900 kW for the ANN. The ANN's tendency to overshoot the peak load is also apparent in Figure 6. Nonetheless, when considering the rule-based controller, the annuities decline by 2 % for the ANN, increase by 0.5 % for GB and by 0.5 % for the RF.

https://doi.org/10.52202/069564-0287

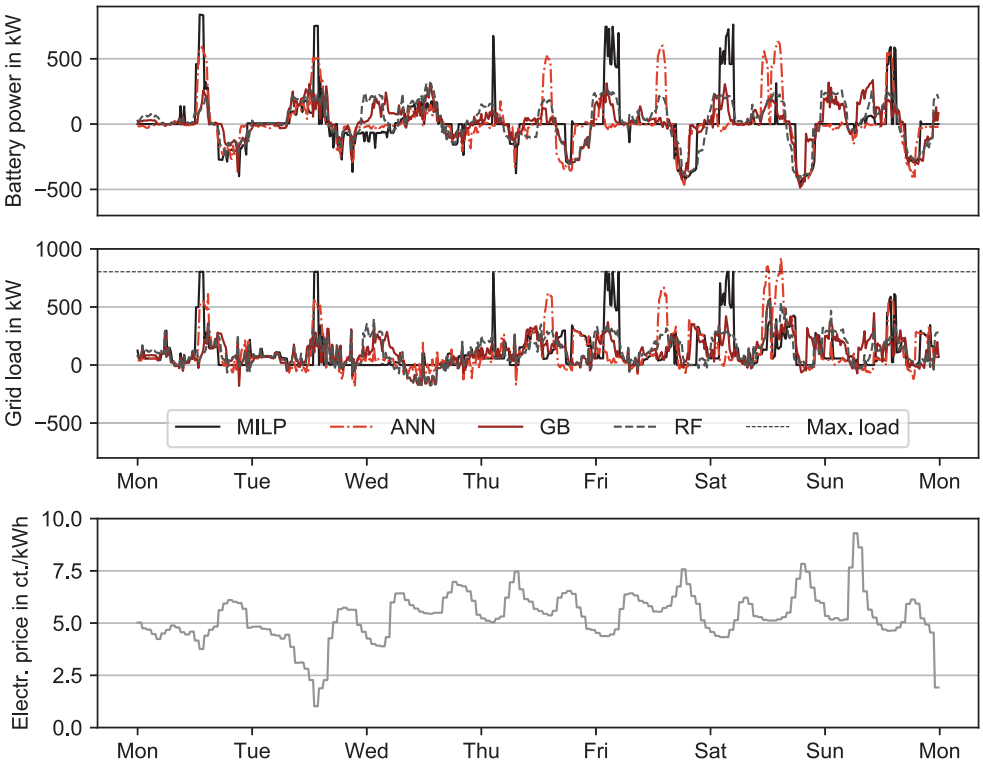

Figure 6: Closed-loop simulation results for the first week in June. The top plot shows the manipulated variables, i.e., the set battery power. In the center, the resulting grid load is depicted. Here, we also mark the optimized peak load as a reference. The bottom plot presents the EEX electricity tariff. The approximators control the system without the rule-based backup controller.

4. Discussion

The results of section 3. highlight that the relation between open- and closed-loop performance is difficult to assess. While a purely time-series-based comparison as shown in Figure 4 supports the conclusion that there is no correlation between open- and closed-loop performance, the subsequent analysis of the MPC-related objectives partially proves the opposite. When considering the annuity, the GB-based controller outperforms the ANN- and RF-based one (see Figure 5). Yet, the GB's resulting battery operation significantly differs from the MILP-based one, which is further supported by the low R^2 of 0.34 and the high *MAE* of 101 kW as illustrated in Figure 4. The ANN realizes the highest conformity with the teacher MPC's set points, which is apparent in Figure 6. Nonetheless, it tends to overshoot the optimized electrical peak due to the correlation with the electricity price. The other machine-learning-based approximators do not adopt this behavior as strongly. This is also why integrating the rule-based backup controller as shown on the right-hand side of Figure 5 has the most significant effect on the ANN-based controller's peak load.

Furthermore, it becomes evident that despite the integration of a rule-based backup controller that reduces the peak load, the optimized peak load of 803 kW is still surpassed by the approximators (see Figure 5). This behavior is explained by the rule-based controller's dependence on past measurements. Consequently, if the building load is higher than the load of the previous timestep, the battery's charging and discharging power in addition to the building load might still exceed the set upper limit. This disadvantage could be mitigated if the rule-based controller was based on building load predictions rather than past measurements. However, integrating these predictions would further increase the approximators' complexity. Still, we see great potential in developing a hybrid approach with a more sophisticated backup controller.

Another aspect is that the GB-based controller seems to slightly outperform its teacher MPC when taking the annuity as a basis (left plot in Figure 5). However, this behavior is only possible because the GB-based controller does not consider the MILP's constraints. A pre-optimization is carried out to obtain the optimal electricity peak. Subsequently, the MILP-based MPC uses this as an upper boundary and manipulates the

charging and discharging powers accordingly. Consequently, the GB extrapolates the defined solution space. Even though the GB-based controller does indeed achieve the lowest annuity, it is still not favorable compared to the MILP from an operator's point of view. It exceeds the optimal peak by approximately 100 kW and results in higher OPEX. Its low annuity mainly derives from low investment costs. The low investment costs are primarily caused by a longer battery lifetime due to reduced cycling and calendrical aging (see subsection 2.5.). When solely focusing on the operation, the ANN- and MILP-based controller are favored.

5. Conclusions

In the present study, we developed an approximate model predictive control approach for a PV-battery system in an office and laboratory building. The study is based on the preceding development of a hierarchical MPC in [27]. The hierarchy is based on prediction horizons so that different time scales can be considered. In the case of the PV-battery system, the MPC aims at optimizing the electrical peak load in a first step and incorporate it in the battery energy management in the second step (see subsection 2.2.). The hierarchical MPC serves as teacher to generate training data for machine learning models. These machine learning models are the approximators that replace the optimization problem after successful training. I.e., they function as controllers instead of the optimization program leading to a speed-up in processing time and lower hard- and software requirements in practice. Since the scientific literature is still missing a comparison of advanced machine learning models in the context of approximate MPC, we investigate an artificial neural network as well as the sophisticated tree-based models Random Forest and Gradient Boosting (see section 1.). In addition to that, we detect a research gap in in-depth analyses of the open- and closed-loop performance of the approximators. In the context of this study, the open-loop performance refers to the purely time-series based comparison of the true, i.e., the MPC output and the predicted, i.e., the approximators' outputs without any system interaction. The closed-loop performance is calculated based on the system interaction between the MPC or the approximators with the target system.

Our results prove that the machine-learning-based approximators all result in relatively high open-loop performances (see Figure 4). While the Gradient Boosting model slightly outperforms the other approximators open-loop, the same trend is not apparent when comparing the closed-loop performance solely based on the respectvie time-series. Here, we compare the MPC's output, i.e., the manipulated variables, with the approximator's one without evaluating the control-oriented metrics as defined in subsection 2.5.. Here, the Artificial Neural Network results in the highest accuracy metrics. When taking into consideration the PV-battery system's annuity, however, the Gradient Boosting model even surpasses the MPC results. The effect that the approximator outperforms its original teacher, is explained by the machine learning models' missing constraints. The MPC considers the maximum electrical peak as a constraint within the optimization. Hence, the MPC does not exceed the pre-optimized peak of 803 kW (see right-hand side of Figure 5). Nonetheless, the MPC performance is still considered the best when taking into account all relevant operation metrics. This behavior can be mitigated by an addtional rule-based backup controller that has been implemented in this study (see right-hand side of Figure 5). The integration of the backup controller results in a decrease in the electrical peak load while achieving almost the same annuity as the purely approximator-controlled system.

All in all, all three approximators result in a reliable closed-loop performance. Even though their control output clearly differs from the MPC's one, they outperform a benchmark model regarding the peak load. No clear relation between open- and closed-loop performance is apparent considering that the approximators all perform best regarding other metrics. Nonetheless, we detect that high open-loop performance supports a robust closed-loop performance. Finally, we go back to the stated research questions in subsection 1.1. as follows:

1. Advanced tree-based machine learning models like Gradient Boosting and Random Forest can outperform Artificial Neural Networks as controllers in an approximate MPC application. In addition, they are favourable from a comprehensibility point of view as they resemble the conventional "if-condition-then-action" structure of rule-based controllers.

2. The open-loop performance is an indicator for a good closed-loop performance, however, no generalizable conclusions can be drawn. A purely time-series-based comparison of controller outputs, i.e., the control trajectory, can be misleading regarding the closed-loop performance.

6. Outlook

The simulative assessment of the PV-battery system is based on a hierarchical MPC with a rolling horizon of 1 h. I.e., the control loop is repeated every hour with a timestep of 15 min and a prediction horizon of 16 h. The approximators, however, learn the in- to output relations based on every timestep, i.e., 15 min. This means that the approximators inheretly learn the model uncertainty of the MPC's process model for 3 out of 4 timesteps. Thus, we recommend future work to use a 1-step MPC to mitigate the model uncertainty. We expect the open-loop performance to increase for this case.

https://doi.org/10.52202/069564-0287

Apart from that, the overall profitability of the PV-battery system is not existent when taking the annuity into account. This is caused by the provided regulatory framework from Germany. Future work should focus on additional sources of revenue such as frequency control reserves. In addition, the interaction with the building energy system should be more closely investigated and synergy effects between the battery and the HVAC system identified.

Acknowledgment

We gratefully acknowledge the financial support by Federal Ministry for Economic Affairs and Climate Action (BMWK), promotional reference 03EN3026C.

References

[1] 2022 Global status report for buildings and construction: Towards a zero-emissions, efficient and resilient buildings and construction sector. Nairobi; 2022.

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

[3] Oldewurtel F, Parisio A, Jones CN, Gyalistras D, Gwerder M, Stauch V, et al. Use of model predictive control and weather forecasts for energy efficient building climate control. Energy and Buildings. 2012 Feb;45:15-27. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0378778811004105.

[4] Ma Y, Borrelli F, Hencey B, Coffey B, Bengea S, Haves P. Model Predictive Control for the Operation of Building Cooling Systems. IEEE Transactions on Control Systems Technology. 2012;20(3):796-803.

[5] Sturzenegger D, Gyalistras D, Morari M, Smith RS. Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost–Benefit Analysis. IEEE Transactions on Control Systems Technology. 2016 Jan;24(1):1-12. Available from: http://ieeexplore.ieee.org/document/7087366/.

[6] Širokỳ J, Oldewurtel F, Cigler J, Prívara S. Experimental analysis of model predictive control for an energy efficient building heating system. Applied energy. 2011;88(9):3079-87.

[7] Cigler J, Gyalistras D, Široký J, Tiet VN, Ferkl L. Beyond theory: the challenge of implementing model predictive control in buildings. Proceedings of 11th Rehva world congress, Clima. 2013;(250).

[8] Prívara S, Cigler J, Váňa Z, Oldewurtel F, Sagerschnig C, Žáčeková E. Building modeling as a crucial part for building predictive control. Energy and Buildings. 2013;56:8-22.

[9] May-Ostendorp PT, Henze GP, Rajagopalan B, Corbin CD. Extraction of supervisory building control rules from model predictive control of windows in a mixed mode building. Journal of Building Performance Simulation. 2013;6(3):199-219.

[10] Domahidi A, Ullmann F, Morari M, Jones CN. Learning decision rules for energy efficient building control. Journal of Process Control. 2014;24(6):763-72.

[11] Drgoňa J, Picard D, Kvasnica M, Helsen L. Approximate model predictive building control via machine learning. Applied Energy. 2018;218:199-216.

[12] May-Ostendorp P, Henze GP, Corbin CD, Rajagopalan B, Felsmann C. Model-predictive control of mixed-mode buildings with rule extraction. Building and Environment. 2011;46(2):428-37.

[13] Klaučo M, Drgoňa J, Kvasnica M, Di Cairano S. Building Temperature Control by Simple MPC-like Feedback Laws Learned from Closed-Loop Data. IFAC Proceedings Volumes. 2014;47(3):581-6.

[14] May-Ostendorp PT, Henze GP, Rajagopalan B, Kalz D. Experimental investigation of model predictive control-based rules for a radiantly cooled office. 2013;19(5):15.

[15] Bursill J, O'Brien L, Beausoleil-Morrison I. Multi-zone field study of rule extraction control to simplify implementation of predictive control to reduce building energy use. Energy and Buildings. 2020 Sep;222:110056. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0378778819335078.

[16] Piscitelli MS, Brandi S, Gennaro G, Capozzoli A, Favoino F, Serra V. Advanced Control Strategies For The Modulation Of Solar Radiation In Buildings: MPC-Enhanced Rule-Based Control. In: Corrado V, Fabrizio E, Gasparella A, Patuzzi F, editors. Proceedings of Building Simulation 2019: 16th Conference of IBPSA. Building Simulation Conference proceedings. IBPSA; 2020. p. 869-76.

[17] Drgoňa J, Picard D, Helsen L. Cloud-based implementation of white-box model predictive control for a GEOTABS office building: A field test demonstration. Journal of Process Control. 2020;88:63-77.

[18] Le K, Bourdais R, Guéguen H. From hybrid model predictive control to logical control for shading system: A support vector machine approach. Energy and Buildings. 2014;84:352-9.

[19] Yang S, Wan MP, Chen W, Ng BF, Dubey S. Experiment study of machine-learning-based approximate model predictive control for energy-efficient building control. Applied Energy. 2021;288:116648.

[20] Lohr Y, Monnigmann M, Klauco M, Kaluz M. Mimicking Predictive Control with Neural Networks in Domestic Heating Systems. In: 2019 22nd International Conference on Process Control (PC19). IEEE; 2019. .

[21] Kuhn M, Johnson K. Applied Predictive Modeling. New York, NY: Springer New York; 2013. Available from: http://link.springer.com/10.1007/978-1-4614-6849-3.

[22] Owerko D, Gama F, Ribeiro A. Optimal power flow using graph neural networks. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE; 2020. p. 5930-4.

[23] Zamzam AS, Baker K. Learning optimal solutions for extremely fast AC optimal power flow. In: 2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm). IEEE; 2020. p. 1-6.

[24] Canyasse R, Dalal G, Mannor S. Supervised learning for optimal power flow as a real-time proxy. In: 2017 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT). IEEE; 2017. p. 1-5.

[25] Pan X, Zhao T, Chen M, Zhang S. Deepopf: A deep neural network approach for security-constrained dc optimal power flow. IEEE Transactions on Power Systems. 2020;36(3):1725-35.

[26] de Jongh S, Steinle S, Hlawatsch A, Mueller F, Suriyah M, Leibfried T. Neural Predictive Control for the Optimization of Smart Grid Flexibility Schedules. In: 2021 56th International Universities Power Engineering Conference (UPEC). IEEE; 2021. .

[27] Maier LM, Kühn L, Mehrfeld P, Müller D. Time-based economic hierarchical model predictive control of all-electric energy systems in non-residential buildings. RWTH Aachen University;.

[28] Henn S, Richarz J, Maier L, Ying X, Osterhage T, Mehrfeld P, et al. Influences of usage intensity and weather on optimal building energy system design with multiple storage options. Energy and Buildings. 2022;270:112222.

[29] Hart WE, Watson JP, Woodruff DL. Pyomo: modeling and solving mathematical programs in Python. Mathematical Programming Computation. 2011;3(3):219-60. Available from: https://link.springer.com/article/10.1007/s12532-011-0026-8.

[30] Deutscher Wetterdienst. Wetter und Klima - Deutscher Wetterdienst - Leistungen - Testreferenzjahre (TRY); 18.05.2022. Available from: https://www.dwd.de/DE/leistungen/testreferenzjahre/testreferenzjahre.html.

[31] SMARD — Marktdaten visualisieren; 19.05.2022. Available from: https://www.smard.de/home/marktdaten/?marketDataAttributes=%7B%22resolution%22:%22year%22,%22from%22:1514674800000,%22to%22:1577833199999,%22moduleIds%22:%5B8004169%5D,%22selectedCategory%22:null,%22activeChart%22:true,%22style%22:%22color%22,%22categoriesModuleOrder%22:%7B%7D,%22region%22:%22DE%22%7D.

[32] Sandels C, Brodén D, Widén J, Nordström L, Andersson E. Modeling office building consumer load with a combined physical and behavioral approach: Simulation and validation. Applied Energy. 2016;162:472-85.

[33] Pedregosa F, Varoquaux G, Gramfort A, Vincent M, Thirion B. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research. 2011;12:2825-30. Available from: https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf?ref=https://githubhelp.com.

[34] Rätz M, Javadi AP, Baranski M, Finkbeiner K, Müller D. Automated data-driven modeling of building energy systems via machine learning algorithms. Energy and Buildings. 2019;202:109384.

https://doi.org/10.52202/069564-0287