Tailoring the Heat Pump System Controller to the Building: Assessment of Adapted Heating Curves using Data-Driven Methods

Florian Will[1]*, Jonas Klingebiel[1], Christian Vering[1], Dirk Müller[1]

[1]RWTH Aachen, Institute for Energy Efficient Buildings and Indoor Climate, Aachen, Germany

Contact Information: florian.will@eonerc.rwth-aachen.de

* Corresponding Author

ABSTRACT

Space heating accounts for a large share of energy consumption in Germany due to fossil fuel combustion. To reduce combustion-related emissions, space heating electrification with heat pumps is promising. However, heat pumps increase the power demand related to the electrical grid. Therefore, minimizing the additional electrical power consumed by heat pumps is crucial, which can be achieved by either lowering the heating demand or enhancing the overall efficiency. While lowering the heating demand due to improvements to the building envelope is very cost-intensive, increasing the heat pump efficiency is substantially more cost-effective.

To achieve high efficiencies with heat pumps, low supply temperatures are required. Today, static heating curves (HCs) connect the heat pump with a specific building. Conventionally, the HC correlates the ambient temperature with the supply temperature. This method captures the fundamental principle that colder ambient conditions necessitate higher supply temperatures. While HCs are favored for their simplicity in implementation, inefficiencies arise as parameters and form of HC are usually chosen based on simplified heuristics which cannot capture building specifications and user behavior. Hence, a method to tailor the HC to the building and its usage is promising to increase efficiency while maintaining thermal comfort.

This paper introduces a method that optimizes the parameters and form of the HC by using data-driven models. Therefore, relevant data is collected in a system identification process using a conventional HC. The data is used to generate data-driven models that autonomously detect the underlying patterns between ambient conditions and supply temperature. From these models, a self-optimized HC (soHC) is derived using the DICE algorithm. We assess the performance of the soHC with different heating systems in annual building energy system performance simulations. The soHC outperforms the conventional HC in terms of efficiency with a decrease in energy usage of up to 7.3 % while maintaining thermal comfort. In future research, enhancements can be achieved by incorporating specific conditions like solar irradiation and presence detection into the approach.

1. INTRODUCTION

Heat pumps are a crucial technology to reach climate goals and slow down climate change. They need to be more (cost)-efficient than conventional heating technologies like gas boilers to be used more frequently. To achieve this, the control of the heat pump must be improved. One way to improve the control of heat pumps is to find the optimal supply temperature for the heat pump, that is as low as possible and as high as needed to heat the building. The optimal supply temperature depends on various factors. These factors are for example the heat losses of the building at different air temperatures, the type of the heating element, and disturbances like internal gains or solar radiation. Conventionally, the supply temperature is set by the heating curve (HC). The HC sets the supply temperature based on the air temperature. There are problems with conventional HCs that cause non-optimal supply temperatures. There are cases where the HC is horizontal. This independence of the air temperature causes too high flow temperatures. Also, in conventional heat pump controllers, expert knowledge is needed to set the HC accordingly (Vaillant, 2024), (Viessmann, 2024). There are rules how to set the HC, however, to apply them insight in the building energy system is needed as well as knowledge of the local controller of the heat pump. The needed expert-knowledge is a hurdle for heating engineers and inhabitants, which often leads to a HC that is set too high. A HC which is set too high ensures thermal comfort but leads to an inefficient operation of the heat pump. So, to increase the efficiency of heat pumps, the HC needs to be adapted to the building. To integrate this adaption process into practice, it needs to be done automatically to decrease the hurdles for heating engineers and inhabitants.

An automatic adaption strategy is presented by Potočnik and Govekar (2019), who change two grid points of a linear HC in their paper. The algorithm uses the error between the set and actual room temperature to adapt the HC to the heat demand. They apply the algorithm on two buildings and show low thermal discomfort with the adaptive strategy. However, they do not consider multiple zones or a thermostatic valve for the heater.
Ionesi et al. 2015) look at a building energy system (BES) with four rooms, each with a radiator and thermostatic valve. Their adaptive strategy is similar to the one by Potočnik and Govekar (2019) as they adapt the heating curve with up to 62 grid points for the error between set and actual room temperature. They achieve a reduction in energy consumption by 5 % compared to a baseline study while maintaining thermal comfort. In their studies, they neglect solar radiation in the simulation model, which simplifies the problem significantly.
Huchtemann and Müller (2013) investigate a BES with four rooms, each with a radiator and valves, a buffer tank, and a heat pump. They adapt the set supply temperature based on the position of the valves and can save 6.8 % primary energy compared to a standard control. However, their control algorithm needs to have information about the valve positions. Also, they only look at an on/off controlled heat pump.
There are also more complex control strategies for BES. Namely, there is reinforcement learning (RL) and model predictive control (MPC). MPC uses predictions of the system dynamics that are outputs of process models of the system to calculate optimal set values. The process models can be anything between white box and black box models. MPC is applied successfully to BES in multiple studies like (Baumann et al., 2023), (Göbel, Stoffel, et al., 2023) or (Kuboth et al., 2020). The studies demonstrate reduced electric energy consumption between 4.1 % and 20.3 % compared to benchmark controllers. MPC, however, needs more computational resources than a HC, because optimization needs to be carried out regularly. Another hurdle for using MPC in praxis is the modeling effort needed to create the process models (Stoffel et al., 2023).
RL-based control strategies do not need to carry out an optimization to get the set values. During the training of the RL-Agent, who makes the decisions about the set values, a model is generated that finds the optimal set values based on the state (can be past and/or future) of the BES. The optimal set values are indirectly defined by a reward function, that often includes thermal comfort and energy efficiency. RL is applied on BES in for example these studies: (Huang et al., 2022; Rohrer et al., 2023), a comparison to a benchmark controller however is not carried out in the two studies. A shortcoming of RL is that it needs a lot of training time, during which thermal comfort is not guaranteed (Shaqour & Hagishima, 2022).
RL and MPC can both use machine-learning methods. Machine-learning methods are suitable for control applications in BES because of the increasing sensor data availability (Stoffel et al., 2024), that can be used as training data for the machine learning models. Due to the high individuality of buildings, data-based models are a way to reduce the modeling expense in comparison to physical models.
In this paper, we use a method for a more efficient operation of BES that uses data-driven modeling to adapt the heating curve of a building. The method is less complex and less computationally intensive than RL or MPC but uses the advantages of machine learning-based control of BES that are described before.
In **section 2**, the method to generate the HC from the data-driven models will be described as well as some necessary foundations of machine learning methods. In **section 3**, the case study of this paper will be presented. **Section 4** presents a simulative comparison of the soHCs with a benchmark HC. The results will be discussed, **section 5** gives a conclusion.

2. METHOD

In this section, the methods used in this paper will be described. The method is depicted in Figure 1. First, the simulation models on which the studies will be carried out are shown. After that, an algorithm for systematic system identification (1) that is used to generate data for the data-driven models is described. In the next subsection, the data-driven modeling methods of this paper are shown (2). The final subsection describes the algorithm to extract the HC from the data-driven models (3).

2.1 Simulation Modeling

This paper carries out a simulative study. The simulation takes place in Modelica and uses the open-source library BESMod (Wüllhorst et al., 2022). The BES with the control components is shown in Figure 2. The building is generated with TEASER, an open-source library for building performance simulations (Remmen et al., 2018). The building model has a nominal heating load of 11 kW at -12 °C outdoor temperature and consists of one zone that is heated by either an underfloor heating (UFH) or a radiator. The heat is provided by a heat pump. The corresponding

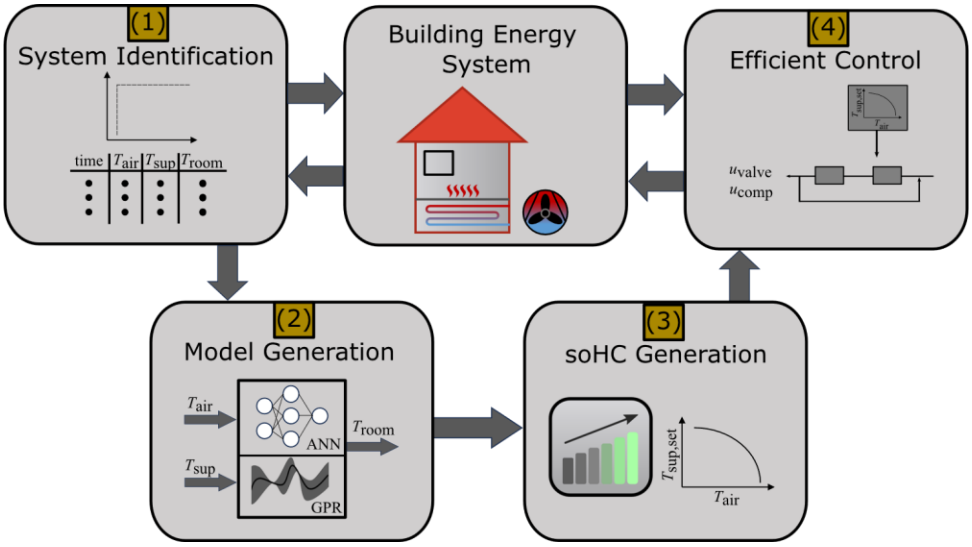

Figure 1: The method to generate soHCs used in this paper. First, system identification is carried out (1). Then, data-driven models of the system are created (2) and a soHC is generated with these data-driven models (3). The soHC is used for efficient control of the BES (4).

heat pump model is calibrated and validated on real heat pump data and described in (Göbel, Waiz, et al., 2023). It has a maximal condenser capacity of 13.6 kW.

The control of the BES is carried out by a HC and two PI-controllers. The HC defines a $T_{\mathrm{sup,set}}$ based on T_{air}. The PI-controller sets the compressor speed u_{comp} of the heat pump based on the difference between $T_{\mathrm{sup,set}}$ and T_{sup}. With u_{comp}, the condenser heat is set. In the other control loop, the PI-controller sets the valve opening u_{valve} based on $T_{\mathrm{room,set}}$ and T_{room}. With the valve opening, the mass flow to the building is controlled to account for disturbances like internal gains or solar radiation, that are not accounted for in the heating curve.

2.2 System Identification

To generate data for the data-driven models, system identification needs to be done. The data should contain as much information as possible, so in this context different combinations of T_{sup} and T_{air}. If a system with a fixed HC is used for system identification, there will only be the combinations of T_{sup} and T_{air} that the HC specifies. We use an algorithm for system identification that uses different HCs to increase the amount of information in the data. The algorithm is explained in the following.

The system identification is carried out with the valve completely opened to eliminate this additional disturbance. First, the BES is simulated with 2 linear HCs (eq. 1) with different slopes (k) over 4 days each.

$$T_{\mathrm{sup,set}} = k \cdot (20\ °\mathrm{C} - T_{\mathrm{air}}) + 20\ °\mathrm{C} \tag{1}$$

The room temperature deviation ($T_{\mathrm{room,dev}} = T_{\mathrm{room,set}} - T_{\mathrm{room}}$) is measured and then averaged over 8 hours, to account for the thermal inertia of the system. The averaged $T_{\mathrm{room,dev}}$ are assigned to the corresponding k. We then do a linear regression with $T_{\mathrm{room,dev}}$ as y-value and k as x-value. In the obtained linear equation, we set $T_{\mathrm{room,dev}} = 0$ and calculate k. This k is used as the slope for the next HC, which runs for another 7 days. The data generated with the new HC we then use for another linear regression and to calculate the next k. The process is repeated 6 times, so that there is a total of 50 training days.

With this algorithm, many different combinations of T_{sup} and T_{air} are present in the training data, because HCs with different slopes are used. Also, the algorithm converges to a reasonably fitted k, so that the thermal discomfort for the inhabitants is kept low.

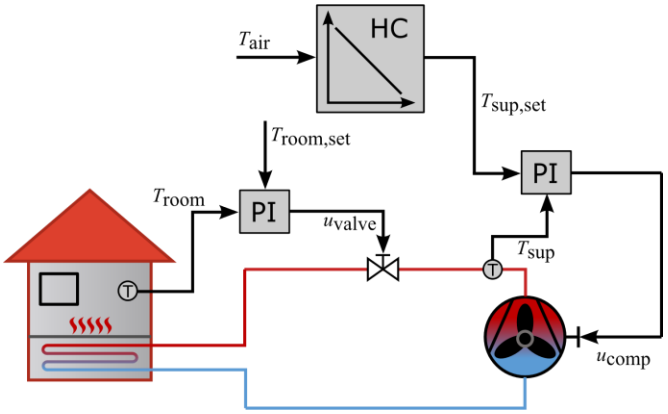

Figure 2: BES with control setup. There are two PI-Controllers. One controls u_{valve} to influence the mass flow into the heating system and therefore T_{room}. The other controls u_{comp} to get a T_{sup} according to the HC.

2.3 Data-Driven Modeling

In this section, the models that are later used to generate the soHCs are described. We use 2 different machine-learning model architectures in this paper, artificial neural networks (ANNs) and gaussian process regressors (GPRs). These 2 architectures are chosen, because they are both well-suited for regression tasks. The main difference between ANNs and GPRs lies in the model structure and the training process. ANNs have a fixed number of hyperparameters (once the number of layers and neurons per layer are chosen), that are optimized during the training. The computational complexity of ANNs doesn't grow with the number of training data points. With GPRs, the computational complexity increases cubically with the number of training data points (Wang, 2022). One advantage of GPRs compared to ANNs is, that they provide the variance of an output along with the value. For a more detailed (mathematical) description of ANNs and GPRs, read here: (Nielsen, 2015; Wang, 2022). For training of the ANNs we use TensorFlow (TensorFlow Developers, 2024), for generating the GPRs we use Scikit-learn (Pedregosa et al., 2011).
The inputs for the models are T_{sup} and T_{air}, the output is $T_{room,dev}$. They are averaged over 8 hours before they are used as inputs for the models to account for the thermal inertia of the system. The hyperparameters for the ANN and GPR are determined iteratively and based on experience. The ANN has 2 hidden layers with 6 neurons each. The GPR has a radial basis function kernel that is scaled by a constant kernel. It is also added a white kernel to account for noise in the data. Another important hyperparameter is the length scale of the radial basis kernel. If the length scale is too low, the model is not smooth enough to generate a HC. So, the minimal bound of the length scale is set to be 8, which is an iteratively determined value. The training data for the models consists of 50 days, generated as described in the section before.

2.4 Algorithm to extract the HC

From the step above we have obtained a model that is able to predict $T_{room,dev}$ based on T_{sup} and T_{air}. In this step, these models are used to generate the soHC. The basic principle is as follows: The model gets a fixed T_{air} as an input and looks for a T_{sup} so that $T_{room,dev} = 0$ for T_{air} in range [-8 °C, 15 °C] in steps of 1 °C. The T_{air} are the grid points of the soHC. After a T_{sup} is found for each T_{air}, the points are linearly interpolated to get a piecewise linear function, that is our soHC.
To extract the soHC from the model, the "Diverse Counterfactual Explanations" (DICE) (Mothilal et al., 2019) method is used. DICE is used to generate post-hoc explanations of machine learning models. For that, it looks for model inputs to get a desired model output. In our case, it searches for a T_{sup} with a fixed T_{air} so that $T_{room,dev} = 0$. DICE solves an optimization problem to generate these counterfactuals. The algorithm is parametrized as follows: The permitted range for T_{sup} is [15 °C, 60 °C], the desired range for $T_{room,dev}$ is [0.15 K, 0.25 K] and we set up the genetic algorithm of DICE.

3. CASE STUDY

A case study is carried out to evaluate the soHCs. For that, the simulation model described in subsection 2.1 is used. The evaluated key performance indicators (KPIs) are described in subsection 3.1. The benchmark controller is described in subsection 3.2 and a detailed look of the generated soHCs is given in subsection 3.3.

3.1 KPIs

To evaluate the performance of the soHCs, two KPIs are defined. The **thermal discomfort (*TD*)** is a measure for the following of temperature comfort bounds. In this paper, every T_{room} below 20 °C is considered as thermal discomfort. Because we have discrete timeseries data, *TD* is defined as a sum:

$$TD = \sum_{i=1}^{n} d_i \cdot I \cdot (20\,°\mathrm{C} - T_{\mathrm{room}}) \tag{2}$$

Where *n* is the number of datapoints in the timeseries, d_i is 1 if ($T_{\mathrm{room}} < 20$ °C) and else 0, and *I* is the time interval between two steps that are evaluated.
Another KPI is the efficiency of the heat pump. For that, the *SCOP* is used as a measure. It is defined as follows:

$$SCOP = \frac{Q}{W_{\mathrm{el}}} \tag{3}$$

Where Q is the generated heat by the heat pump and W_{el} is the used electric energy.

3.2 Benchmark Controller

As a benchmark to compare the soHCs to, linear HCs are used. The formula for the linear HC is described above in equation 1. The linear form of the HC with the footing point at (20 °C, 20 °C) is based on (Ruhnau et al., 2019). The slope *k* of the linear HC is varied between [0.4, 0.75] by steps of 0.05 for the UFH system and [1.1, 1.8] by steps of 0.1 for the radiator system. The slope needs to be higher for the radiator system because higher T_{sup} are needed. For each soHC and each linear HC, yearly simulations are carried out. In the yearly simulations, both PI-controllers that control u_{valve} and u_{comp} are active.

3.3 Generated Heating Curves

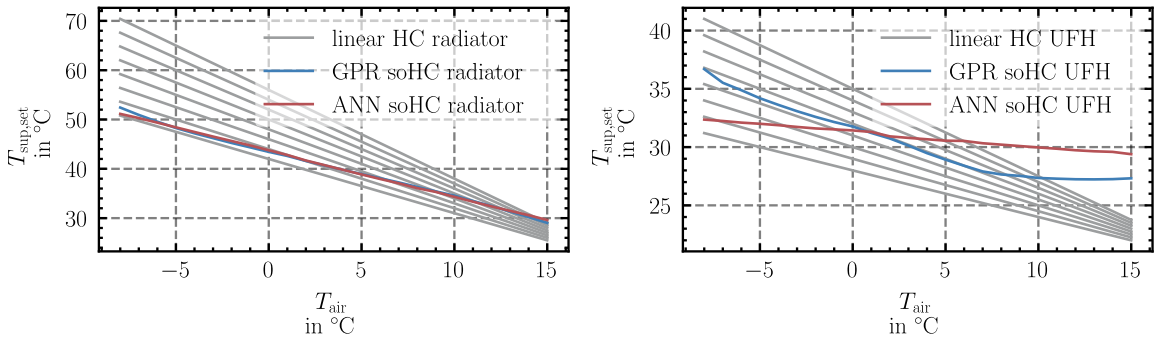

Figure 3: The soHCs for the radiator system and the UFH system, generated from ANN and GPR. For comparison the benchmark linear HCs are shown.

Figure 3 shows the soHCs from both data-driven models for the radiator and UFH system compared to the benchmark linear HCs.
For the radiator system, the soHCs generated with GPR or ANN are similar. The shape of both soHCs are nearly linear, but the footing point is set higher compared to the linear HCs and the slope is lower, which leads to higher $T_{\mathrm{sup,set}}$ at low T_{air} and lower $T_{\mathrm{sup,set}}$ at high T_{air}. For the UFH system, the soHCs differ between the models. It is difficult to explain the reason for the different outcomes because the models are black-boxes. One reason could be the higher thermal inertia of the UFH system. The high thermal inertia lowers the quality of the training data, because the response to a change in $T_{\mathrm{sup,set}}$ takes a lot of time to be seen in T_{room}.

All in all, the data-driven models generate senseful outputs. The ANN predicts a linear dependency between $T_{\mathrm{sup,set}}$ and T_{air} with a $T_{\mathrm{sup,set}} = 29\ °C$ at $T_{\mathrm{air}} = 15\ °C$ and $T_{\mathrm{sup,set}} = 32\ °C$ at $T_{\mathrm{air}} = -8\ °C$. The soHC generated from the GPR is more like the linear HCs, but it has a lower slope starting from $T_{\mathrm{air}} = 7\ °C$.
In the next section, these soHCs are compared to the linear HCs in the previously described case study.

4. RESULTS

In this section, the soHCs will be analyzed with yearly simulations regarding the KPIs and compared to the linear HCs. Subsection 4.2 will give a deeper insight into the operation of the heat pump and the BES. In subsection 4.3, the results are discussed.

4.1 Controller Comparison

We first evaluate the yearly simulations with Figure 4. It shows the thermal discomfort in dependence on the *SCOP* of the heat pump, for the radiator system on the left and the UFH system on the right. The soHCs are compared to linear HCs with a slope of 1.1 to 1.8 for the radiator system and a slope of 0.4 to 0.75 for the UFH system. A spline regression between the points of the linear HCs is performed to make the differences between the soHCs and the linear HCs more visible. The area with the soHCs is zoomed in to have a detailed look at the *SCOP* and thermal discomfort for the soHCs. It is to be noted that in this section, the soHCs are compared to the best possible linear HC on the regression curve, so the benchmark controller is already well fitted to the BES.

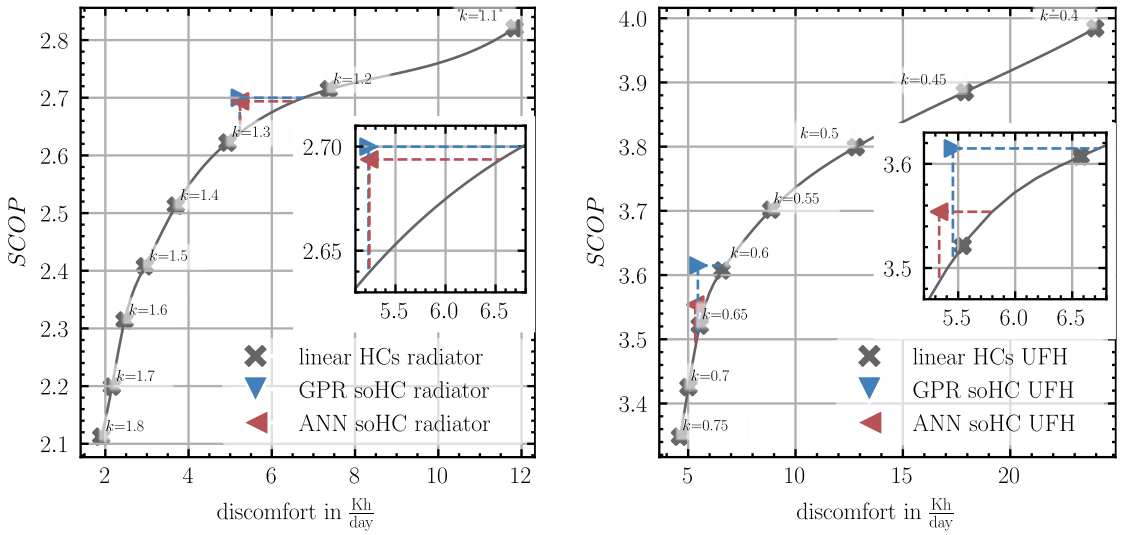

Figure 4: *SCOP* and thermal discomfort for the benchmark HCs (grey), the soHC from the GPR Model (blue) and the ANN Model (red) for the radiator and UFH system.

First, we analyze the radiator system with the left figure. For the linear HCs, the *SCOP* increases with decreasing slope from 2.1 for a slope of 1.8 to a *SCOP* of 2.8 for a slope of 1.1, which is due to the increasing $T_{\mathrm{sup}}$with increasing slope. The heat pump operates less efficiently with increasing temperature difference between T_{air} and T_{sup}The thermal discomfort increases from 2.0 $\frac{\mathrm{Kh}}{\mathrm{day}}$ to 12.0 $\frac{\mathrm{Kh}}{\mathrm{day}}$ with decreasing slope because the buildings heat demand can't be supplied with low values for T_{sup} . Both soHCs have very similar values in *SCOP* and thermal discomfort, which is due to their similar form. Their tuples of thermal discomfort and *SCOP* are to the left of the regression curve, which means a higher *SCOP* and lower thermal discomfort. Next, the KPIs of the soHCs are compared to values on the regresiion curves. Compared to the regression curve, the soHC generated from the GPR increases the *SCOP* by 2.3 % for a discomfort of 5.5 $\frac{\mathrm{Kh}}{\mathrm{day}}$ and decreases the thermal discomfort by 22.7 % for an *SCOP* of 2.7.
For the UFH system, the absolute *SCOP*s are higher compared to the radiator system because of the lower T_{sup}. However, the tendencies are the same as for the radiator system with increasing *SCOP* and increasing thermal

discomfort for increasing slopes of the linear HCs. The thermal discomfort is for both soHCs at around 5.0 to 5.5 $\frac{\text{Kh}}{\text{day}}$, which is like the thermal discomfort values of the soHCs of the radiator system. In the zoomed-in part of the figure, it can be seen that the soHC generated from the GPR reaches an increase of 2.8 % of the *SCOP* at a thermal discomfort of 5.5 $\frac{\text{Kh}}{\text{day}}$ and a decrease of 19.0 % in thermal discomfort at a *SCOP* of 3.6 compared to the linear HCs. The soHC generated from the ANN increases the *SCOP* by 1.7 % at a thermal discomfort of 5.3 $\frac{\text{Kh}}{\text{day}}$ and decreases the thermal discomfort by 8.1 % at a *SCOP* of 3.55 compared to the linear HCs.

4.2 Analysis of the Operation

In the following, the soHCs are compared in more detail to the linear HC with the next lowest thermal discomfort to the soHCs. This is the HC with a slope of 1.3 for the radiator system and the linear HC with a slope of 0.7 for the UFH system. These HCs are chosen as a comparison because in reality HCs are set to high most of the time to ensure low thermal discomfort (Ionesi et al., 2015). For the soHCs, the GPR model is chosen because it performed the best in the KPIs thermal discomfort and *SCOP*.

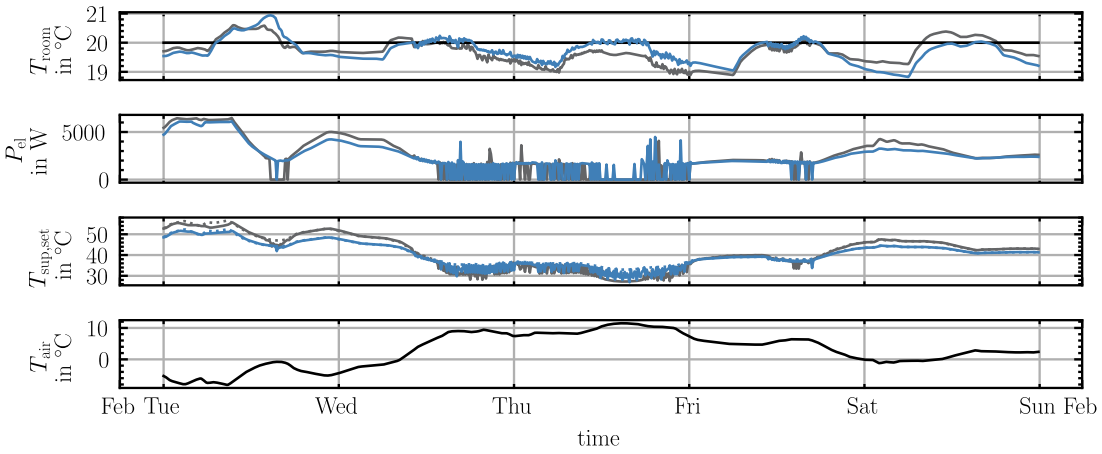

Figure 5: Timeseries data of T_{room}, P_{el}, $T_{\text{sup,set}}$ and T_{air} of a linear HC (grey) and the soHC (blue) for the radiator system over 5 days in February.

For a detailed analysis of the operation, Figure 5 shows the timeseries data of relevant variables for 5 exemplary days in February for the radiator system. They are chosen because of their wide range in T_{air} from -8 °C to 10 °C. In blue, the soHC is shown, in grey the linear HC.

T_{room} fluctuates for both HCs because of the disturbances on the building like internal gains and solar radiation. Due to the thermal inertia of the system, the valve cannot compensate these disturbances. However, at $T_{\text{air}} < 0$ °C, T_{room} fluctuates around the set temperature of 20 °C, both HCs are well set. For T_{air} around 10 °C, T_{room} of the linear HC is up to 1.7 K below the set temperature, because the linear HC is not well set at these T_{air}. The soHC however has higher $T_{\text{sup,set}}$ at these T_{air}, so the deviation from 20 °C is lower.

Figure 6 gives a detailed insight into the operation hours in the yearly simulation of T_{sup} for the two systems and the HCs and soHCs. In the left figure, it can be seen that for the linear HC, the heat pump operates T_{sup} less hours in the range around 40 °C. However, high T_{sup} above 48 °C are much more present for the linear HC. This is due to the higher $T_{\text{sup,set}}$ of the linear HC for low T_{air}.

Now the operation hours of the heat pump for the UFH system, seen in Figure 6 on the right, are analyzed. The ANN soHC almost only operates in T_{sup} range from 30 °C to 32 °C due to the low slope of the ANN soHC. For the GPR soHC and the linear HC, similar tendencies are present as for the radiator system. The heat pump with the GPR soHC operates more hours in the range $T_{\text{sup}} < 34$ °C. With the linear HC, the heat pump operates more hours in the range $T_{\text{sup}} > 34$ °C, due to the higher $T_{\text{sup,set}}$ in for low T_{air}.

The higher T_{sup} at low T_{air} for the linear HCs for both systems lead to a less efficient operation of the heat pump. The thermal discomfort however is not lower compared to the soHCs, because $T_{\text{sup,set}}$ is above the needed T_{sup} to compensate the heat losses of the building.

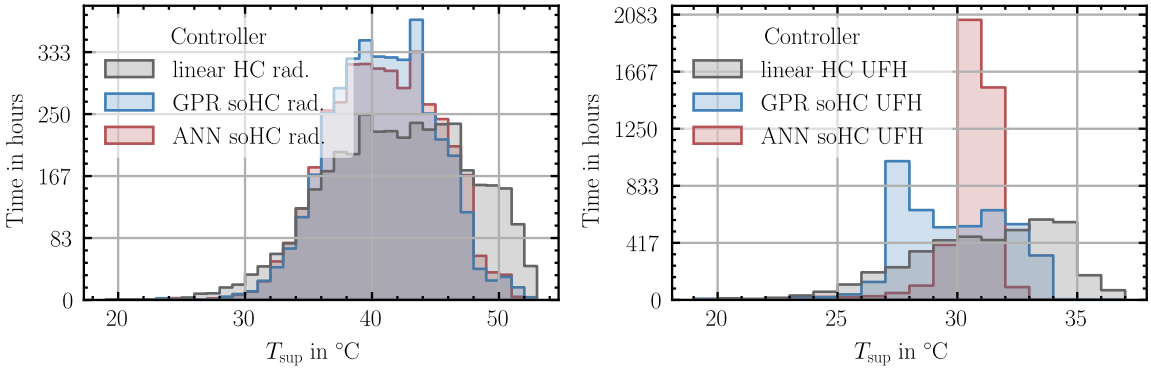

Figure 6: Histograms that show the operation hours of T_{sup} for the linear HC (grey) and the soHCs (blue for GPR and red for ANN) for the radiator (left) and the UFH (right) system. Only operation points where the heat pump is on depicted.

Table 1 shows the KPIs for a comparison between the linear HC and the GPR soHC for both systems. The electric energy W_{el} is reduced by 5.3 % for the radiator system and by 7.7 % for the UFH system. Further, the thermal discomfort is increased for both systems compared to the chosen linear HCs. The average T_{sup} is decreased by around 1.4 K with both soHCs, which leads to an increase in *SCOP*.

Table 1: KPIs to compare the linear HC with the soHC for the radiator and the UFH system.

System	**Radiator**		**UFH**	
HC	**linear ($k = 1.3$)**	**soHC (GPR)**	**linear($k = 0.7$)**	**soHC (GPR)**
W_{el} in kWh	10368	9815	11508	10621
TD in $\frac{Kh}{day}$	4.9	5.2	5.0	5.5
average T_{sup} in °C	42.3	40.9	31.1	29.8
SCOP	2.62	2.7	3.43	3.61

4.3 Discussion of the Results

The results demonstrate that the method to find optimized HCs for different BES by using data-driven models is successful. Due to the adapted supply temperatures, a more efficient operation and less thermal discomfort can be achieved.

Using the defined setup, there are limitations for discussion:

- **Building Energy System:** The model of the BES consists of one zone, which simplifies the system. In future work, method should be used for a BES with multiple zones.
- **System Identification:** To create the data-driven models, system identification over multiple weeks must be carried out, where thermal comfort is not guaranteed. The system identification process could be made more efficient, so that less training data is needed.
- **Disturbances:** Disturbances like solar radiation, internal gain or thermal inertia of the building influence the BES. The data-driven models in this paper however only consider the ambient air temperature as a disturbance. The low number of inputs leads to inaccuracies in the data-driven models and to therefore inaccurate soHCs. In future work, more disturbances should be used as inputs to the data-driven models. This way, a more dimensional soHC will be created, which leads to more accurate supply temperatures.
- **Benchmark:** The benchmark controller in this paper is a group of linear HCs. The soHC is then compared to the best fitted linear HC, which is an unrealistically good benchmark that probably won't exist in the real world. In future work, the benchmark HC should be chosen based on field test data.

5. CONCLUSIONS

The paper shows the potential of the self-optimized heating curves compared to benchmark heating curves. A method is developed, where data from a system identification process is used to generate data-driven models. From the data-driven model, the self-optimized heating curve is generated using the DICE algorithm.
The method is demonstrated on two building energy system simulations with a radiator and an under-floor heating system, where differently high T_{sup} are needed. The self-optimized heating curves not only adapted to these different levels but also to the heat losses of the building at different T_{air}. Compared to a well-fitted linear heating curve, an improvement in *SCOP* of up to 2.3 % for the radiator system and up to 2.8 % for the under-floor heating system at the same thermal discomfort could be reached. An improvement in thermal discomfort at a fixed *SCOP* of 22.7 % and 19 % for the two systems could be made. Compared to a more realistically used linear heating curve, that is designed more conservatively to ensure thermal comfort, energy savings of 5.3 % for the radiator system and 7.7 % for the under-floor heating system can be achieved.
In conclusion, the demonstrated method to generated self-optimized heating curves is successfully demonstrated on two different building energy systems. However, in future work the building energy system simulation model should be made more complex to include multiple zones. Also, the data-driven models should include more disturbances to generated more accurate heating curves which leads to a more efficient operation. To exploit the entire potential of using data-driven models to generate heating curves, the next steps are increasing the complexity of the data-driven models.

NOMENCLATURE

Subscript

air	air
dev	deviation
el	electric
i	index for sum
n	number of datapoints
room	room
set	set value
sup	supply

REFERENCES

Andricciola, A. (2018). *Development of an algorithm for the automatic adjustment of the heating curve of a heat pump heating system*. https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-240070

Baumann, C., Huber, G., Alavanja, J., Preißinger, M., & Kepplinger, P. (2023). Experimental validation of a state-of-the-art model predictive control approach for demand side management with a hot water heat pump. *Energy and Buildings*, *285*, 112923. https://doi.org/10.1016/j.enbuild.2023.112923

Göbel, S., Stoffel, P., Will, F., Vering, C., & Müller, D. (2023, Mai 15). Maximizing operational efficiency of heat pumps with Model Predictive Control: An experimental case study for residential application. *14th IEA Heat Pump Conference 2023 (HPC2023) Conference Proceedings – Full Papers*. 14th IEA Heat Pump Conference, Chicago, IL, USA. https://heatpumpingtechnologies.org/publications/presentation-no-805-maximizing-operational-efficiency-of-heat-pumps-withmodel-predictive-control-an-experimental-case-study-for-residential-application-14th-iea-heat-pump-conference-chicago/

Göbel, S., Waiz, K., Christian, V., & Dirk, M. (2023). The Impact of Controller Settings in Heat Pumps: Numerical Findings and Experimental Verification. *36th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2023)*, 827–838. https://doi.org/10.52202/069564-0075

Huang, C., Seidel, S., Paschke, F., & Braunig, J. (2022). A reinforcement learning approach for optimal heating curve adaption. *2022 IEEE 27th International Conference on Emerging Technologies and Factory Automation (ETFA)*, 1–4. https://doi.org/10.1109/ETFA52439.2022.9921461

Huang, C., Seidel, S., Pruvost, H., & Bräunig, J. (2023). Ein Reinforcement-Learning-Ansatz für die Optimierung von Heizkurven: Vorlauftemperaturanpassung mittels Q-Learning. *atp magazin*, *65*(4), 70–77. https://doi.org/10.17560/atp.v65i4.2648

Huchtemann, K., & Müller, D. (2013). Simulation study on supply temperature optimization in domestic heat pump systems. *Building and Environment*, *59*, 327–335. https://doi.org/10.1016/j.buildenv.2012.08.030

Ionesi, A., Jradi, M., Eric Thorsen, J., & Veje, C. (2015, Dezember 7). Simulation of An Adaptive Heat Curve for Automatic Optimization of District Heating Installation. *Proceedings of Building Simulation 2015: 14th Conference of IBPSA*. 2015 Building Simulation Conference. https://doi.org/10.26868/25222708.2015.2823

Kuboth, S., Weith, T., Heberle, F., Welzl, M., & Brüggemann, D. (2020). Experimental Long-Term Investigation of Model Predictive Heat Pump Control in Residential Buildings with Photovoltaic Power Generation. *Energies*, *13*(22), 6016. https://doi.org/10.3390/en13226016

Mothilal, R. K., Sharma, A., & Tan, C. (2019). *Explaining Machine Learning Classifiers through Diverse Counterfactual Explanations*. https://doi.org/10.48550/ARXIV.1905.07697

Nielsen, M. A. (2015). *Neural Networks and Deep Learning* (Bd. 25). Determination press. http://neuralnetworksanddeeplearning.com

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., & Dubourg, V. (2011). Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, *12*, 2825–2830.

Potočnik, P., & Govekar, E. (2019). Adaptive optimization of heating curves in buildings heated by a weather-compensated heat pump. *Science and Technology for the Built Environment*, *25*(10), 1380–1393. https://doi.org/10.1080/23744731.2019.1616984

Remmen, P., Lauster, M., Mans, M., Fuchs, M., Osterhage, T., & Müller, D. (2018). TEASER: An open tool for urban energy modelling of building stocks. *Journal of Building Performance Simulation*, *11*(1), Article 1. https://doi.org/10.1080/19401493.2017.1283539

Rohrer, T., Frison, L., Kaupenjohann, L., Scharf, K., & Hergenröther, E. (2023). Deep Reinforcement Learning for Heat Pump Control. In K. Arai (Hrsg.), *Intelligent Computing* (Bd. 711, S. 459–471). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-37717-4_29

Ruhnau, O., Hirth, L., & Praktiknjo, A. (2019). Time series of heat demand and heat pump efficiency for energy system modeling. *Scientific Data*, *6*(1), 189. https://doi.org/10.1038/s41597-019-0199-y

Shaqour, A., & Hagishima, A. (2022). Systematic Review on Deep Reinforcement Learning-Based Energy Management for Different Building Types. *Energies*, *15*(22), 8663. https://doi.org/10.3390/en15228663

Stoffel, P., Berktold, M., & Müller, D. (2024). Real-life data-driven model predictive control for building energy systems comparing different machine learning models. *Energy and Buildings*, *305*, 113895. https://doi.org/10.1016/j.enbuild.2024.113895

Stoffel, P., Henkel, P., Rätz, M., Kümpel, A., & Müller, D. (2023). Safe operation of online learning data driven model predictive control of building energy systems. *Energy and AI*, *14*, 100296. https://doi.org/10.1016/j.egyai.2023.100296

TensorFlow Developers. (2024). *TensorFlow* (v2.15.1) [Software]. [object Object]. https://doi.org/10.5281/ZENODO.4724125

Vaillant. (2024). *Vorlauftemperatur: Fußbodenheizung & Heizanlage richtig einstellen*. https://www.vaillant.de/heizung/heizung-verstehen/tipps-rund-um-ihre-heizung/vorlauf-rucklauftemperatur/

Viessmann. (2024). *Heizkurve richtig einstellen*. https://www.viessmann.de/de/wissen/anleitungen-und-tipps/heizkurve-einstellen.html

Wang, J. (2022). *An Intuitive Tutorial to Gaussian Processes Regression* (arXiv:2009.10862). arXiv. http://arxiv.org/abs/2009.10862

Wüllhorst, F., Maier, L., Jansen, D., Kühn, L., & Hering, D. (2022, Oktober 26). BESMod—A Modelica Library providing Building Energy System Modules. *Proceedings of the American Modelica Conference 2022*. The American Modelica 2022 Conference, Dallas, TX, US. https://doi.org/10.3384/ECP211869

ACKNOWLEDGEMENT

We gratefully acknowledge the financial support by the German Federal Ministry for Economic Affairs and Climate Action (BMWK), promotional references 03EN4076A.