Energy & Buildings 374 (2027) 118448

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

Adaptive heating curves using data-driven methods to improve the energy efficiency of buildings

Florian Will*, Matthias Mersch, Christian Vering, Dirk Müller

Institute for Energy Efficient Buildings and Indoor Climate, E.ON Energy Research Center, RWTH Aachen University, Mathieustr. 10, Aachen, 52074, Germany

ARTICLE INFO

Keywords:
Heat pump
Heating curve
Optimisation
Machine learning
MPC

ABSTRACT

Heat pumps are a key technology for electrifying and decarbonising heating in the building sector. Their efficiency is strongly coupled to the heat supply temperature: the lower the supply temperature, the higher the coefficient of performance. In practice, the supply temperature is governed by the heating curve, which relates the supply temperature setpoint to the outdoor air temperature. This heating curve is often poorly configured, as proper parametrisation requires testing and detailed knowledge of the specific building. This paper presents a method to optimise the heating curve automatically using data-driven techniques, which can be implemented as a cloud-based service, thereby eliminating the need for manual parametrisation by a technician. The method first generates diverse operational data through an automated system-identification procedure. A gaussian process regressor is then trained to predict the room temperature deviation. The regressor is finally used to derive a self-optimising heating curve. The approach is evaluated in a simulation study using a high-order, multi-zone model of a single-family home coupled to an air-source heat pump, across three climate locations and two heat transfer systems (radiators and underfloor heating). The results show that the method generates suitable heating curves for all cases, achieving thermal comfort and heat pump efficiency comparable or better to reference heating curves that were parametrised optimally with full system knowledge. A model predictive controller outperforms the heating curves in almost all cases, but requires higher development effort and online optimization.

1. Introduction

Heat pumps are the leading technology to decarbonise domestic heating [1], as they can achieve high efficiencies and are locally emission-free. However, compared to conventional heating systems such as gas boilers or oil burners, they require a more careful design of the heating system to reach their potential. In particular, the heat pump coefficient of performance (COP) depends on the source temperature, T_c, and the heat supply temperature (sink temperature), T_s, according to the Carnot efficiency:

$$COP = \eta_c \cdot \frac{T_s}{T_s - T_c}, \tag{1}$$

where η_c is the Carnot efficiency factor describing how much the real system deviates from the ideal one. Although Eq. (1) provides a simplified correlation for the COP, it shows that the temperature lift, *i.e.*, the difference between source and sink temperature, should be as small as possible to achieve maximum efficiency. The source temperature is typically given by the operating conditions, *i.e.*, the ambient air temperature for air-source heat pumps (ASHPs) or the ground temperature for ground-source heat pumps. While the source temperature can be increased by using more efficient heat exchangers that reduce the temperature difference between the refrigeration cycle and the source, the potential is limited. On the heat sink side on the other hand, optimized heat transfer systems and controllers can reduce the required supply temperature significantly and thus aid a more efficient operation of the heat pump.

The required heat supply temperature is typically determined from a heating curve, which, in the most simple case, defines a linear correlation between outdoor air temperature, T_{oda}, and supply temperature setpoint, $T_{s,set}$:

$$T_{s,set} = k \cdot (T_{ro,set} - T_{oda}) + b. \tag{2}$$

In this equation, k represents the slope of the heating curve, b the linear offset, and $T_{ro,set}$ the target room temperature. Determining suitable coefficients for the heating curve requires testing and detailed information on the specific building [2]. Therefore, in practice often quite high values are chosen to ensure that the thermal power is sufficient to maintain the desired indoor temperature [3]. However, as explained above, this leads to a significantly reduced heat pump efficiency. Therefore, it is desirable to optimize the heating curve for each specific building in

* Corresponding author.
E-mail address: florian.will@eonerc.rwth-aachen.de (F. Will).

https://doi.org/10.1016/j.enbuild.2026.118448
Received 24 June 2026; Received in revised form 20 October 2026; Accepted 26 October 2026
Available online 29 October 2026
0378-7788/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).
Open access funding enabled and organized by Projekt DEAL.

Nomenclature

Abbreviations

ANN	Artificial neural network
ASHP	Air-source heat pump
COP	Coefficient of performance
DHW	Domestic hot water
DiCE	Diverse Counterfactual Explanations
GPR	Gaussian process regressor
GSHP	Ground-source heat pump
KPI	Key performance indicator
MPC	Model predictive control
PI	Proportional-integral
RBF	Radial basis function
SCOP	Seasonal coefficient of performance
soHC	Self-optimizing heating curve
TD	Thermal discomfort
UFH	Underfloor heating

Greek Letters

η_c	Carnot efficiency factor
ΔT	Temperature difference (K)

Symbols

b	Linear offset of the heating curve (°C)
$c_1 \dots c_8$	Empirical coefficients of the heat pump performance map
COP	Coefficient of performance
k	Slope of the heating curve
n	Curvature parameter of the heating curve
n_c	Compressor speed
P	Power (W)
$\dot{Q}$	Heat flow (W)
Q_{rel}	Relative heating load
t	Time (s)
T	Temperature (K or °C)
w_i	Weight of cluster i

Subscripts

Add	Additive offset
c	Cold / source
con	Condenser
dev	Deviation
el	Electrical
HeaThr	Heating threshold
n	Nominal
oda	Outdoor air
pre	Predicted
rel	Relative
ro	Room
s	Supply / sink
set	Setpoint
i	Cluster index
r	Room index

order to maximize the heat pump efficiency while maintaining thermal comfort at all times.

1.1. Literature review

The question of optimizing the heating system controller has become more important with the increase of relevance of heat pumps. Huchtemann and Müller [4] investigated the impact of different heating curve parametrisations on the seasonal coefficient of performance (SCOP) of an on-off controlled heat pump in a single-family home. They investigated parallel shifts of the heating curve (*i.e.*, changing parameter b), and changed the hysteresis of the on-off controller. Additionally, the authors developed an adaptive heating curve that is adjusted based on the position of the most-open thermostatic valve in the heating circuit. This adaptive control strategy was able to save up to 6.8 % of energy compared to an optimized but fixed heating curve for the simulated building. However, it requires information on the position of all valves to be sent to the controller regularly. More recently, Potthoff et al. [5] investigated the option to optimize heating curves based on easy-to-apply outdoor temperature-dependent guidelines. These guidelines allow users to fine-tune their heating curve, or manufacturers to implement automatic fine-tuning programs. These guidelines were found to perform significantly better than standard parallel-shift approaches, reducing the annual energy consumption by 4 % on average across a sample of about 4000 heat pump installations while also improving thermal comfort. Cholewa et al. [6] adjusted the heating curves of two heat pumps supplying heat to multi-family homes in Poland and monitored their performance over two heating seasons (November 2023 to May 2025). The authors compared adjustments based on the technical documentation of the buildings to adjustments based on monitored energy consumption. The former results in energy savings of 4 % for both buildings, while the latter strategy is able to reduce energy consumption by 10 % and 15 %, respectively, compared to the initial heating curve.

Ionesi et al. [7] proposed an algorithm to adapt heating curves for district heating networks based on the real behavior of the building in question. An initial heating curve is separated into a number of segments based on the ambient temperature. Then, if the room temperature differs from the setpoint, the required supply temperature to reach the desired room temperature is estimated and the corresponding segment of the heating curve is adjusted. Neighboring segments are adjusted as well to maintain a smooth curve. This approach was tested using a simulation model of three rooms, each modeled through a four-capacity model. In this case study, energy savings of 1 % to 5 % were achieved while maintaining thermal comfort.

Potočnik et al. [8] optimized various heating curves with the aim to minimize thermal discomfort. In addition to the standard heating curve which is only based on the ambient temperature, they also considered the ambient temperature gradient, solar irradiation, and predictions of future temperature and irradiation as inputs to the heating curve. The more complicated heating curves achieved better thermal comfort and a lower electricity consumption, but require much more information. Especially including solar irradiation measurements and temperature and irradiation predictions may proof difficult. This is likely more appropriate for model predictive control approaches rather than heating curves, whose benefits are their simplicity. Potočnik and Govekar [9] then investigated an adaptation strategy similar to that of Ionesi et al. [7], which is based on the offset in room temperature. The adaptation strategy was shown to improve thermal comfort if the parameters of the adaptation strategy are chosen well, but can also result in worse thermal comfort compared to the optimal static heating curve if the parameters are unsuitable.

Rolando and Madani [10] also investigated different heating curve adaptation strategies. The first approach involved integrating the predicted daily solar radiation into the heating curve [11]. This resulted in significant energy savings especially during spring and autumn, leading to total annual energy savings of 9 %. However, similar to the approach of Potočnik et al. [8], a drawback of this approach is the need for additional data. The authors also tried a similar approach as Ionesi et al. [7], discretising the heating curve and adapting it based on the room temperature and/or user feedback.

Le et al. [12] proposed a data-driven self-learning algorithm, which is not further specified, to adjust heating curves based on inputs such as the flow rate, supply and return temperatures, and the control signal. The algorithm was applied to a gas boiler in a field test, showing that the supply temperature setpoint could be reduced by an average of 8 K.

Beyond improved or even adaptive heating curves, more complex control strategies for building energy systems are becoming increasingly

popular in research. Agyekum et al. [13] observe that research on smart buildings is shifting towards data-driven and adaptive energy systems to improve energy efficiency. They identify explainable AI and real-time adaptive control as promising directions for future research. Model predictive control (MPC) for example is being extensively researched. Göbel et al. [14] performed dynamic Hardware-in-the-loop experiments to compare different levels of MPC of a heat pump to the built-in controller using a heating curve. The MPC was able to reduce the electrical energy consumption by 19 % on a cold winter day and by over 50 % on a sunny day in autumn, while also being able to reduce the amount of compressor starts. However, implementing MPC is quite complex and requires much more data and modeling effort as well as computational resources compared to heating curve adjustments.

1.2. Research gap and contributions

In this paper, we use data-driven machine-learning approaches to generate self-optimizing heating curves based on Gaussian process regressors (GPRs). This approach avoids the significant modeling effort needed for MPC algorithms based on detailed building energy system models. Within this paper, the workflow is applied to a multi-zone building hull, equipped with either a radiator or an underfloor heating system, and evaluated for three climates. Evidence for its transferability to other buildings comes from our preliminary study [15], in which the same method was applied to a different building with a different heat pump model and a different thermal mass. That study used a simplified single-zone building model, so it does not cover the multi-zone case with individual thermostatic valves considered here, but it shows that the workflow is not tied to one specific building.

Compared to the studies on adaptive heating curves by Ionesi et al. [7] and Potočnik and Govekar [9], we consider both solar irradiation and wind, and apply the methodology to multiple heated rooms with individual thermostatic valves in our analysis. Compared to the data-driven approach of Le et al. [12], we investigate ASHPs rather than gas boilers and consider different locations. As an applied contribution, the novelty of this work lies in the combination and integration of automated system identification, machine learning with GPRs and optimisation to generate self-optimizing heating curves, and in demonstrating this workflow across different heat transfer systems and climates.

The remainder of this paper is structured as follows: Section 2 discusses the methodology, focussing on system identification, data-driven modeling and the algorithm to generate the self-optimizing heating curves. Section 3 presents the results and Section 4 discusses their implications, before Section 5 concludes the paper.

2. Methodology

The following section describes the methodology by introducing the simulation models and the machine learning method for generating the self-optimizing heating curves.

2.1. Simulation models

The dynamic simulations of the heat pump and building models are conducted using the object-oriented modeling language Modelica. The system is implemented utilizing the BESMod library [16], a Modelica library specifically designed for building energy systems which builds upon the libraries IDEAS [17], Buildings [18], IBPSA [19], AixLib [20], and BuildingSystems [21]. The modeled building represents a newly built single-family home with a usable floor area of 158 m^2. It is modeled using the high-order building model from the AixLib library, which has been described and validated in previous publications [22,23]. The building envelope is parametrized according to the 1995 German Thermal Insulation Ordinance (WSVO95), resulting in a nominal heating load of 53 W/m^2 at a design ambient temperature of −14 °C, calculated in accordance with the EN 12,831 standard.

Table 1
Room usage and corresponding target temperatures.

room ID	Usage	Target Temperature (°C)
1	Living Room	20
2	Dining Room	20
3	Utility Room	20
4	Hallway	17
5	Kitchen	20
6	Bathroom 1	20
7	Hallway	17
8	Child's Room 1	20
9	Child's Room 2	20
10	Bathroom 2	24
11	Bedroom	20
12	Hallway	17

To accurately capture the thermal dynamics and the spatial distribution of the heat demand, the building is divided into 12 distinct heated rooms representing different rooms of the building (e.g., kitchen, bathroom, living room). Each room is assigned an individual, constant temperature setpoint, as detailed in Table 1 after DIN EN 12831-1 [24]. Notably, the primary bathroom (room 10) operates with the highest temperature setpoint of 24 °C, making it consistently the room that needs the highest T_s throughout the simulation.

The study uses two different types of heat transfer systems: radiators and underfloor heating (UFH) systems. The radiator system is designed for nominal supply and return temperatures of 55/45 °C, whereas the UFH system operates at lower nominal temperatures of 38/31 °C. A combined 900-liter buffer storage tank is utilized to provide both space heating and domestic hot water (DHW). The daily hot water demand is assumed to be 200 liters with a required temperature of 60 °C. It is modeled using deterministic draw-off events over a 24-h cycle after DIN EN 15450 [25].

The required thermal energy is provided by an inverter-controlled ASHP using propane (R290) as the refrigerant. The heat pump is modeled based on the approach proposed by Wüllhorst et al. [26]. This approach utilizes a performance map that determines the efficiency of the heat pump as a function of the compressor speed (n_c), T_{oda}, and T_s. The electrical power demand (P_{el}) and the condenser heat flow ($\dot{Q}_{con}$) required to calculate the efficiency are calculated using the following polynomial equations fitted to manufacturer measurement data:

$$P_{el} = (c_1 \cdot T_s^2 + c_2 \cdot T_s + c_3) \cdot n_c^2 + (c_4 \cdot T_s^2 + c_5 \cdot T_s + c_6) \cdot n_c; \quad (3)$$

$$\dot{Q}_{con} = (c_7 \cdot T_{oda} - T_s + c_8) \cdot n_c, \quad (4)$$

where c_1 to c_8 are empirical coefficients derived from the measurement data.

As an example, Fig. 1 shows the COP of the heat pump model depending on the operating conditions for a fixed compressor frequency of 50 Hz. It can be seen that the COP increases for decreasing T_s and increasing T_{oda}.

The overall control logic of the heating system is integrated into the Modelica model. The baseline $T_{s,set}$ is determined by an ambient temperature-dependent heating curve, implemented according to Lämmle et al. [27]. The relative heating load (Q_{rel}) and the resulting $T_{s,set}$ are calculated as follows:

$$Q_{rel} = \frac{T_{HeaThr} - T_{oda}}{T_{HeaThr} - T_{oda,n}}; \quad (5)$$

$$T_{s,set} = T_{ro} + \left(\frac{T_{s,n} + T_{r,n}}{2} - T_{ro} \right) \cdot Q_{rel}^{\frac{1}{n}} + \frac{T_{s,n} - T_{r,n}}{2} \cdot Q_{rel} + T_{Add}, \quad (6)$$

where T_{HeaThr} is the heating threshold temperature, T_{Add} is an additive temperature offset and T_{ro} is the room temperature.

Based on this $T_{s,set}$ and the temperature at the top of the buffer storage tank, the heat pump is turned on or off. Both the space heating and DHW storage controls utilize a hysteresis deadband of ±3 K. When the heat pump is active, an internal proportional-integral (PI) controller ad-

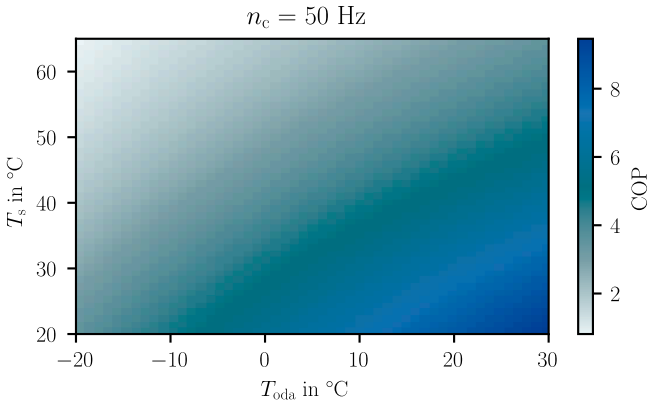


Fig. 1. COP calculated from the heat pump model depending on supply temperature (y-axis) and ambient temperature (x-axis) for a constant compressor speed of 50 Hz.

justs the compressor speed based on the deviation between the actual T_s and $T_{s,set}$. The heat pump is constrained by its operating envelope and by minimum on- and off-times which mirrors a realistic heat pump controller. DHW charging is managed by a hysteresis controller that actuates a three-way valve to prioritize the DHW tank when the temperature drops below the threshold. Additionally, the water pump of the heating circuit is regulated by a PI controller designed to maintain a defined temperature difference across the condenser on the water side: 5 K during space heating operation and 8 K during DHW loading. Actuator delays are included, as valves, pumps and the hydraulic network are physically simulated with actuator rise times enabled.

The primary goal of this paper is to improve and adapt the heating curve. As the heating curve directly dictates $T_{s,set}$, any changes made to it will have a direct influence on both the thermal comfort within the building and the overall operational efficiency of the heat pump. These effects are evaluated using the simulation models to assess the performance of different heating curves.

2.2. Automated building system identification

The objective of the system identification is to generate comprehensive operational data, which is essential for developing self-optimizing heating curves. Training data-driven models requires a prior system identification of the building to capture its thermal dynamics. To ensure high training data quality, the dataset must contain as many different combinations of T_s and T_{oda} as possible, thereby maximizing the information content. Relying on a single, fixed heating curve for system identification restricts the available T_s and T_{oda} combinations to an operational band defined by that specific curve. To overcome this limitation, an algorithm employing varying heating curves with different slopes is utilized to increase the diversity of the training data.

Under ideal operating conditions with the optimal T_s, the thermostatic valve of the critical room, which is here the primary bathroom (room 10) with the highest room temperature setpoint of 24 °C, would remain completely open. To replicate this ideal behavior and eliminate the valve position as an additional disturbance variable during data generation, the valve of the critical room is fixed at a 100 % open position. Meanwhile, the valves in the remaining rooms are controlled to maintain their respective $T_{ro,set}$.

The data generation algorithm begins with an initial 8-day phase. During this period, the building energy system is simulated using four different linear heating curves, each active for two days. The slopes (k) for these initial curves are chosen randomly within bounds that depend on the heat transfer system (radiators or UFH). The linear heating curve is defined by the following equation:

$$T_{s,set} = k \cdot (24\ °C - T_{oda}) + 24\ °C, \tag{7}$$

where the base temperature of 24 °C corresponds to the target setpoint of the critical room. The room temperature deviation in the critical room is defined as:

$$T_{ro,dev} = T_{ro} - T_{ro,set}. \tag{8}$$

To account for the thermal inertia of the building and the heating system, the continuous measurements of $T_{ro,dev}$ are averaged over discrete, non-overlapping 8-h blocks. Each averaged deviation value is then linked to the specific slope k that was active during that 8-h window.

Following the initial phase, a linear regression is performed using the averaged $T_{ro,dev}$ as the dependent variable and the corresponding slope k as the independent variable. By setting $T_{ro,dev} = 0$ in the resulting regression equation, a new slope k is calculated, representing the slope required to match the room temperature to its setpoint. A new heating curve utilizing the calculated k is then simulated for 7 days, generating additional training data. This iterative process is repeated 9 times. In each iteration, the linear regression is updated using the entire pool of accumulated historical data—comprising the initial 8 days plus all subsequent 7-day periods—to calculate the next slope.

The complete system identification spans a total duration of 71 days (8 initial days followed by 9 iterations of 7 days). In this paper, the system identification always starts on the 1st of January. This is arbitrary, in a real application it would be the time of heat pump installation. In case of installation in the summer, the system identification should start with the beginning of the heating period. The system identification yields a highly diverse training dataset, as the variation of the heating curve slopes produces a wide array of T_s and T_{oda} combinations. Furthermore, the iterative regression process ensures that the algorithm converges toward a well-fitted slope k, thereby keeping thermal discomfort on a low level for the building's inhabitants during the system identification phase. While this iterative method is capable of adapting the heating curve to different environmental conditions, it remains constrained by the predefined linear equation. The following section details how a more flexible heating curve is generated from data-driven models utilizing the training data generated through this procedure.

2.3. Data driven modelling to obtain self-optimizing heating curves

The generation of the self-optimizing heating curves (soHC) in this paper relies on data-driven models, specifically GPRs. GPRs are suitable for regression tasks and offer the advantage of providing the predicted variance alongside the predicted mean value. The models are implemented utilizing the Scikit-learn machine learning library [28]. The inputs to the GPR model are T_s and T_{oda}, while the target output is the room temperature deviation ($T_{ro,dev}$). To account for the thermal inertia of the building energy system, all input and output variables are averaged over discrete 8-h intervals before being fed into the model.

The GPR utilizes a composite kernel structure consisting of a Radial Basis Function (RBF) kernel scaled by a Constant kernel, with the addition of a White kernel to account for measurement and system noise in the training data. During the model configuration, the lower bound for the RBF length scale is set to 10. This limit was determined iteratively; allowing a lower length scale results in a regression surface that is sensitive to local data variations. This leads to a rough regression surface that cannot generate a stable, monotonic heating curve. The time it takes to train the GPR for the given amount of datapoints (189) ranges between 0.15 s and 0.3 s on an Intel Xeon processor with 2.9 GHz on a single core.

Once the GPR is trained on the 71-day dataset, the principle for generating the soHC is to interrogate the model by fixing T_{oda} and searching for an optimal T_s that yields a target $T_{ro,dev}$. To extract this inverse relationship, the Diverse Counterfactual Explanations (DiCE) method by [29] is used. DiCE generates post-hoc explanations by framing the input extraction as an optimization problem, finding the necessary model inputs to achieve a desired output. The evaluation grid for T_{oda} is established by identifying the minimum and maximum outdoor temperatures

present in the training data, rounding these boundaries, and creating a discrete array with 1 K increments.

For the case shown here, DiCE is functionally equivalent to a grid search, a root-finding algorithm or similar algorithms. It is used because the inverse problem is no longer one-dimensional once further controllable inputs are considered. When several inputs are varied simultaneously, different input combinations achieve the same target ($T_{ro,dev}$), and the counterfactual formulation returns feasible solutions that change as few inputs as possible. Retaining this formulation keeps the workflow extensible to the multi-input case without reimplementation.

For each T_{oda} grid point, the DiCE algorithm utilizes a genetic algorithm optimizer to search for one counterfactual T_s value. The search space for T_s is restricted to a permissible operational range of 20°C to 65°C. The optimization is configured to find a T_s that results in a $T_{ro,dev}$ within the range of 0.25 K to 0.75 K. This positive target range acts as a thermal safety buffer, ensuring the room temperature remains slightly above the setpoint to prevent underheating of the thermal room. For the given 189 datapoints, the DiCE algorithm takes between 3 s and 4 s to complete the optimization for one heating curve on an Intel Xeon processor with 2.9 GHz on a single core.

After the DiCE algorithm identifies the optimal T_s for each T_{oda} grid point, these discrete points are linearly interpolated to construct a continuous, piecewise linear function representing the final soHC. Since the data-driven model is only reliable within the T_{oda} bounds observed during the training period, the heating curve must be extrapolated to handle extreme weather conditions outside this range. This extrapolation is performed linearly, utilizing an extrapolation slope that is calculated as the arithmetic mean of all the individual piecewise slopes between the 1 K grid points of the generated soHC.

2.4. Model predictive controller

MPC serves as the upper benchmark in this study. Compared to the soHC it requires considerably higher development and commissioning effort, since a process model of the building must be identified and forecasts as well as a solver must be available during operation. In return, it can exploit information that a heating curve is unable to use, namely disturbances such as solar irradiation and internal gains, and weather forecasts.

The MPC is implemented using the framework presented in Stoffel et al. [30], which has been applied to heat pump and building control in Stoffel et al. [31] and Göbel et al. [14]. The manipulated variables are the supply temperature setpoint of the heat pump and the opening degrees of the thermostatic valves of the twelve thermal zones. The controlled variables are the corresponding room temperatures. The cost function balances thermal comfort against system efficiency. Comfort is represented by keeping the room temperatures within defined bounds, efficiency by minimizing the supply temperature setpoint. The temperature bounds are formulated as soft constraints whose violation enters the cost function quadratically. The optimal control problem is formulated in Appendix A.5. The controller operates with a step size of 10 min and a prediction horizon of 30 steps, corresponding to 5 h. We assume a perfect forecast of solar irradiation, outdoor air temperature and internal gains. This is not directly applicable to praxis, as especially predictions of internal gains that come from user behavior are difficult to obtain.

Process models

Twelve individual linear regression models are used, one per room, each predicting the change of the respective room temperature over the next time step. The inputs comprise state, manipulated and disturbance variables, each considered with an individual lag. A lag of n denotes the n most recent samples including the current one. The inputs are summarized in Appendix A.5 Table A.2.

Training data generation

The linear regression models are trained with training data obtained from the simulation model. To ensure that the transient heating behavior is represented in the data, the room temperature setpoints are varied randomly by ±3 K every five hours. In order to capture the effects of the two manipulated variables both in isolation and in their coupled dynamics, three operating and control variations are simulated. They are described in more detail in Appendix A.4. The training data simulations are carried out for representative periods in winter and spring. Each variation is recorded for one week, so that one period spans $5 + 5 \times 7 = 40$ d of contiguous simulation time. The data generated in this way form the basis for the training of the twelve room-wise regression models.

Table 2
Overview of all conducted simulations. The simulations differ in location, transfer system and controller.

Location	System	HC low	HC high	soHC	MPC
Warm	UFH	TD	TD	TD	TD
	Radiator	TD	TD	TD	TD
Average	UFH	TD, FY	TD	TD	TD
	Radiator	TD, FY	TD	TD	TD
Cold	UFH	TD	TD	TD	TD
	Radiator	TD	TD	TD	TD

TD = Typical Days, FY = Full Year

3. Results

In this paper we present results from testing the described algorithm to generate a soHC in 3 different locations and using 2 different heating systems, radiators and UFH systems. An overview of the different tested variations is shown in Table 2. The locations are determined from clustering of German weather data for test reference years [32]. From the clustering, we chose a warm location (Gelsenkirchen), an average location (Langerringen), and a cold location (Schönau am Königssee). Additional information on the weather data is in Fig. A.1 in Appendix A.1.

As references, we also simulate 2 conventional heating curves in addition to the soHC for each location. These are indicated as “HC low” and “HC high”. The reference heating curves are based on the heating curve described in Eq. (6). The parametrization of the reference heating curves is shown in Table B.1. The difference between the low and the high reference heating curve is the parameter T_{Add}, which is a parallel shift of the heating curve.

In addition to the three heating-curve-based controllers, the MPC described in Section 2.4 is simulated for every location and transfer system. It serves as an upper benchmark, since it uses disturbance information and forecasts that are not available to the heating curves.

To keep the total time needed for the simulations manageable, we perform the simulations on typical days for every variation. These typical days are determined from the clustering of a whole year of weather data. We use the python package tsam [33] with hierarchical clustering. The clustering algorithm calculates 8 clusters that are similar in their weather characteristics (outdoor air temperature, relative humidity and solar radiation). Because the heat pumps operation is mainly influenced by the outdoor air temperature, the weight on the outdoor air temperature in the clustering algorithm is the highest (= 0.8), while the relative humidity and solar radiation are weighted lower (= 0.15 and = 0.05). The Key Performance Indicators (KPIs) calculated for the typical days are used to predict the KPIs for the complete year. This is done by simulating the 8 typical days and multiplying the daily results by specific weighting factors, which correspond to the number of days in the year that each cluster represents. Table A.1 in the appendix shows the typical days with their weighting factors for every location. To validate this approach, we do simulations of full years for the average location with the low reference heating curve and compare it to the results predicted from the typical-day approach.

In total, we simulate 4 different controllers for every location and heating system, resulting in 24 different simulations.

3.1. Self-optimizing heating curves

Based on the data-driven model developed in Section 2.3, the soHC is generated and evaluated in comparison with the reference heating curves (HC low and HC high) for every system variation detailed in Table 2. The upper subplots of Fig. 3 illustrate these heating curves for the location with average outdoor air temperatures, plotting $T_{s,set}$ against T_{oda}. The results for the radiator transfer system are displayed on the left-hand side (Fig. 3(a)), while the UFH system is shown on the right-hand side (Fig. 3(b)). An analysis of the heating curves reveals that the soHC correctly dictates higher $T_{s,set}$ values for the radiator system compared to the UFH system, demonstrating that the data-driven models successfully learned the required temperature levels for each transfer system. The soHC exhibits a slope similar to the reference heating curves, showing that the physical dependency between $T_{s,set}$ and T_{oda} was accurately captured by the data-driven model. Overall, the soHC closely aligns with the reference heating curves. This alignment confirms the efficacy of the proposed method, especially considering that the reference heating curves used in this study are already parametrized for the actual transfer systems. In practical real-world applications, heating curves are typically configured manually by technicians and are rarely well-tuned as described in Potthoff et al. [5].

The lower three subplots of Fig. 3 display the predicted deviations of the room temperature ($T_{ro,dev,pre}$) for each of the three heating curves. These predicted deviations are generated by feeding the respective $T_{s,set}$, T_{oda} value pairs into the data-driven model described in Section 2.3. Since GPRs are utilized, the model outputs not only the predicted mean value but also the standard deviation for each input pair. This is visualized as a 95 % confidence interval, represented by a grey shaded area around the predicted mean values. For the soHC, $T_{ro,dev,pre}$ consistently aligns with the targeted safety buffer of 0.25 K to 0.75 K. This behavior is expected, as the DiCE algorithm was constrained to generate the soHC within this comfort band. The data-driven model predicts larger deviations for the reference heating curves. HC low for example is parametrized too conservatively, resulting in negative temperature deviations that indicate insufficient heating. Additionally, the 95 % confidence intervals widen towards the lowest and highest outdoor temperatures shown in Fig. 3. This widening indicates a decrease in model reliability in these regions, as the model is forced to extrapolate. This is due to a lack of very low and very high temperatures in the 71-day training dataset.

The results for the cold and warm locations are presented in Figs. 2 and 4, respectively. The overall trends remain consistent across these different climates, demonstrating that the data-driven method functions effectively under varying weather dynamics. However, the specific climate impacts the available training data. For instance, in the cold location, the training period did not include outdoor temperatures above 5 °C. Consequently, the 95 % confidence interval becomes wider for $T_{oda} > 5$ °C. This paper addresses data scarcity at the extremes by extrapolating the heating curve linearly based on the mean slope. Alternative strategies could be implemented in these highly uncertain regions, such as reverting to a conservative baseline heating curve when the predicted model variance exceeds a certain threshold.

The results of a sensitivity analysis of the soHC generation are presented in Appendix B.2 Figs. B.1 and B.2 exemplarily for the average location with radiator and UFH system. A variation of the comfort bounds (base setting [0.25 K, 0.75 K]) shows that the level of the soHC increases with increasing bounds. The variation of the averaging interval of the training data (base setting 8 h) shows little influence on the generated soHC. The sensitivity of the soHC generation to the amount of training data (base setting 71 days) is examined. For the radiator system, 43 days of training data are already sufficient to generate a soHC. For the UFH system, the variant with 43 days of training data does not yield a sensible soHC. This is likely due to the higher thermal capacity of the UFH system, which requires a larger amount of training data to be captured compared to the radiator system. Lastly we examined the influence of the length scale of the GPR, the results show that it does not have meaningful influence on the generated soHC. The training time of the GPR increases with increasing number of training data points. The maximum training time observed in the sensitivity analysis is 1.8 s for the largest amount of training data points that we looked at (390). The execution time of the DiCE algorithm is not sensitive to the number of datapoints in the application of this paper, because the number of datapoints enter primarily through the cost of the GPR mean prediction, which is computationally negligible compared to the time the optimizer takes. The forward pass of a GPR with 390 training data points takes ca. 140 μs on our hardware (Intel Xeon processor with 2.9 GHz).

The genetic algorithm used by DiCE is stochastic, so repeated runs on the same GPR yield slightly different soHCs. To quantify this effect, the soHC for the average location with the UFH system was generated 25 times with different random seeds while keeping the trained GPR fixed. Across the runs, the mean standard deviation over all outdoor temperatures is 0.14 K, so the stochasticity has only a minor effect on the generated soHC. The results are shown in Appendix B.3 Fig. B.3.

Table 3
Comparison of KPIs obtained from full-annual simulations versus predictions based on 8 typical days for the average climate location with the reference heating curve HC low.

Transfer System	KPI	Annual Simulation	Typical-day approach	Deviation
UFH	SCOP	3.86	3.95	2.28 %
UFH	TD in Kh/d	23.74	24.94	4.8 %
Radiator	SCOP	3.08	3.17	2.84 %
Radiator	TD in Kh/d	22.4	22.67	1.19 %

3.2. Annual simulations

This section validates the accuracy of this typical day approach by comparing its predictions against a full annual simulation. The validation is demonstrated using the average climate location with the low heating curve as a reference, as previously defined in Table 2. The evaluation focuses on two primary KPIs: the SCOP of the heat pump and the cumulative thermal discomfort (TD), measured in Kelvin-hours per day (Kh/d), which is the sum of instantaneous thermal discomforts per room over the whole year. The calculation formulas for the SCOP and the thermal discomfort are provided in Appendix A.3.

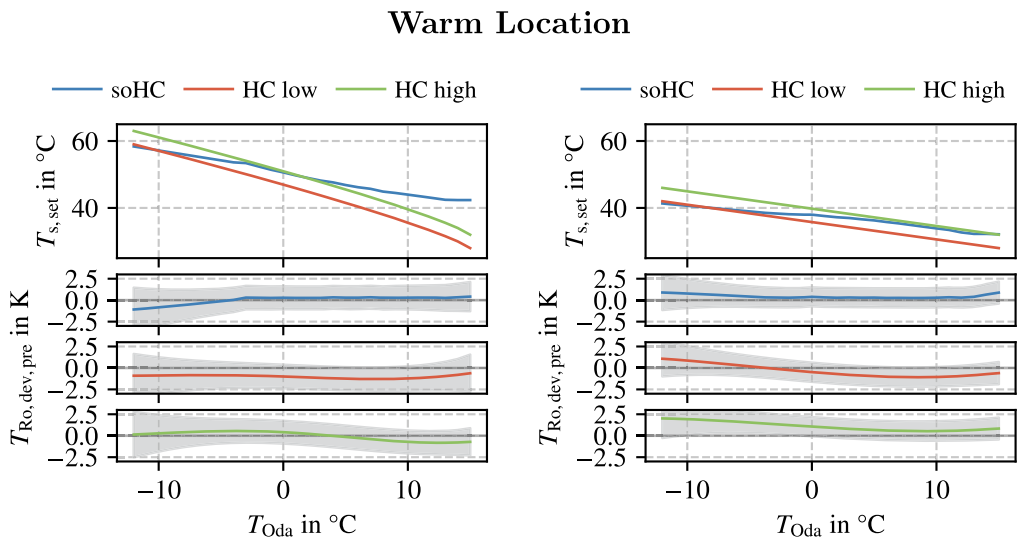
The results of this comparison for both the UFH and radiator heat transfer systems are summarized in Table 3.

As shown in Table 3, the predicted SCOP values are close to the full-year results for both the UFH and radiator systems, with relative deviations of 2 % to 3 %. Regarding Thermal Discomfort, the prediction for the Radiator system is also highly accurate, showing a deviation of 1.19 %. The largest discrepancy is observed in the TD prediction for the UFH system, which has a deviation of 4.8 % compared to the full annual simulation.

The accuracy of the predictions across both efficiency and comfort metrics is rated sufficiently high. The reduction in computational time outweighs the loss in precision. From this point forward, all subsequent results presented in this paper are derived from typical day simulations.

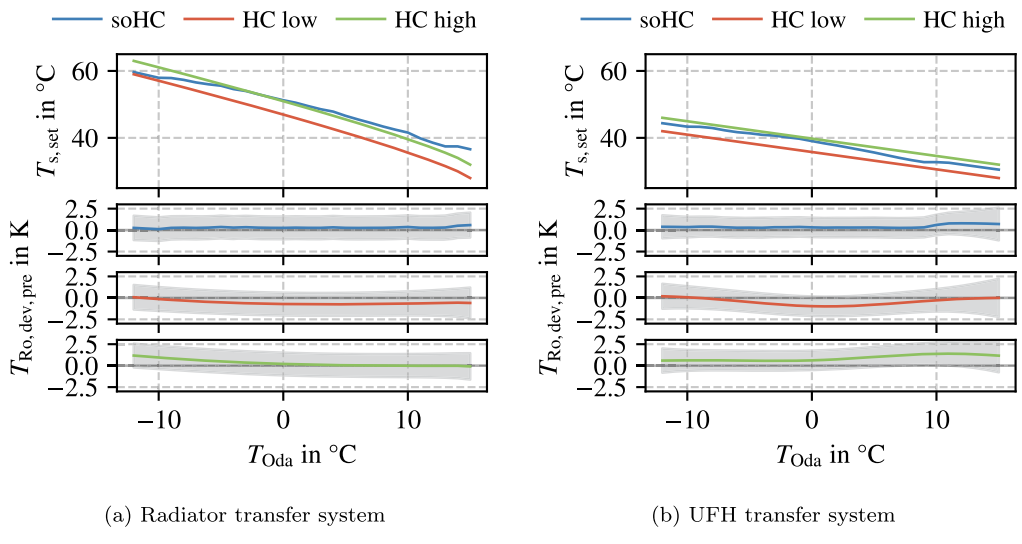
3.3. Typical day simulations

This section analyzes the results of the typical day simulations for the warm, average and cold location with the UFH and radiator transfer systems. Table B.2 shows the numerical results of the typical day simulations. Figs. 5–7 illustrate the thermal discomfort (y-axis) of the three different locations over the SCOP (left subplot) and over the relative savings of electrical energy ΔW_{el} (right subplot). Each figure contains subfigures separating the results for the radiator (left) and UFH (right) transfer systems. In these scatter plots, the soHC is depicted in blue, HC



(a) Radiator transfer system (b) UFH transfer system

Fig. 2. soHC (blue) and reference heating curves (red and green) for the warm location using (a) a radiator system; and (b) an UFH system. The top subplot shows the heating curves, the bottom subplots show the predicted room temperature deviations for each HC. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



(a) Radiator transfer system (b) UFH transfer system

Fig. 3. soHC (blue) and reference heating curves (red and green) for the average location using (a) a radiator system; and (b) an UFH system. The top subplot shows the heating curves, the bottom subplots show the predicted room temperature deviations for each HC. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

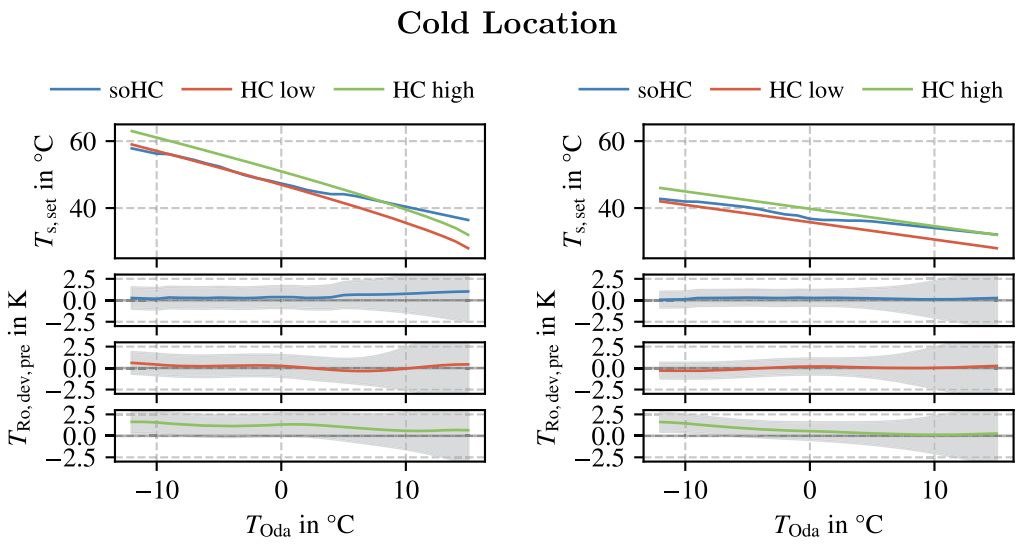
low in red, HC high in green and the MPC in purple. Additionally, a dashed line represents the linear interpolation between the two reference heating curves. The savings of electrical energy are evaluated with respect to HC high and are defined as

$$\Delta W_{el} = \frac{W_{el,HC\,high} - W_{el}}{W_{el,HC\,high}} \cdot 100\ \%, \tag{9}$$

where W_{el} is the electrical energy consumed by the heat pump over the weighted typical days.

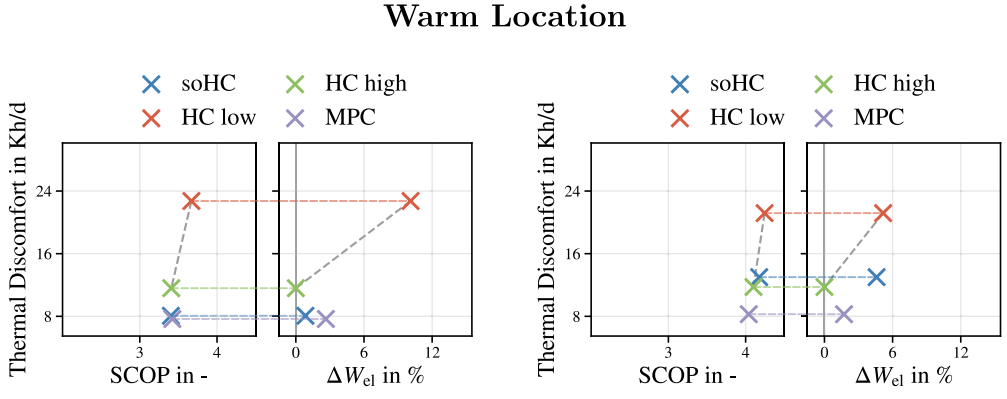
For all variations, the thermal discomfort is dominated by the thermal discomfort of the primary bathroom (room 10), as this room has the highest $T_{ro,set}$ of 24°C. This is demonstrated in Fig. B.4 Appendix B.4.

One key aspect of the analysis is the trade-off between thermal discomfort and SCOP. Implementing a higher heating curve (HC high) generally minimizes thermal discomfort by providing more heat to the building; however, it simultaneously reduces the SCOP. This occurs because a higher heating curve necessitates higher $T_{s,set}$. This results in a higher required temperature lift of the heat pump and consequently a lower efficiency, as dictated by Eq. (1). This trade-off is seen across all locations and transfer systems: HC high consistently yields lower discomfort but also a lower SCOP compared to HC low. Warmer locations show higher SCOP values due to higher average outdoor temperatures. Similarly, the UFH system achieves higher SCOPs than the radiator system, as its larger heat transfer area enables operation at lower T_s. The results also highlight the influence that the choice of heating curve has



(a) Radiator transfer system (b) UFH transfer system

Fig. 4. soHC (blue) and reference heating curves (red and green) for the cold location using (a) a radiator system; and (b) an UFH system. The top subplot shows the heating curves, the bottom subplots show the predicted room temperature deviations for each HC. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



(a) Radiator transfer system (b) UFH transfer system

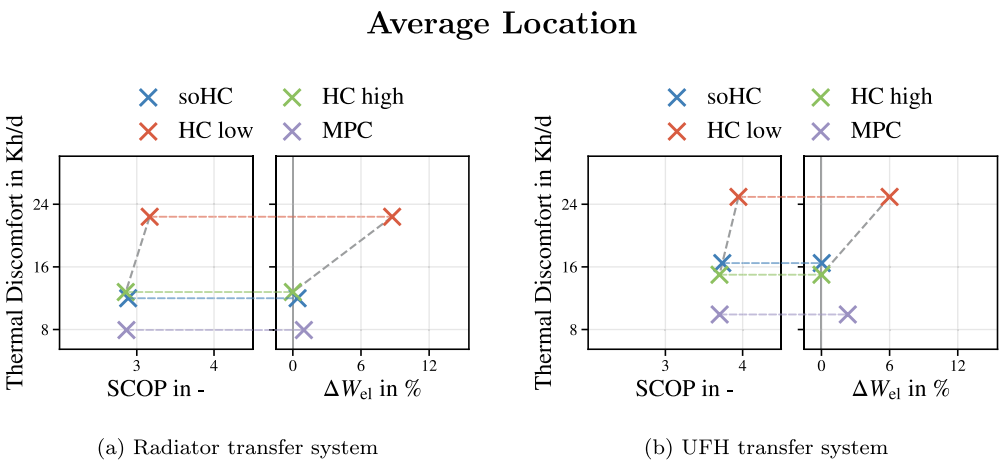
Fig. 5. Thermal discomfort (y-axis) over the SCOP (left subplot) and over the savings of electrical energy compared to HC high (right subplot) for the different controllers in the warm location using (a) radiator-based heating; and (b) UFH systems.

on the efficiency of the heat pump. For instance, in the cold location utilizing radiators, there is a more than 10 % difference in SCOP between the highest and lowest heating curves, but also an over 40 % reduction in thermal discomfort.

We evaluate ΔW_{el} to quantify the electricity demand of the different controllers. For the average location with the radiator system, the soHC and the MPC reach practically the same SCOP (2.89 and 2.87), while the MPC consumes 0.6 percentage points less electrical energy. The cold location with the radiator system shows the same effect, with SCOPs of 2.53 and 2.54 and savings of 9.2 % and 10.9 %. The reason is that the SCOP is a ratio and therefore does not penalize heat that is supplied although it is not needed. During mild periods, the heating-curve-based controllers keep the heat pump in operation. Because of the high outdoor air temperatures, this operation takes place at a small temperature lift and thus at a high COP, which raises the SCOP. The MPC switches the heat pump off in these periods, which can lead to similar SCOPs, but lower electricity demand.

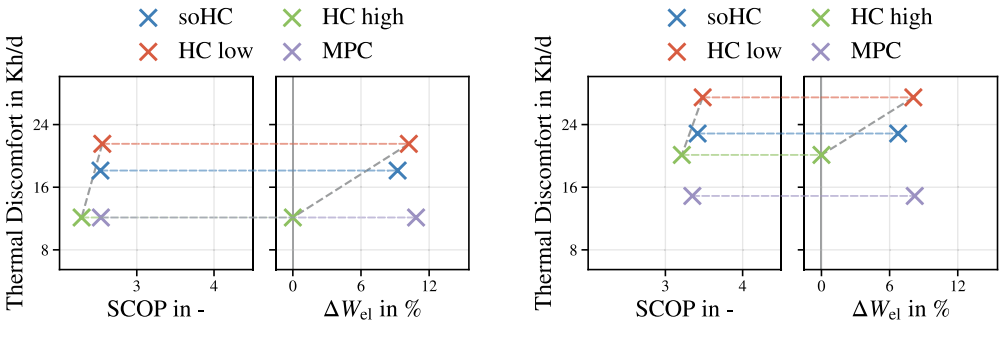
The dashed interpolation lines serve as a benchmark for evaluating the soHCs. A data point falling to the right of this line indicates an improved heating curve (yielding a better SCOP for a given level of discomfort), whereas a point to the left indicates a worse heating curve. The figures show that the soHC lies to the right of this interpolation line for all locations and transfer systems, except for the average climate location with UFH transfer system, where it is on the line. This indicates that the soHC has an improved form, so an improved dependency between T_{oda} and $T_{s,set}$, compared to the reference heating curves. Given that the reference curves were already expertly parametrized for the specific transfer system, this is a positive outcome. In the SCOP subplot, the MPC lies to the right of the interpolation line for all variations except the warm location with the UFH system, where it lies on the line. The same interpretation of the interpolation line of the SCOP subplot applies to the ΔW_{el} subplot. There, the MPC lies to the right of the line for all six variations, and the soHC lies to the right of it for all variations except the average location with the UFH system. In this variation, the soHC consumes practically the same electrical energy as HC high at a slightly higher thermal discomfort, so its form is not improved compared to the reference heating curves. This is caused by the training data of the system identification not containing enough information for this variation.

The MPC saves electrical energy compared to the heating-curve-based controllers because it raises $T_{s,set}$ only as far as it is required at



(a) Radiator transfer system (b) UFH transfer system

Fig. 6. Thermal discomfort (y-axis) over the SCOP (left subplot) and over the savings of electrical energy compared to HC high (right subplot) for the different controllers in the average location using (a) radiator-based heating; and (b) UFH systems.



(a) Radiator transfer system (b) UFH transfer system

Fig. 7. Thermal discomfort (y-axis) over the SCOP (left subplot) and over the savings of electrical energy compared to HC high (right subplot) for the different controllers in the cold location using (a) radiator-based heating; and (b) UFH systems.

each point in time. Solar radiation and internal gains are inputs of the process models of the MPC, whereas a heating curve depends on T_{oda} only. The MPC therefore reduces $T_{s,set}$ whenever solar radiation or internal gains contribute to the heat supply of the rooms. A heating curve cannot react to these disturbances and keeps $T_{s,set}$ at the level dictated by T_{oda}. This results in supply temperatures that are higher than needed and forces the thermostatic valves to close. The MPC sets $T_{s,set}$ such that the valve of the critical room remains completely open most of the time, as shown in Fig. B.5 in Appendix B.4. This effect is strongest in the cold location. Cold and sunny days occur frequently, whereas the cold days of the milder locations stay at low solar irradiation. This effect leads to the higher savings of the MPC in the cold location, because the MPC can exploit the solar radiation on the cold and sunny days to lower the supply temperature.

For the warm and average locations, the soHC achieves low thermal discomfort, comparable to the performance of HC high. For the warm climate location and the average climate location in combination with the radiator transfer system, the soHC is right of the interpolation line, which means better SCOP compared to the reference HCs for the same discomfort. For the average climate location in combination with the UFH heating system , the soHC falls on the interpolation line, which corresponds to a similar SCOP for a given discomfort compared to the reference HCs.

Overall, for the warm and average location, the low discomfort aligns with the design objective of the soHC, which was targeted to maintain a low room temperature deviation. The residual discomfort observed in these scenarios is attributable to transient thermal effects—such as sudden internal heat gains or fluctuating solar radiation—that the systems' PI controllers cannot compensate for due to the thermal inertia.

In the cold location, while the soHC is to the right of the interpolation lines, lower SCOPs and slightly increased thermal discomfort are observed overall. This occurs due to the maximum thermal capacity of the air-source heat pump, that decreases as the outdoor air temperature drops. During cold days, the available heating capacity of the heat pump becomes insufficient to meet the building's demand. This is especially pronounced in the UFH system; coupled with the high thermal inertia of the floor mass, the reduced heat pump capacity prevents the system from reaching and maintaining the desired $T_{ro,set}$.

For the cold climate location, the thermal discomfort is higher compared to the reference heating curves than it is for the other locations. Minimizing the thermal discomfort did not work as well as for the other locations. This is because of the limited range of T_{oda} and $T_{s,set}$ in the training data for the cold climate. In Fig. 4 it can be seen, that the standard deviation output of the GPR is higher around higher T_{oda}. This indicates lower quality of the soHC in this temperature range ($T_{oda} > 5°C$) which leads to higher thermal discomfort.

Overall, the typical day simulations confirm that the soHC provides low thermal discomfort across the three distinct climatic locations for both UFH and radiator systems. These results demonstrate that the proposed data-driven method is a robust and effective approach for au-

tomatically generating adapted heating curves. Due to the data-driven character of the method, no expert knowledge is needed for setting the heating curve. In this paper, we always chose to generate the soHC to provide low discomfort. By decreasing the target room temperature deviation in the soHC generation, the soHC can also be tuned to provide high SCOPs rather than low discomfort.

In summary, the soHC improves the shape of the heating curve in almost all variations. In the warm and average locations it saves between 0.03 % and 4.6 % of electrical energy at a comfort level comparable to HC high. In the cold location the savings are larger (9.2 % and 6.8 %), but thermal discomfort rises by 2.7 K h d^{-1} to 6 K h d^{-1}. The MPC saves between 1.0 % and 10.9 % and outperforms the soHC except in the warm location with UFH, where the soHC already provides a supply temperature close to the optimal level. It also attains the lowest or equal-lowest thermal discomfort in every variation, as only the MPC can lower T_s in response to solar irradiation and internal gains.

4. Discussion

The simulation results demonstrate the high potential of the self-optimizing heating curve. Still there are limitations and practical considerations that are to be addressed.

The building model utilized in this study incorporates a set of realistic disturbances, including solar radiation, wind speed, and internal heat gains. Similarly, the heat pump model is fitted to experimental data and employs realistic system-level control logic, such as PI controllers for the compressor speed and the secondary water pump. However, physical heat pumps possess internal control algorithms that can introduce additional dynamic disturbances not captured in simulation. Therefore, while the simulation environment provides a proof-of-concept, future work should include Hardware-in-the-Loop (HiL) testing and field deployments in real buildings to fully validate the robustness of the method under real operating conditions.

The practical implementation of this data-driven method assumes the presence of smart room thermostats. These devices are required to communicate real-time room temperatures and setpoints back to the central heat pump controller. While this necessitates a modern smart-home infrastructure, such networked systems are becoming standard in new heat pump installations. Furthermore, this hardware infrastructure provides a solution to one of the study's underlying assumptions: the prior knowledge of the critical room (Room 10 in this study). In a real-world application, the critical room can be determined empirically during an initial monitoring phase by tracking the smart room thermostats and identifying the room where the valve is most frequently at the 100 % open position.

In this study, the training data generation always started on the 1st of January, ensuring exposure to a wide distribution of outdoor air temperatures. In reality, a heat pump might be installed during the summer. A summer system identification would yield a training dataset lacking the cold temperature data necessary to construct a complete soHC for the winter. This scenario shows an advantage of utilizing GPRs as the data-driven model to construct the heating curve from. The GPR outputs a predictive variance alongside its mean prediction. If the model is forced to extrapolate into unobserved cold temperature ranges, the system can revert to a conservative baseline heating curve. As the heating season progresses and more diverse T_{oda} data is collected, the soHC can be continuously refined through online learning. This fallback strategy is not evaluated in the present study and is left for future work. The required duration of the system identification depends on the transfer system. It can be determined by continuing the identification until the predicted standard deviation of the GPR falls below a threshold.

In the present study, the supply temperature is the only controllable input of the GPR. For this problem, DiCE is functionally equivalent to simpler optimisation approaches. Its benefit is expected once additional controllable inputs are considered. Future work should therefore extend the framework to multiple controllable inputs.

The soHC results from a single identification phase followed by an offline optimisation. Over the life of a heat pump, the thermal behavior of a building changes, for instance through renovation or different user behavior, so that a heating curve parametrised at commissioning eventually no longer fits. If events that lead to decreased thermal comfort occur, the algorithm should re-enter identification mode, refit the GPR and derive an updated soHc. This workflow can run repeatedly over the lifetime of the system rather than once. Suitable thresholds to trigger a retraining are individual to the specific use-case.

The results of the MPC depend on the weighting of the comfort and the efficiency term in the cost function. A different weighting shifts the MPC along the trade-off between thermal discomfort and electricity demand, so the presented values represent one operating point out of a range that the MPC can reach. Furthermore, the MPC uses perfect forecasts of outdoor air temperature, solar radiation and internal gains. The reported savings are therefore an upper benchmark that an implementation with real forecasts will not fully reach.

Finally, this study does not attempt to provide generalized, absolute values for SCOP improvements or thermal discomfort reductions compared to real-world reference heating curves. In practice, manually parametrized heating curves vary and are dependent on the installing technician. Nevertheless, the results demonstrate that our approach eliminates the need for manual parametrization. The method autonomously generates a suitable heating curve, ensuring both efficient heat pump operation and thermal comfort without relying on human intervention. As Potthoff et al. [5] point out, heating curve parameters are in practice based on the experience of the installer or on rough estimates, so that no single parametrisation can be regarded as industry standard. A comparison against commercially available adaptive controllers was therefore not possible within this study. The two reference curves used here are theoretical heating curves that differ by a parallel shift, representing a conservatively low and a generously high heating curve.

5. Conclusion

This paper introduces a novel, data-driven method for the automated generation of heating curves for heat pump systems. The proposed approach was developed and evaluated using a high-order building model that features multiple thermal rooms and incorporates realistic environmental and operational disturbances, such as solar radiation and internal heat gains. To show the method's robustness, the system was simulated using typical days across three climatic locations, considering both UFH and radiator heat transfer systems.

The core of the methodology consists of utilizing historical operational data to train a data-driven model. Once trained, this model captures the thermal dynamics of the building and can predict expected room temperature setpoint violations based on the applied heating curve and outdoor weather conditions. By coupling this predictive model with an optimization algorithm, self-optimizing heating curve (soHC) that is tailored to the building's thermal characteristics can be generated.

The simulation results demonstrate that the generated soHCs perform well across all tested climates and transfer systems. Utilizing the soHC results in low thermal discomfort. The automatically generated soHCs achieve a performance level comparable or better to expertly configured, system-specific reference heating curves. The primary advantage of this approach lies in its complete automation, eliminating the need for manual parameterization by technicians.

An MPC serves as an upper benchmark. It consumes less electrical energy than the soHC in five of six variations and attains the lowest or equal-lowest thermal discomfort in every variation, since it can lower the supply temperature setpoint in response to solar irradiation and internal gains, which a heating curve cannot. Even though the MPC outperforms all heating curves in this study, this comes at the price of high development and modeling expense aswell as the need to perform online optimization.

Overall, this study shows that data-driven methods can be used to simplify the generation of heating curves. By autonomously adapting to the specific building and transfer system, the proposed method shows efficient heat pump operation and low thermal discomfort. The next step is to demonstrate the simulative findings from this study experimentally in an occupied building.

CRediT authorship contribution statement

Florian Will: Writing – original draft, Visualization, Software, Methodology, Investigation, Conceptualization; **Matthias Mersch:** Writing – original draft, Investigation; **Christian Vering:** Writing – review & editing, Supervision, Project administration, Methodology, Conceptualization; **Dirk Müller:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Anthropic Claude in order to improve grammar, spelling and readability of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We gratefully acknowledge the financial support by the Federal Ministry for Economic Affairs and Energy (BMWE), promotional reference 03EN1096B. This study was conducted in collaboration with the project partners Chair of Building Energy Systems and Heat Supply, Dresden University of Technology and Kermi GmbH.

Data availability

Data will be made available on request.

Appendix A. Additional methodology

A.1. Weather data

A.2. Data on typical days

Table A.1
Chronological list of typical days and their corresponding weight (number of days) for the warm, average, and cold locations.

Warm		Average		Cold	
Date	Weight	Date	Weight	Date	Weight
Jan 6	59	Jan 21	85	Feb 22	46
May 2,	48	Mar 28	85	Mar 4	20
Jun 30	37	Apr 25	35	Mar 28	65
Aug 14	41	Apr 30	38	Aug 21	22
Sep 20	60	Jul 13	39	Aug 25	43
Oct 15	28	Aug 22	43	Sep 28	39
No. 5	38	Sep 23	25	Sep 29	71
No. 22	54	Dec 19	15	No. 8	59

A.3. Calculation of KPIs from typical-day data

The SCOP is calculated as the ratio of the total weighted thermal energy provided by the heat pump to the total weighted electrical energy consumed over the year:

$$\mathrm{SCOP} = \frac{\sum_{i=1}^{8} w_i \int_{\mathrm{day}_i} \dot{Q}_{\mathrm{con},i}(t)\,dt}{\sum_{i=1}^{8} w_i \int_{\mathrm{day}_i} P_{\mathrm{el},i}(t)\,dt} \tag{A.1}$$

where:

- w_i is the weight, i.e., the number of days represented by cluster i.
- $\dot{Q}_{\mathrm{con},i}(t)$ is the condenser heat flow during typical day i.
- $P_{\mathrm{el},i}(t)$ is the electrical power demand during typical day i.

The Thermal Discomfort (TD) is calculated by integrating the positive deviations of the actual room temperature from the setpoint temperature (which correspond to lower actual room temperatures than the set room temperature). This integration is performed for all 12 rooms, weighted by the cluster days, and then averaged over the year:

$$\mathrm{TD} = \frac{1}{3600 \cdot 365} \sum_{i=1}^{8} w_i \sum_{r=1}^{12} \int_{\mathrm{day}_i} \max\left(0, T_{\mathrm{ro,set},r}(t) - T_{\mathrm{ro},r}(t)\right) dt \tag{A.2}$$

where:

- $T_{\mathrm{ro,set},r}(t)$ is the target setpoint temperature for room r.
- $T_{\mathrm{ro},r}(t)$ is the actual measured temperature in room r.
- The $\max(0, \dots)$ function ensures that only underheating (temperatures dropping below the setpoint) contributes to the discomfort metric.
- The factors 3600 and 365 convert the integrated seconds into hours and average the annual total over the days of the year, respectively.

A.4. Training data generation for the MPC

- **Variation 0** (constant valves, dynamic supply temperature): the supply temperature is governed by a PI controller using the temperature of room 10 (bathroom) as reference, as this room has the highest setpoint requirement and therefore determines the supply temperature demand. The valve opening degrees of all rooms are simultaneously held constant at $40\,\%$ and $100\,\%$.
- **Variation 1** (constant supply temperature, dynamic valves): the supply temperature setpoint is fixed at three levels while the room temperatures are controlled by individual PI controllers acting on the respective zone valves. The levels are adapted to the heat emission system, 40°C and 60°C for the radiator system and 25°C and 45°C for the underfloor heating, since exciting a low-temperature system at radiator level would generate data outside its attainable operating range.
- **Variation 2** (coupled dynamics): both the supply temperature, led by room 10, and the zone valves are controlled by PI controllers simultaneously, in order to record the coupled system behavior.

A.5. MPC optimal control problem

Let $\mathcal{J} = \{1, \dots, 12\}$ denote the set of considered zones and $\mathcal{K} = \{0, \dots, N-1\}$ the prediction horizon with $N = 30$ steps. The manipulated variables are the supply temperature setpoint $T_{\mathrm{sup,set}}$ and the thermostatic valve opening degrees $u_{\mathrm{val},j}$, collected in $\mathbf{u} = \left[T_{\mathrm{s,set}}, u_{\mathrm{val},1}, \dots, u_{\mathrm{val},12}\right]^{\top}$. At each time step the following optimal control problem is solved:

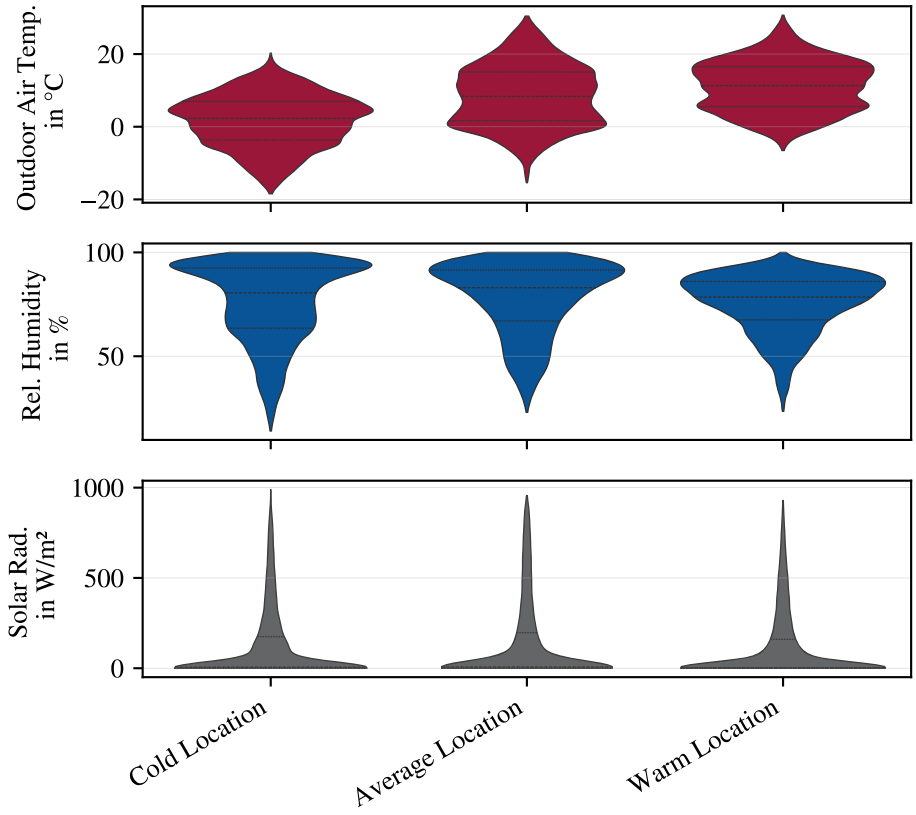


Fig. A.1. Outdoor air temperatures, relative humidities and solar radiations of the yearly weather data for all locations. The dashed lines show the 25 %, the 50 % and the 75 % quartiles.

$$\min_{\mathbf{u}} \quad \sum_{k\in\mathcal{K}} \Bigg(w_T\, T_{\mathrm{s,set},k} + w_{\Delta T}\, \Delta T^2_{\mathrm{s,set},k} + \sum_{j\in\mathcal{J}} \left(w_{\epsilon,j}\, \epsilon^2_{j,k} + w_v\, u_{\mathrm{val},j,k} + w_{\Delta v}\, \Delta u^2_{\mathrm{val},j,k} \right) \Bigg) \tag{A.3a}$$

$$\text{s.t.} \quad \Delta T_{\mathrm{ro},j,k} = f_{\mathrm{lr},j}\left(\tilde{x}_{j,k}, \tilde{u}_{j,k}, \tilde{d}_k\right) \tag{A.3b}$$

$$T_{\mathrm{ro},j,k+1} = T_{\mathrm{ro},j,k} + \Delta T_{\mathrm{ro},j,k} \tag{A.3c}$$

$$T_{\mathrm{lb},j} - \epsilon_{j,k} \le T_{\mathrm{ro},j,k} \le T_{\mathrm{ub},j} + \epsilon_{j,k} \tag{A.3d}$$

$$\epsilon_{j,k} \ge 0 \tag{A.3e}$$

$$T_{\mathrm{s,min}} \le T_{\mathrm{s,set},k} \le T_{\mathrm{s,max}} \tag{A.3f}$$

$$u_{\mathrm{val,min}} \le u_{\mathrm{val},j,k} \le u_{\mathrm{val,max}} \tag{A.3g}$$

$$\text{for all } j \in \mathcal{J},\ k \in \mathcal{K}.$$

Here $\Delta T_{\mathrm{s,set},k}$ and $\Delta u_{\mathrm{val},j,k}$ denote the change of the respective manipulated variable between two consecutive steps, and $\tilde{x}_{j,k}$, $\tilde{u}_{j,k}$, $\tilde{d}_k$ the lagged state, manipulated and disturbance inputs according to Table A.2. Eq. (A.3d) formulates the comfort bounds as soft constraints. The violation is penalized quadratically. Efficiency is addressed by the linear terms on $T_{\mathrm{s,set},k}$, which drives the supply temperature towards the lowest values compatible with comfort. Quadratic terms on the input changes suppress oscillations of the actuators. The resulting problem is solved using the solver IPOPT [34].

The comfort bounds are identical for all zones relative to their individual target temperature $T_{\mathrm{ro,set},j}$, with $T_{\mathrm{lb},j} = T_{\mathrm{ro,set},j}$ and $T_{\mathrm{ub},j} = T_{\mathrm{ro,set},j} + 3\,\mathrm{K}$. The weights and the bounds are summarized in Table A.3.

Table A.2
Inputs of the room-wise linear regression models, where $\dot{q}_{\mathrm{sol}}$ is solar irradiation and $\dot{Q}_{\mathrm{int},j}$ are roomwise internal gains. The lag denotes the number of considered time steps including the current one. At a step size of 10 min a lag of n covers the preceding $(n-1)\cdot 10$ min.

Category	Symbol	Description	Lag
State	$T_{\mathrm{ro},j}$	room temperature of zone j	2
Manipulated	$T_{\mathrm{s,set}}$	supply temperature setpoint	4
	$u_{\mathrm{val},j}$	valve opening degree of zone j	4
Disturbance	T_{oda}	ambient dry bulb temperature	2
	$\dot{q}_{\mathrm{sol}}$	global solar irradiation	2
	$\dot{Q}_{\mathrm{int},j}$	internal gains of zone j	2
Output	$\Delta T_{\mathrm{ro},j}$	change of room temperature of zone j	–

Table A.3
Weighting factors and bounds of the optimal control problem per heat transfer system.

Symbol	Description	Radiator	UFH
N	prediction horizon in steps	30	30
$w_{\epsilon,j}$	comfort weight zones	8	8
w_T	supply temperature	0.4	0.4
$w_{\Delta T}$	supply temperature change	1	1
w_v	valve opening	1	1
$w_{\Delta v}$	valve opening change	7	7
$T_{\mathrm{sup,min}}$	lower supply temperature bound in °C	20	20
$T_{\mathrm{sup,max}}$	upper supply temperature bound in °C	65	45
$u_{\mathrm{val,min}}$	lower valve bound	0	0
$u_{\mathrm{val,max}}$	upper valve bound	1	1

Appendix B. Additional results

B.1. Heating curve parameters

Table B.1
Configuration parameters for the heating curves of the four variations. The underlying equation is described in Eq. (6). All temperatures are given in Kelvin.

Parameter	Radiator		UFH	
	HC low	HC high	HC low	HC high
T_{HeaThr}	288.15	288.15	288.15	288.15
$T_{\mathrm{oda,n}}$	261.15	261.15	261.15	261.15
T_{ro}	297.15	297.15	297.15	297.15
$T_{\mathrm{s,n}}$	328.15	328.15	311.15	311.15
$T_{\mathrm{r,n}}$	318.15	318.15	304.15	304.15
n	1.24	1.24	1.0	1.0
T_{Add}	4.0	8.0	4.0	8.0

B.2. Sensitivity analysis on the heating curve algorithm

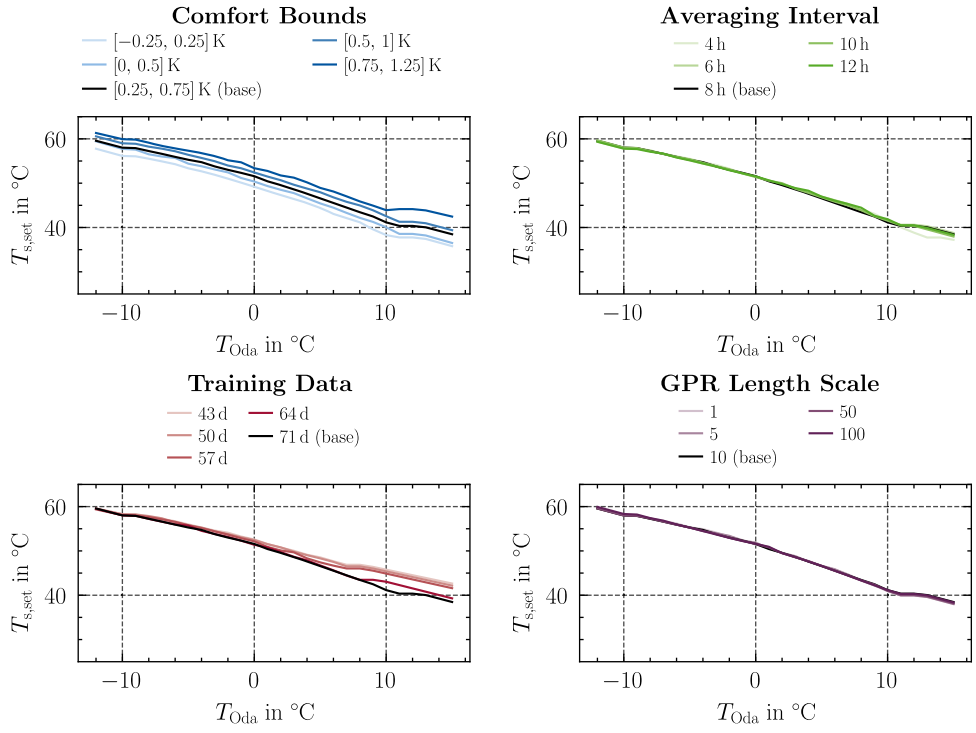


Fig. B.1. Sensitivity analysis of the soHC generation for the average location with radiator system. Each column varies one parameter while the others are held at their base setting (black): the comfort bounds ([0.25 K, 0.75 K]), the averaging interval of the training data (8 h), the amount of training data (71 days) and the length scale bound of the GPR (10). The subplots show the resulting heating curves.

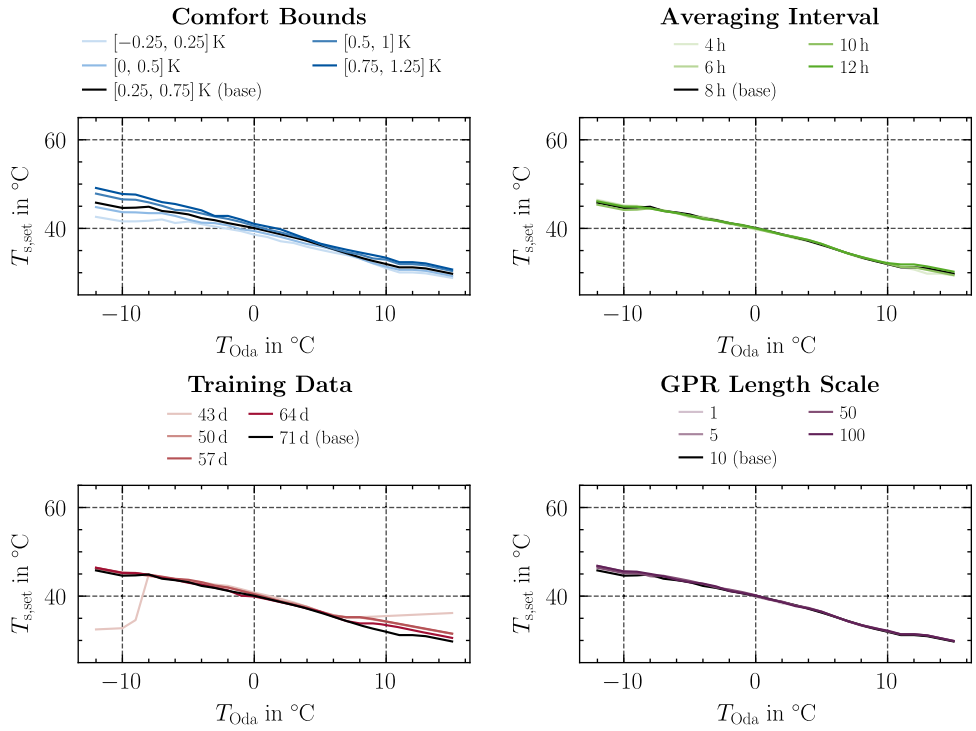


Fig. B.2. Sensitivity analysis of the soHC generation for the average location with UFH system. Each column varies one parameter while the others are held at their base setting (black): the comfort bounds ([0.25 K, 0.75 K]), the averaging interval of the training data (8 h), the amount of training data (71 days) and the length scale bound of the GPR (10). The subplots show the resulting heating curves.

B.3. Stochasticity of the DiCE genetic algorithm

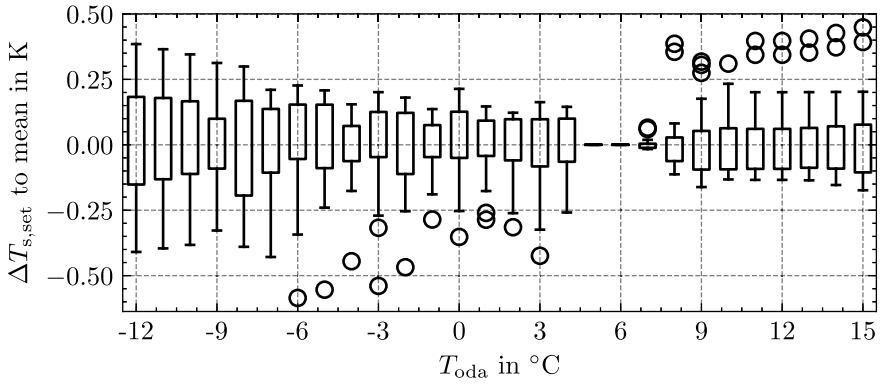


Fig. B.3. Deviation of $T_{s,set}$ of 25 soHCs, generated with different random seeds of the DiCE genetic algorithm, from their mean for the average location with the UFH system. Boxes show the interquartile range, whiskers extend to 1.5 times the interquartile range, and circles mark outliers.

B.4. Typical day simulation results

Table B.2
Evaluation of typical day simulations: Seasonal Coefficient of Performance (SCOP), thermal discomfort (TD) and savings of electrical energy ΔW_{el} with respect to HC high across different locations, controllers and transfer systems.

Climate	Controller	Radiator			UFH		
		SCOP in –	TD in Kh/d	ΔW_{el} in %	SCOP in –	TD in Kh/d	ΔW_{el} in %
Warm	HC low	3.67	22.73	10.10	4.25	21.18	5.17
	HC high	3.41	11.59	0.00	4.10	11.74	0.00
	soHC	3.40	8.07	0.84	4.18	13.00	4.59
	MPC	3.42	7.66	2.65	4.04	8.27	1.72
Average	HC low	3.17	22.40	8.74	3.95	24.94	6.00
	HC high	2.86	12.79	0.00	3.70	15.00	0.00
	soHC	2.89	12.00	0.42	3.74	16.48	0.03
	MPC	2.87	7.94	0.97	3.70	9.92	2.32
Cold	HC low	2.55	21.54	10.21	3.48	27.47	8.09
	HC high	2.29	12.13	0.00	3.21	20.13	0.00
	soHC	2.53	18.13	9.22	3.42	22.86	6.75
	MPC	2.54	12.13	10.85	3.35	14.89	8.21

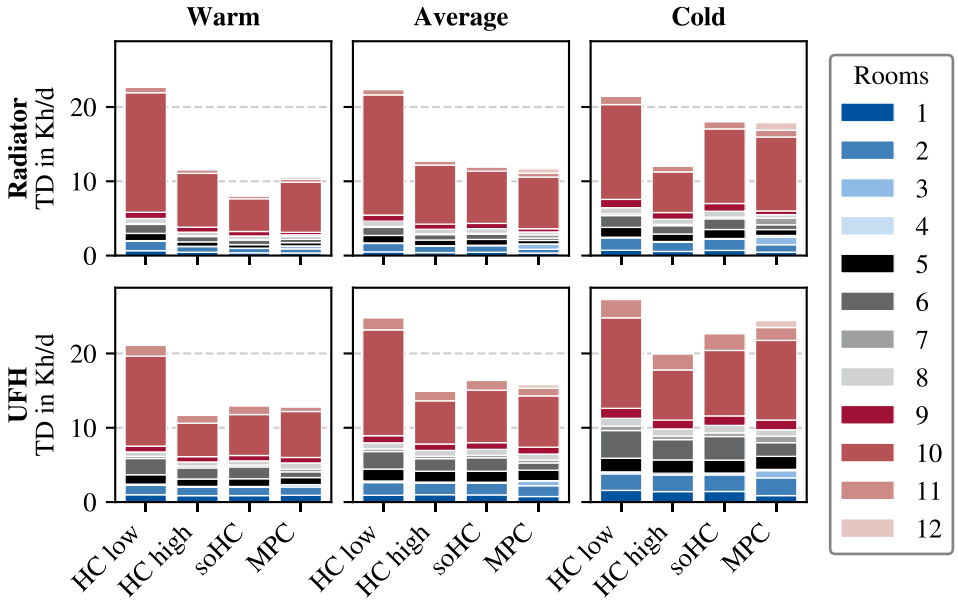


Fig. B.4. Detailed analysis of the thermal discomfort of individual rooms.

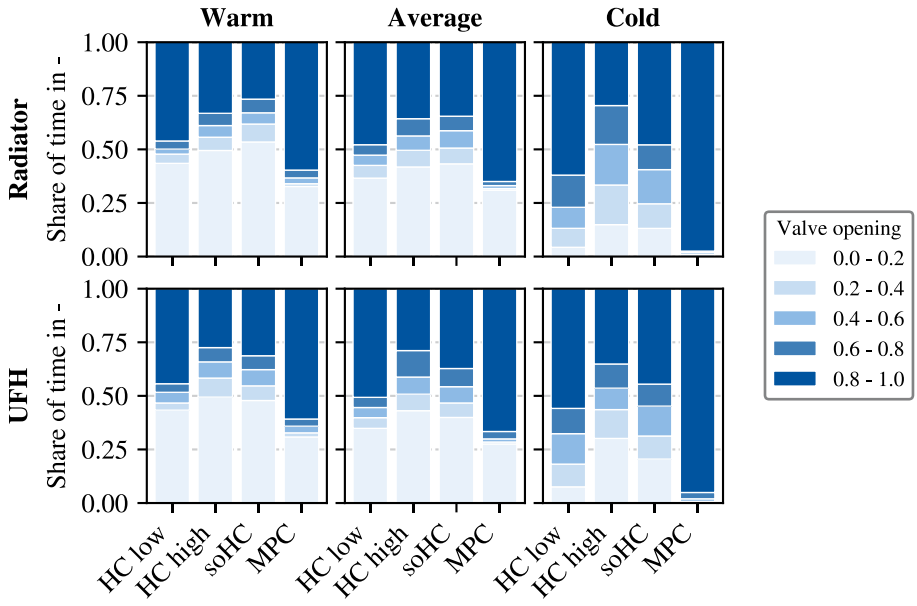


Fig. B.5. Histograms of the valve openings of the primary bathroom (Room 10) for all simulated variations.

References

[1] A.V. Olympios, P. Sapin, M. Mersch, A.M. Maghrabi, C.N. Markides, A review of recent progress in the design and integration of domestic heat pumps, Next Energy 5 (2024) 100163. https://linkinghub.elsevier.com/retrieve/pii/S2949821X24000681, https://doi.org/10.1016/j.nxener.2024.100163

[2] SwissEnergy, Set the Heating Curve Correctly, Technical Report, Federal Office of Energy (SFOE), Bern, Switzerland. https://pubdb.bfe.admin.ch/en/publication/download/9982

[3] OpenEnergyMonitor, Analysis of electrification of heat trial data, 2023. https://docs.openenergymonitor.org/heatpumpmonitor/eoh.html

[4] K. Huchtemann, D. Müller, Simulation study on supply temperature optimization in domestic heat pump systems, Buildi. Environ. 59 (2013) 327–335. https://linkinghub.elsevier.com/retrieve/pii/S0360132312002405, https://doi.org/10.1016/j.buildenv.2012.08.030

[5] U. Potthoff, T. Brudermueller, K. Hopf, F. Wortmann, Optimization of heating curves for heat pumps in operation: outdoor temperature ranges for energy-efficient heating curve shifts, Appl. Energy 389 (2025) 125725. https://linkinghub.elsevier.com/retrieve/pii/S0306261925004556, https://doi.org/10.1016/j.apenergy.2025.125725

[6] T. Cholewa, A. Siuta-Olcha, M. Bocian, E. Korzeniowska, B. Bomba, B. Babiarz, A.J. Werner-Juszczuk, C.A. Balaras, B. Sadowska, P. Rynkowski, A. Święcicki, L. Lichołai, J. Krasoń, P. Miąsik, M. Kłopotowski, D.A. Krawczyk, Adjustment of heating curve in space heating system to real energy needs as easily applicable and effective measure to increase energy efficiency in existing buildings, Energy Build. 352 (2026) 116856. https://linkinghub.elsevier.com/retrieve/pii/S0378778825015865, https://doi.org/10.1016/j.enbuild.2025.116856

[7] A. Ionesi, M. Jradi, J. Eric Thorsen, C. Veje, Simulation of an adaptive heat curve for automatic optimization of district heating installation, 2015. https://doi.org/10.26868/25222708.2015.2823

[8] P. Potočnik, B. Vidrih, A. Kitanovski, E. Govekar, et al., Analysis and optimization of thermal comfort in residential buildings by means of a weather-controlled air-to-water heat pump, Build. Environ. 140 (2018) 68–79. https://linkinghub.elsevier.com/retrieve/pii/S036013231830310X, https://doi.org/10.1016/j.buildenv.2018.05.044

[9] P. Potočnik, E. Govekar, Adaptive optimization of heating curves in buildings heated by a weather-compensated heat pump, Sci. Technol. Built Environ. 25 (10) (2019) 1380–1393. https://doi.org/10.1080/23744731.2019.1616984

[10] D. Rolando, H. Madani, Smart Control Strategies for Heat Pump Systems, Technical Report, 2018. https://varmtochkallt.se/wp-content/uploads/Projekt/EffsysExpand/P18_Project_Report_final_reviewed.pdf

[11] D. Rolando, H. Madani, Improving the traditional heat pump system control through prediction of daily solar radiation, Heat Pumping Technol. Mag. (2022). Publisher: Heat Pump Centre, https://doi.org/10.23697/FYFZ-JD71

[12] H.T. Le, T. Scheunemann, Adaptation of the heating curve for heating systems, 2025. https://publications.ibpsa.org/conference/paper/?id=bs2025_1141, https://doi.org/10.26868/25222708.2025.1141

[13] E.B. Agyekum, B. Tarawneh, S. Praveenkumar, F. Odoi-Yorke, F.L. Rashid, et al., Review of emerging smart techniques and their role in enhancing energy efficiency in residential and commercial buildings, Energy Explor. Exploit. 44 (5) (2026) 2996–3036. https://doi.org/10.1177/01445987261470169

[14] S. Göbel, F. Will, P. Stoffel, C. Vering, D. Müller, Optimizing energy efficiency in buildings: experimental findings on the role of heat pump control interfaces in model predictive control, Energy Conver. Manage. X 28 (2025) 101306. https://linkinghub.elsevier.com/retrieve/pii/S2590174525004386, https://doi.org/10.1016/j.ecmx.2025.101306

[15] F. Will, J. Klingebiel, C. Vering, D. Müller, Tailoring the Heat Pump System Controller to the Building: Assessment of Adapted Heating Curves using Data-Driven Methods, Purdue, United States, 2024.

[16] F. Wüllhorst, L. Maier, D. Jansen, L. Kühn, D. Hering, D. Müller, et al., BESMod - A Modelica Library providing Building Energy System Modules, 2023, pp. 9–18. https://ecp.ep.liu.se/index.php/modelica/article/view/619, https://doi.org/10.3384/ECP211869

[17] F. Jorissen, G. Reynders, R. Baetens, D. Picard, D. Saelens, L. Helsen, et al., Implementation and verification of the IDEAS building energy simulation library, J. Build. Perfor. Simul. 11 (6) (2018) 669–688. https://doi.org/10.1080/19401493.2018.1428361

[18] M. Wetter, W. Zuo, T.S. Nouidui, X. Pang, et al., Modelica buildings library, J. Build. Perfor. Simul. 7 (4) (2014) 253–270. https://doi.org/10.1080/19401493.2013.765506

[19] M. Wetter, D. Blum, J. Hu, Modelica IBPSA library v1, 2019. https://www.osti.gov/doecode/biblio/25252

[20] L. Maier, D. Jansen, F. Wüllhorst, M. Kremer, A. Kümpel, T. Blacha, D. Müller, et al., AixLib: an open-source modelica library for compound building energy systems from component to district level with automated quality management, J. Build. Perfor. Simul. (2023) 1–24. https://doi.org/10.1080/19401493.2023.2250521

[21] C. Nytsch-Geusen, J. Huber, M. Ljubijankic, J. Rädler, Modelica buildingsystems - eine modellbibliothek zur simulation komplexer energietechnischer Gebäudesysteme, Bauphysik 35 (1) (2013) 21–29. https://doi.org/10.1002/bapi.201310045

[22] A. Constantin, R. Streblow, D. Müller, et al., The modelica HouseModels library: presentation and evaluation of a room model with the ASHRAE standard 140, in: Proceedings of the 10th International Modelica Conference, March 10–12, 2014, Lund, Sweden, Linköping Electronic Conference Proceedings, Linköping University Electronic Press, 2014, pp. 293–299. 1 citations (Crossref) [2021-03-08]. https://doi.org/10.3384/ECP14096293

[23] K. Xanthopoulou, P. Mehrfeld, P. Kangowski, F. Lanzerath, D. Müller, Validation of a building model as part of the AixLib modelica library for dynamic plant and building performance simulations, Energy and Buildings 250 (2021) 111248. https://linkinghub.elsevier.com/retrieve/pii/S0378778821005326 https://doi.org/10.1016/j.enbuild.2021.111248

[24] D.I.f. Normung, DIN EN 12831-1: Energetische Bewertung von Gebäuden - Verfahren zur Berechnung der Norm-Heizlast - Teil 1: Raumheizlast, Modul M3-3; Deutsche und Englische Fassung prEN 12831-1:2025, 2025.

[25] D.I.f. Normung, DIN EN 15450:2007, Heizungsanlagen in Gebäuden - Planung von Heizungsanlagen mit Wärmepumpen; Deutsche Fassung EN 15450:2007, 2007. https://www.beuth.de/de/norm/din-en-15450/98862901, https://doi.org/10.31030/9854436

[26] F. Wüllhorst, D. Jansen, P. Mehrfeld, D. Müller, et al., A modular model of reversible heat pumps and chillers for system applications, in: Proceedings of 14th Modelica Conference 2021, Linköping, Sweden, September 20–24, 2021, 2021, pp. 561–568. https://ecp.ep.liu.se/index.php/modelica/article/view/234, https://doi.org/10.3384/ecp21181561

[27] M. Lämmle, C. Bongs, J. Wapler, D. Günther, S. Hess, M. Kropp, S. Herkel, et al., Performance of air and ground source heat pumps retrofitted to radiator heating systems and measures to reduce space heating temperatures in existing buildings, Energy 242 (2022) 122952. https://linkinghub.elsevier.com/retrieve/pii/S0360544221032011, https://doi.org/10.1016/j.energy.2021.122952

[28] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: machine learning in python, J. Mach. Learn. Res. 12 (2011) 2825–2830.

[29] R.K. Mothilal, A. Sharma, C. Tan, Explaining machine learning classifiers through diverse counterfactual explanations, in: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 2020, pp. 607–617.

[30] P. Stoffel, D. Abel, D. Müller, Learning Strategies for Data-Driven Model Predictive Control of Building Energy Systems, Ph.D. thesis, E.ON Energy Research Center, 2024. Number: RWTH-2024-08179, https://publications.rwth-aachen.de/record/992360, https://doi.org/10.18154/RWTH-2024-08179

[31] P. Stoffel, P. Henkel, M. Rätz, A. Kümpel, D. Müller, et al., Safe operation of online learning data driven model predictive control of building energy systems, Energy AI 14 (2023) 100296. https://linkinghub.elsevier.com/retrieve/pii/S266654682300068X, https://doi.org/10.1016/j.egyai.2023.100296

[32] Bundesamt für Bauwesen und Raumordnung (BBR), Deutscher Wetterdienst (DWD), Ortsgenaue Testreferenzjahre von Deutschland für mittlere, extreme und zukünftige Witterungsverhältnisse, Offenbach, Offenbach, 2017. https://www.bbsr.bund.de/BBSR/DE/forschung/programme/zb/Auftragsforschung/5EnergieKlimaBauen/2013/testreferenzjahre/try-handbuch.pdf?_blob=publicationFile&v=

[33] L. Kotzur, P. Markewitz, M. Robinius, D. Stolten, et al., Impact of different time series aggregation methods on optimal energy system design, Renew. Energy 117 (2018) 474–487. https://linkinghub.elsevier.com/retrieve/pii/S0960148117309783, https://doi.org/10.1016/j.renene.2017.10.017

[34] A. Wächter, L.T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Math. Program. 106 (1) (2006) 25–57. https://doi.org/10.1007/s10107-004-0559-y