Energy Reports 13 (2025) 1054–1068

Contents lists available at ScienceDirect

Energy Reports

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

Behaviour tree based control strategies for resilient heat pump operation in residential buildings☆

Piet Urban *, Peter Klement, Sunke Schlüters, Patrik Schönfeldt

DLR Institute for Networked Energy Systems, Carl-von-Ossietzky-Str. 15, 26129, Oldenburg, Germany

ARTICLE INFO

Keywords:
Behaviour tree
Decision tree
CART
Optimisation
MILP
Heat pump
Energy management
Machine learning

ABSTRACT

Behaviour trees are a proven concept in the creation of complex task-switching control and artificial intelligence for robotic systems and non-player characters in the computer games industry. Requirements such as flexibility, maintainability, reusability of functionalities or expandability also apply to the control of decentralised energy systems. Despite this, there is a noticeable research gap regarding the application of behaviour trees in that sector. Based on a foundational heating system, including thermodynamic modelling of a part-load capable heat pump with *TESPy*, tree structures for its control are created using the Python library *py_trees* for implementation. With a view to minimising the annual operational performance indicators electricity price and CO_2 emissions, which reflect the optimal use of renewable shares, several control strategies are compared. We identify and illustrate the principal limitations of decision trees, mixed-integer linear optimisation performed with *oemof-solph*, as well as a classic rule-based approach. The proposed higher-level behaviour tree combines the strengths of such approaches whilst pursuing the additional target of reducing the start-up and associated wear of the heat pump without significantly increasing the computation time.

1. Introduction

Decentralisation and digitisation of our energy system pose major challenges to us, notably the rise in actors and the shift from consumers to prosumers, leading to increasingly complex systems. Furthermore, the inherent intermittency and unpredictability of renewable energies present fundamental issues. This creates a significant demand for flexibility in technologies and their control options, emphasising the importance of system interpretability, along with considerations for maintainability and expandability (Perger et al., 2022; Razmi and Lu, 2022).

Proposing a potential solution to address above requirements of future system and component control, the concept of BTs stands out. BT enable a structure for switching between different tasks in an autonomous agent. They are characterised by high modularity, particularly in the reusable nature of individual subtrees or behaviours, and clarity, along with a high reactivity to changing environmental conditions (Biggar et al., 2020b). In this context, behaviour trees established themselves since the mid-2000s in the computer game industry for non-player character control. Pioneered by the works of Mateas and Stern (2002) as well as Isla (2005) and Champandard (2007), the concept has seen numerous extensions for creating artificial intelligence (AI) behaviour (Rabin, 2013; Colledanchise and Ögren, 2018). Especially the possibility to easily extend behavioural patterns offers a great advantage over the otherwise widespread control with finite state machines (FSMs).

With an increasing number of scientific publications outside industry, further application areas have been opened up for the use of BTs. Example fields, beyond the gaming realm, are humanoid robotics (Marzinotto et al., 2014), automation of production (Guerin et al., 2015) and support in brain surgery (Hu et al., 2015).

Only recently, their utility has also been proposed for application in the energy sector, more specifically in the control of smart grids (Perger et al., 2022) and the operation of microgrids (Jingsong, 2023). Again, the focus is on ensuring stable and economic operation under a wide range of environmental and working conditions. As the demand for control solutions in decentralised energy systems continues to grow, the theoretically well-researched and in diverse fields practically proven concept of behaviour trees emerges as a promising candidate.

☆ This document is the results of the research project ENaQ (project number 03SBE111) funded by German Federal Ministry for Economic Affairs and Climate Action (BMWK) and the Federal Ministry of Education and Research (BMBF) and the research project WWNW (project number 03SF0624) funded by German Federal Ministry for Economic Affairs and Climate Action (BMWK).

* Corresponding author.
E-mail addresses: piet.urban@dlr.de (P. Urban), peter.klement@dlr.de (P. Klement), sunke.schlueters@dlr.de (S. Schlüters), patrik.schoenfeldt@dlr.de (P. Schönfeldt).

https://doi.org/10.1016/j.egyr.2024.12.039
Received 4 September 2024; Received in revised form 3 December 2024; Accepted 15 December 2024
Available online 2 January 2025

Nomenclature

AI	artificial intelligence
BT	behaviour tree
CART	Classification and Regression Trees
DT	decision tree
ENaQ	Energetisches Nachbarschaftsquartier
ES	energy system
FSM	finite state machine
HP	heat pump
KPI	key performance indicator
MILP	mixed integer linear programming
MPC	model predictive control
oemof	open energy modelling framework
TES	thermal energy storage
TESPy	Thermal Engineering Systems in Python
A_{TES}	surface area (TES)
DHI	diffuse horizontal irradiance
DNI	direct normal irradiance
E_{TES}	energy content (TES)
E_{th}	thermal energy
$P_{\mathrm{HP,el}}$	electrical power input (HP)
$Q_{\mathrm{HP,max}}$	nominal, i. e. maximum heat output (HP)
$\varepsilon_{\mathrm{HP}}$	coefficient of performance, COP (HP)
Q_{dem}	heat demand
T_{max}	upper temperature limit (TES) with respect to T_{ref}
T_{min}	lower temperature limit (TES) with respect to T_{ref}
T_{ref}	reference temperature (TES, $E_{\mathrm{TES}} = 0$)
T_{TES}	actual temperature (TES)
V_{TES}	volume (TES)
ΔT_{TES}	temperature spread (TES)
Δt_{ex}	period between two control activation steps
Δt_{upd}	period between regressor updates
β	lossrate (TES)
$\dot{Q}_{\mathrm{loss}}$	heat loss flow (TES)
ρ	density
τ	time increment of loss (TES)
θ_{amb}	ambient temperature
$\varepsilon_{\mathrm{HP}}$	coefficient of performance, COP (HP)
$a_{\mathrm{el,grid}}$	amount of electrical energy (grid supply)
c_{E}	normalised CO_2 emissions associated with heating
c_{P}	normalised heating price
$c_{\mathrm{el,P}}$	specific price (grid supply)
c_{p}	specific heat capacity
n_{on}	start ups per day (HP)
r_{xy}	correlation
u	thermal transmittance (TES)

One relevant application alongside the above-mentioned can be found in the electrification of domestic heating systems using HPs. In view of the growing proportion of renewable energies, smart energy management seeks to reduce not only the price of heat provision but also associated emissions. Moreover, a component-friendly mode of operation should be aimed for in order to reduce environmental costs. The described advantages of BTs motivate to use them for such a purpose.

Our notable contributions, expanding on previous research, are as follows.

- Investigations on the basis of a part-load capable HP model including domestic hot water in addition to space heating,
- Validation of a dynamic DT strategy, its sensitivity to demand profiles, energy standard and heating season,
- Proposal of a superordinate BT control strategy, combining strengths of MILP operation optimisation, DT and static, rule-based hysteresis control,
- Consideration of CO_2 emissions and number of HP start ups as performance indicators in addition to electricity price.

The rest of this study is structured as follows: Section 2 contains relevant information on the methodology, including a description of the energy system model and data basis as well as a detailed explanation of the control structures applied. In Section 3, simulation results of different control strategies from two sample years are shown, comprising the influence of varying demand characteristics on HP operation. Finally, the transition to a resilient approach based on BT is justified here. Key findings and further research opportunities are concluded in the last Section 4.

2. Methodology

This section first presents project information, data basis and the energy system model. An important part is concerned with the optimisation tool used and the development of specific control structures based on BT.

2.1. Background - Energetisches Nachbarschaftsquartier

Simulation data from the project Energetisches Nachbarschaftsquartier (ENaQ) was used as a basis. Within ENaQ possibilities for energy distribution between producers and consumers in the immediate vicinity have been investigated (Grimm et al., 2021; Schmeling et al., 2022; Schönfeldt et al., 2022). The project is linked to the construction of a real neighbourhood on the former airbase in Oldenburg. Following considerations comprise a total of approximately 140 residential units.

2.2. Data basis

The simulation process that was conducted to determine realistic heat load profiles and specific emissions is described by Schmeling et al. (2022). Note that both seasonal space heating and rather season-independent domestic hot water consumption were taken into account. Two distinct sets of hourly data have been used for validating resilience of subsequent control systems:

- 2017: lower efficiency standard with a more balanced daily profile, renewable integration of 36.2%,
- 2020: higher efficiency standard but with larger load peaks, renewable integration of 47.7%, see Fig. A.15.

An overview of the full demand profiles is given in Fig. 1 (a) and of the specific CO_2 emissions $c_{\mathrm{el,E}}$ of the German energy market in Fig. 1 (b). A reduction in cumulative demand from 642 to 389 MW h a^{-1} for the later year can be observed. This corresponds to a daily heat consumption of 12.6 kW h resp. 7.6 kW h, which is realistic for energy-efficient buildings (Wirth, 2023).

Apart from the variation in the assumed energy standards and profiles, an explanation for the significant difference in cumulative heating demand can be given with respect to the applied weather data Table 1 and Fig. 1(c) (Deutscher Wetterdienst, 2023, Measurement station Bremen, ID: 691). Lower temperature extrema and a mean difference of approximately one degree Celsius, but also the higher overall irradiation were assumed to have an influence on the heating behaviour.

Furthermore a clear shift of emissions can be seen, which is confirmed by the trend of the German electricity mix. The share of renew-

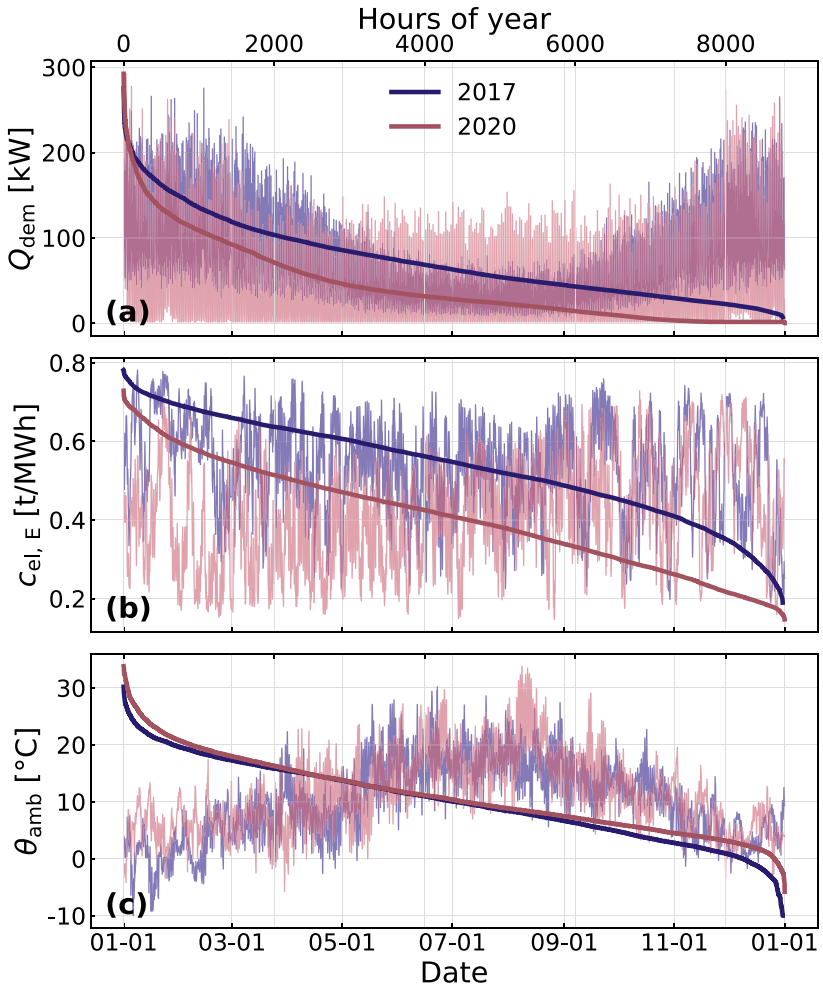

Fig. 1. Hourly data basis: cumulative spacial heating and domestic hot water demands (a), specific CO_2 emissions of the German electricity mix (b) and ambient temperature (c).

Table 1
Overview of utilised weather data: ground level ambient temperature, direct and diffuse solar irradiance (Deutscher Wetterdienst, 2023), specific CO_2 emissions reflecting the share of renewable energy production in Germany (Schmeling et al., 2022).

Data	Minimum		Maximum		Mean	
	2017	2020	2017	2020	2017	2020
θ_{amb} [°C]	−10.0	−5.8	30.2	33.8	10.1	11.1
DNI [W m^{-2}]	–	–	831.0	794.0	46.5	63.6
DHI [W m^{-2}]	–	–	614.0	536.0	64.4	61.1
$c_{el,E}$ [t/MWh]	0.19	0.15	0.78	0.73	0.53	0.41

able energies increased significantly in the later year (Umweltbundesamt, 2022). Using electricity at times when these are low favours the optimal use of renewable energies and increases their share in the heating sector. The fixed specific prices $c_{el,P}$ used for subsequent calculations, on the other hand, hardly differ, 2017: 29.28 ct kW h^{-1}, 2020: 31.81 ct kW h^{-1} (Deutschland, 2023).

2.3. Energy system and component modelling

The energy system (ES) under consideration, Fig. 2 consists of a grid connected HP, a TES and a fixed heating demand Q_{dem} that can be covered by either of them. Flexibility is only allowed in the form of excess, e.g. in case of a full TES. Through the introduction of storage, it is conceptually the smallest system with a sensible application of energy management.

Our choice of sizing was mainly based on an example from the work of Schmeling et al. (2022) where a case study on the above described

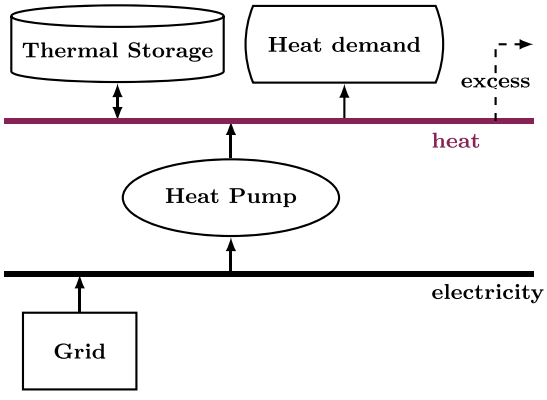

Fig. 2. Bus representation of the considered Heating System.

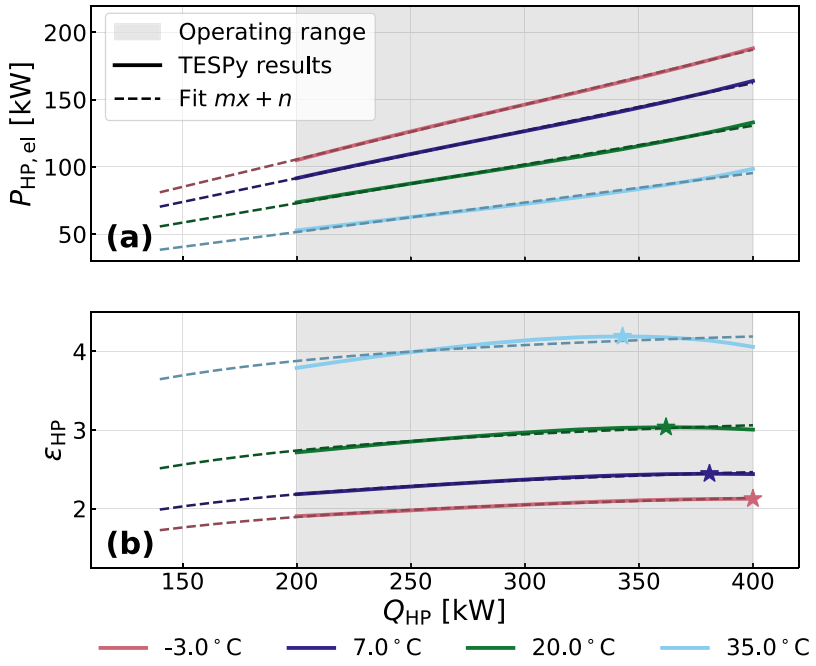

Fig. 3. HP characteristics: It can be seen, that a linearisation with offset of the compressor power (a) approaches the TESPy results well, while deviation of the actual ε_{HP} maximum ($\star$) gets larger for higher ambient temperatures (b).

Table 2
Heat pump specification used to obtain $\varepsilon_{HP}(\theta_{amb}, Q_{HP})$ values from the TESPy model.

Type	Air–water
Working fluid	R290
Feed temperature	67.5 °C
Return temperature	62.5 °C
θ_{amb} range	−10 – 35 °C
Q_{HP} range	0.2 – 0.4 MW

Table 3
Specifications for the TES.

Medium	Water
c_p	4.19 kg kJ K^{-1}
V_{TES}	20 m^3
ΔT_{TES}	35 K
β (τ = 1 h)	1.0%

neighbourhood was conducted. The size of the HP and the TES were adopted directly: $Q_{HP,max}$ = 400 kW, V_{TES} = 20 m^3. However, note that it is only one solution for the optimisation of sizing in that paper. It is thus definitely a large scale system, which is to be seen in distinction to single-family house applications. Consequently, the influence of individual variation in demand patterns is already weakened, making the system statistically more stable against supply and energy availability fluctuations. Further assumptions are made below.

2.3.1. Heat pump

In simplified terms, the HP is merely a converter between electricity part and heating part system as seen in Fig. 2. A key characteristic of the heat pump is its coefficient of performance $\varepsilon_{HP} = Q_{HP}/P_{HP,el}$, which is primarily temperature dependent. To obtain such dependency, component based modelling with TESPy by Witte and Tuschy (2020) was carried out. The full thermodynamic model corresponds to the components of Fig. A.16. Associated input properties are concluded in Table 2.

Now, in addition to the influence of the ambient temperature, part load behaviour of the single components can be taken into account for the design range of operating power, Fig. 3. Note that TESPy generic data was used for the characteristic line of the isentropic efficiency of the compressor. To use this model for a MILP optimiser a linearisation needed to be performed. It was found that fitting with a negative offset (corresponding to a constant base load of the system) matches the TESPy results especially for lower temperatures. Efficiency differences of up to 13% for constant ambient temperatures are reproduced very well, while the position of the maximum deviates for higher temperatures. Linearisation without offset leads to a significantly less computationally intensive optimisation (see Fig. 4), however corresponds less well to the model.

2.3.2. Thermal energy storage

The energy content of the TES can be calculated with respect to the heat capacity of the storage medium using the specifications from Table 3, $E_{th} = \rho c_p V_{TES} \Delta T_{TES}$ = 814.7 kW h.

A fairly basic approach for the storage loss was adopted, according to which the energy content of a subsequent time step $E_{TES}(t + \Delta t)$ can be calculated as follows, Eq. (1).

$$E_{TES}(t + \Delta t) = E_{TES}(t)(1 - \beta)^{\frac{\Delta t}{\tau}} \tag{1}$$

However, fixed absolute and relative losses were neglected. A time sensitive loss rate β results in exponential behaviour, approximating heat flow losses through the surface of the water tank $\dot{Q}_{loss} = uA_{TES}(T_{TES} - T_{ref})$. Assuming a number of smaller storage units and u of 1.21 (Cruickshank and Harrison, 2010), β was estimated at an upper 1.0% for τ = 1 h.

2.3.3. Performance indicators

Normalised price c_P and emissions c_E are calculated as key performance indicators (KPIs) in the following. As the introduced ES only procures electricity from the grid as single external source, those values can be calculated from $c_{el,E}$ and $c_{el,P}$ respectively. The amount of utilised electrical energy per time step is denoted by $a_{el,grid}$.

$$C_i = \sum_t C_i(t) = \sum_t c_{el,i}(t) a_{el,grid}(t) \tag{2}$$

$$c_i = C_i / \sum_t Q_{dem}(t) \tag{3}$$

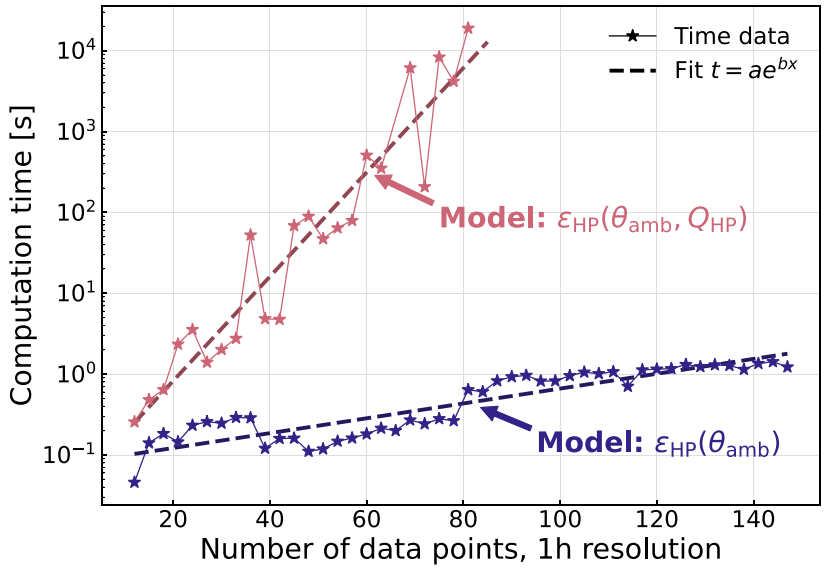

Fig. 4. Exemplary measurement of the computation time for two different HP models. $\varepsilon_{HP}(\theta_{amb})$ vs. $\varepsilon_{HP}(\theta_{amb}, Q_{HP})$.

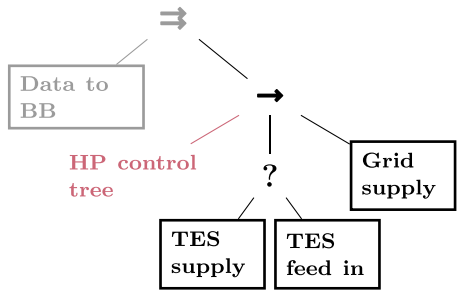

Fig. 5. Superordinate BT structure for the conducted simulations. The HP control node is to be extended by various subtrees.

Another relevant quantity will be the average number of starts of the HP per day n_{on}. This indicator is particularly important as the overall service life of the HP largely depends on that of the compressor, a. o. its number of cycles (Aguilera et al., 2022).

2.4. Model based optimisation

The ES model above gives rise to a mixed-integer linear optimisation problem, when trying to minimise the described costs. In this work the open energy modelling framework (oemof) (Hilpert et al., 2018) package *solph* by Krien et al. (2020) was applied.

With a view to computational costs, i. e. Fig. 4, the performed optimisations were split into clusters of either 24 h or 72 h and processed separately. This exceeds the cycle time of the TES. In addition, the assumption of perfect foresight for optimisation in this temporal range should be somewhat closer to reality, e. g. in a model predictive control (MPC) approach. The increase in complexity when optimising HP operation with varying part load efficiency as done here vs. constant ε_{HP} at a given ambient temperature can be seen in Fig. 4.

Note that, in contrast, for the simulations of BT based strategies frequencies ranged between 700 and 900 execution steps per second. This is far beyond the requirements of real-time applications with periods of several minutes between two decisions (for the operation of a comparably slow reacting thermal system like HPs). As soon as optimisations are incorporated during the creation of BT, this advantage partially vanishes. However, an error due to unpredictably long optimisation durations can be prevented, for example if the BT is updated only upon completion of the optimisation of past data, see Section 2.5.4.

2.5. Derivation of control structures

As described in the introduction, BT possess the main advantages of being reactive and easily expandable (Biggar et al., 2020a; Colledanchise and Ögren, 2018). The general framework and control structures realised with BT should now be introduced. There is a wealth of documentation on the basic working principle of BT (Rabin, 2013; Marzinotto et al., 2014; Colledanchise and Ögren, 2018; Colledanchise and Natale, 2021). An overview of the node types and their signal processing can be found in Table A.9. As commonly used, the organisation of following graphic representations corresponds to execution from left to right and top (root of the tree) to bottom. Final implementation was carried out with the help of the Python package *py_trees* (Stonier, 2023).

2.5.1. Superordinate behaviour tree structure

The superordinate control system sequentially (→) addresses the three ES components, see Fig. 5. As decisions are only to be made actively for the HP, the following passive update of the TES and grid is realised here. However, in actual applications those leaves could conveniently be used e. g. for status checks.

A regular query of environment and system variables is achieved by an upstream parallel node (⇉) that checks for new data but otherwise has no influence on the main activities. Therefore the execution frequency of the BT is not limited to the data frequency, hence the last existing data point is used to perform a control step. Information exchange between different parts, i. e. behaviours of the tree runs via a simple key/value storage, the so called blackboard. Typical execution frequencies of up to 500 Hz in other domains (e.g. robot operating systems) are easily achieved computationally, but are more relevant to simulation time than to the actual application.

In the following only the active HP control will be discussed and expanded. The complete tree structures then result from insertion into the associated node in Fig. 5. This possibility to easily extend the structure illustrates the mayor advantage compared to classical FSM. Behaviour nodes can be edited independently to a great extent. Resilience and efficiency are thus increased on another dimension. By sustaining clarity in increasingly complex systems (a. o. component wise) development and maintenance effort are reduced in particular, as described for example by Iovino et al. (2022).

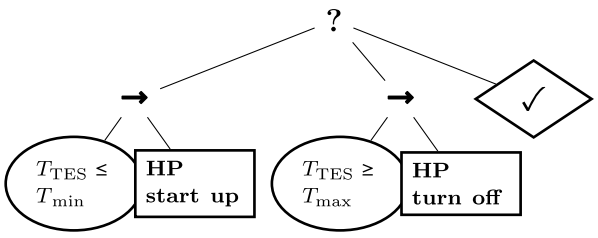

Fig. 6. Realisation of a simple HP hysteresis controller using a BT.

2.5.2. Hysteresis control

The BT realisation of a simple hysteresis control for the HP, based on a fixed set of temperature thresholds T_{min}, T_{max} is shown in Fig. 6.

As soon as one of the lowest level conditions returns a success feedback, the respective HP action (start up/ turn off) is performed. For a TES temperature within the predefined boundaries current operation is to be maintained. This can be achieved using a success decorator (✓).

2.5.3. Decision tree control

Another option resembling the dynamic operating strategy of Luo et al. (2021) is the replacement of the HP control node by a DT. In contrast to hysteresis, this allows for decisions to be made on partial loads increasing the system flexibility. Here, on the other hand the major drawback of missing information back flow in DTs is emphasised. Without observing further boundaries, the coverage of demand cannot be guaranteed. One option to avoid this is the introduction of safety thresholds, such that the decision making is limited with respect to demand Q_{dem} and actual storage content E_{TES}, i. e. temperature, Eq. (4).

$$Q_{HP} = \min \left\{ Q_{dem} + \frac{E_{TES,tot} - E_{TES}}{\Delta t}, \max \left\{ Q_{HP,DT}, Q_{dem} - \frac{E_{TES}}{\Delta t} \right\} \right\} \tag{4}$$

For DTs there exist very mature training algorithms, as concluded by Somvanshi et al. (2016). One that is able to create a DT regressor for continuous data is the so called Classification and Regression Trees (CART) algorithm introduced by Breiman et al. (1984). Here the python package *scikit-learn* (Pedregosa et al., 2011) was used for implementation with a mean-squared error loss function.

As features for training of the DT regressor we used storage content and heating demand, hour of the day and ambient temperature. For the analysis of system emissions, the electricity day-ahead price was furthermore taken into account (correlation with emissions of the energy market 2017/2020: $r_{xy} \approx 0.7$). Optimisation results from previous, i. e. known time periods could then be used for the regression. Details are presented in Section 3.

Exclusive use of a DT for HP control causes difficulties as soon as higher activation frequencies are required (e. g. $n_{on} > 130$ for $\Delta t_{ex} = 60$ s). Rapid fluctuation between tasks as a downside of reactivity is well known from other applications (Biggar et al., 2020b).

Another disadvantage of DT is being addressed in the advancement to BT: The inherent feedback of leaf nodes of the latter, in theory largely improves error handling and backup processes. This is broadly discussed e. g. by Colledanchise and Ögren (2018), however, plays a subordinate role in the simulation of fault-free components as analysed here.

2.5.4. Combined operating strategy

We propose the following combined operating strategy, Fig. 7, to investigate the possibility of resilient integration. Now, the DT is inserted into the above hysteresis control, Fig. 6. Hence, the HP is definitely started or turned off at the respective condition, while the exact power output is defined by the DT and demand coverage is ensured. This does not improve the plausibility or accuracy of the latter decision per se. Adaptation to environmental conditions initially appears reduced compared to pure DT, as the decision space is restricted during certain periods. On the other hand, allowance for higher frequencies in that framework narrows down the time between decisions, i. e. control steps and therefore adaption to a changing environment. Here, the changing storage content itself is included as a decision feature.

There are two options for the actual creation of the DT regressor. Either the system is pre-trained using a reasonable amount of historic data (e. g. of previous years) or it is updated on a regular basis. The latter can be achieved using an upstream selective node ('**?**'), as seen in Fig. 7. Optimisation during run time but no demand predictions are required in that case. As a consequence the computational advantage vanishes, a sensible choice of parameters for the regressor update is inevitable. In the same manner, temperature limits can be derived from the optimisation, using the *scipy-signal* module (Virtanen et al., 2020) for averaged local minima and maxima of the TES.

3. Results

In this section the proposed control strategies should be applied to the introduced ES model to thoroughly examine their functionality. The following comparisons are based on the introduced KPI, see. Section 2.3.3. In 3.1 and 3.3 annual operation of different simulation runs will be considered, while in 3.2 comparison periods vary.

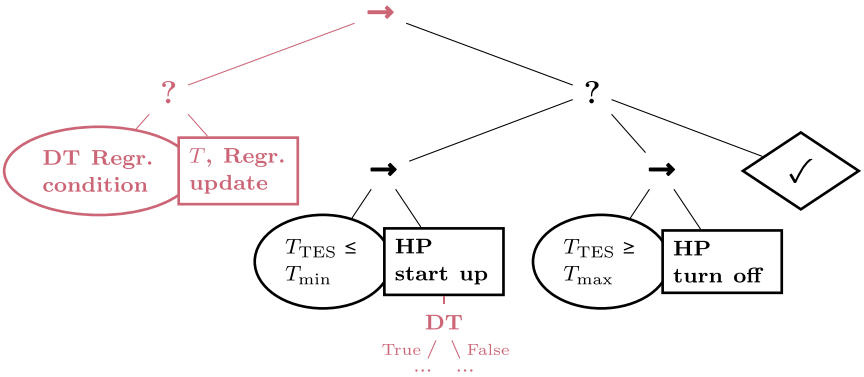

Fig. 7. Complete proposed control structure for the HP using a combination of superordinate BT with decision making DT sub tree. The update of regressor and temperature thresholds is associated with a MILP optimisation using past (known) demand time series.

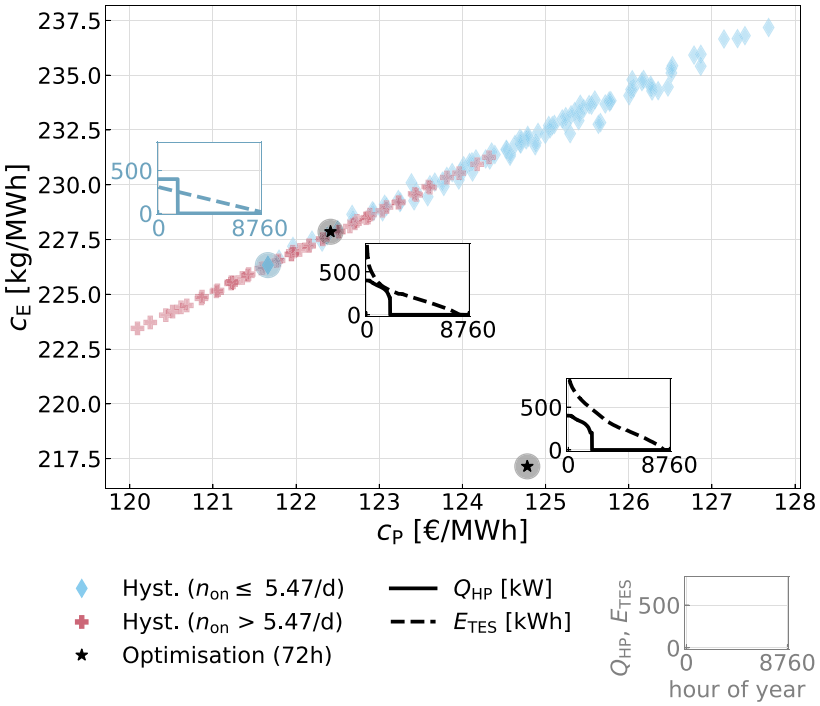

Fig. 8. Results for c_E vs. c_P for annual operation of different hysteresis controllers (♦, +) with 2017 heat demands. Merged operation from optimisation of 24 h intervals (⋆) for comparison. The "feasible" range of starts was defined as $n_{on} \leq 1.1 n_{on,optimisation}$.

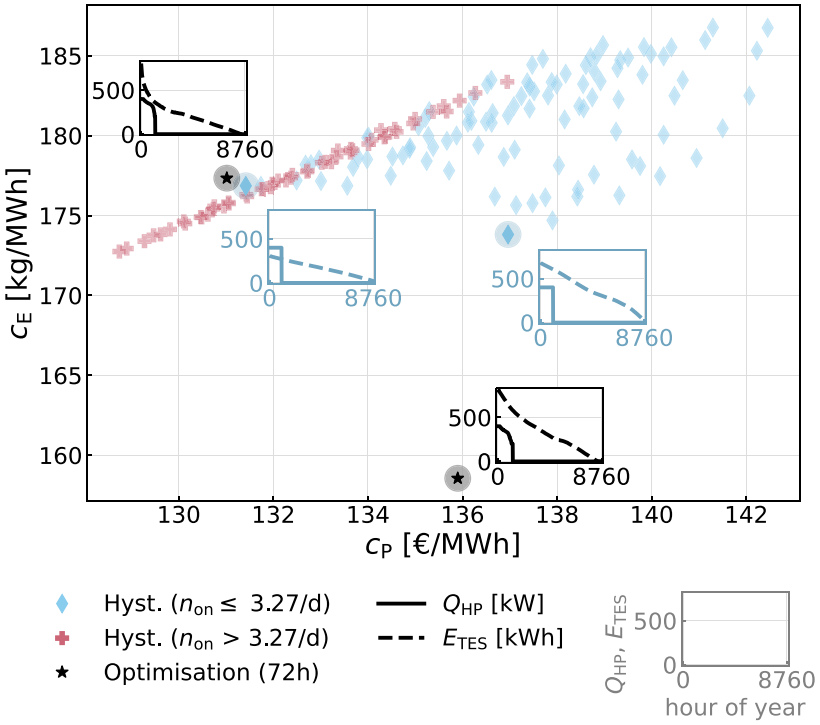

Fig. 9. Results for c_E vs. c_P for annual operation of different hysteresis controllers (♦, +) with 2020 heat demands. Merged operation from optimisation of 72 h intervals (⋆) for comparison. The "feasible" range of starts was defined as $n_{on} \leq 1.1 n_{on,optimisation}$.

3.1. Hysteresis and MILP optimisation

A control system was implemented in correspondence to the explanations in Section 2.5.2, varying temperature limits of the TES between 1 K and 10 K for the lower threshold T_{min} and between 5 K and 35 K (2 K steps) for the upper T_{max} with respect to T_{ref} (storage minimum). Activation frequency Δt_{ex} of the BT was set to 60 s. For the optimisation, however, the choice of an hourly resolution and the application to sliced data proved appropriate with regard to computational limitations, see. Section 2.4.

Figs. 8 and 9 conclude the results of both years, 2017 and 2020. For the optimisation and best hysteresis runs, load duration curves without temporal profile of HP and TES are shown additionally.

Table 4
Results of the MILP and hysteresis optimisation of price and CO_2 emissions. Note that temperature thresholds were inherently fixed for the latter, while average values are given for the first.

KPI	Optimisation			Hysteresis (Δt_{ex} = 60 s)	
	2017 (24 h)	2020 (24 h)	2020 (72 h)	2017	2020
Price [ct kW h^{-1}]	**12.24**	13.25	**13.10**	**12.17** (−0.6%)	**13.14** (+0.3%)
$\overline{T}_{min}/\overline{T}_{max}$ [K]	2.0/12.8	2.1/12.6	1.8/13.8	1/13	1/13
start ups [/d]	4.97	3.22	2.97	4.76	2.96
CO_2 Em. [g kW h^{-1}]	**217.15**	162.07	**158.55**	**226.33** (+4.2%)	**173.81** (+9.6%)
$\overline{T}_{min}/\overline{T}_{max}$ [K]	4.8/18.5	4.7/18.0	4.2/19.3	1/13	1/29
start ups [/d]	4.11	2.86	2.61	4.76	1.32

In 2017 the 1 K/13 K hysteresis presents the best option with respect to price as well as CO_2 emissions, while in 2020 the picture is more diverse. Note, that the selection of feasible parameters was linked to the switching behaviour for comparability, i. e. n_{on}(hyst.) $\leq$ 1.1n_{on}(opti.). The results of optimisation and hysteresis are very close to each other, for 2017 operation of the latter even undercuts the lower resolution optimisation. At first, there are two reasons for this:

- Split of the optimisation: storage content is forced to a predefined value (here 30% or 10.5 K) at the transition. An estimate of the impact of different split lengths can be derived from Table 4 for 2020, i. e. the difference between 24 h and 72 h intervals. However, the assumption of perfect foresight is if at all more realistic in this time range.
- Optimisation is limited to hourly load decisions, while BT hysteresis adapts to its environment by the minute.

The best results are concluded in Table 4. A number of conclusions can be drawn from that.

- In 2017, c_P and c_E show an exclusively linear dependency. The higher overall demand is connected to more operating hours of the HP, where-by an averaging effect is more pronounced. As a consequence storage losses are more decisive for the result of both price and emissions. A lower TES temperature level (limited by requirements for n_{on}) is aimed for, which can also be seen in the duration curves of price optimisation and best hysteresis.
- Decoupling in 2020, in contrast, can be explained by the overall lower demand and less operating hours of the HP. Full exploitation of the storage is favourable to lower emissions, although the effect is comparatively small.
- Average temperature thresholds of the price optimisation resemble that of the best hysteresis. Higher values connected to higher storage losses are accepted in the emission optimisation in order to provide flexibility at time dependent ecological costs $c_{el,E}$.
- This results in a comparatively higher potential for reducing emissions.

Ultimately, the results presented are intended to serve as benchmark values for further analyses. It should be mentioned that the selection of the best hysteresis already represents a form of optimisation, i. e. of the TES temperature thresholds.

3.2. Hourly decision tree control

For the simulation results shown below, DT decisions were introduced into the control system as described in Section 2.5.3. Essentially, the heating output is decided as soon as new data is available, i. e. once an hour, instead of being fixed to nominal as in the hysteresis.

Here it was crucial to determine time periods for training of the respective DT. In each case, the optimisation results of that period (likewise in hourly resolution) served as the basis. Two approaches were investigated:

- Subsequent training, i. e. comparison periods of one month and DT regression using the previous 14, 28 or 56 days.
- Training with 2017 optimisation and comparison within the same date range of 2020.

Maximum depth of the DT regressor was also varied between 20 and 40 (Δ = 5). Fig. 10 shows the results of subsequent training using 2017 demand data.

It can be seen that during summer (June to August) the strategy might be favourable, as optimisation results can be approached. The difference to the hysteresis results is mainly due to the lower heating demand, promoting part load operation of the HP. Sensible decisions on Q_{HP} will have a higher influence compared to winter months. At lower θ_{amb}, the maximum efficiency furthermore approaches operation at nominal load, see Fig. 3, thus making the hysteresis more competitive resp. advantageous. Extending the training period leads to more stable results that are fairly independent of the max. tree depth and even closer to the optimisation results during summer months.

Repeating those simulations for the different demand profiles and higher building efficiency standard of 2020, see Section 2.2, draws a completely different picture, as shown in Fig. 11. Now extension of the regression period is proving to be disadvantageous and c_E of operation during summer months much higher when compared to the hysteresis. This can only partially be explained by the different hysteresis benchmark (1 K/29 K) and possibly higher weather fluctuations. The stronger peaks in the demand profile, see. Fig. A.15 are probably the main reason for the deviating results.

As a consequence also the second strategy of using pre-trained DT with optimisation results of the lower efficiency standard 2017 data shows operation of mixed quality, see Fig. 12. Especially for the shorter periods ($\leq$ 4 months) and for the beginning of both years CO_2 emissions clearly exceed that of the optimal hysteresis. This indicates high influence of the varying weather conditions (see Fig. 1). Furthermore the max. depth of the tree has an unpredictable effect. With a maximum of training data, one full year, hysteresis results are barely approached.

The results show, that the application of pure DT to the presented control problem can have a positive influence on the KPI of emissions in the system. This in principal matches results obtained by Luo et al. (2021) for a similar ES. On the other hand, efficiency standard, demand profile characteristics and season may have a significant impact, making the proposed strategies rather volatile. As an additional problem, switching behaviour of pure DTs at higher execution frequency was already mentioned in Section 2.5.3. The question remains if such temporary advantages can nevertheless be employed in a more versatile control?

3.3. Behaviour tree based strategy of higher frequency

In the above consideration, pure DT proofed to be advantageous during certain periods at least partially when using training resp. comparison data in succession, here from the same year. The hysteresis itself represents a reasonable choice at a higher execution frequency Δt_{ex}. A major difficulty in application could be setting of the temperature thresholds in advance. Here, finally both approaches should be combined, corresponding to Section 2.5.2.

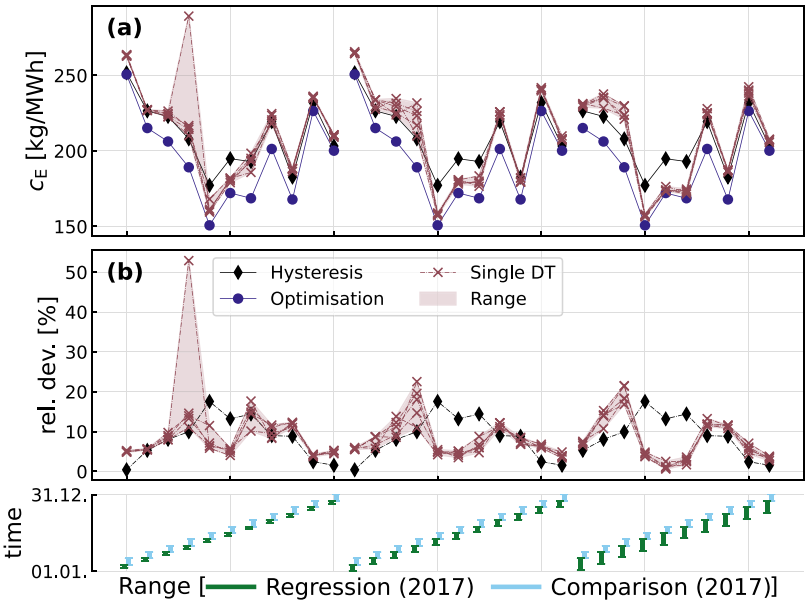

Fig. 10. Strategy of subsequent training (14, 28 or 56 days) and comparison (full month): Hourly DT operation (varied max. depth) is compared to static 1 K/13 K hysteresis and emission optimisation for 2017. Absolute values (a) and relative deviation from optimisation (b).

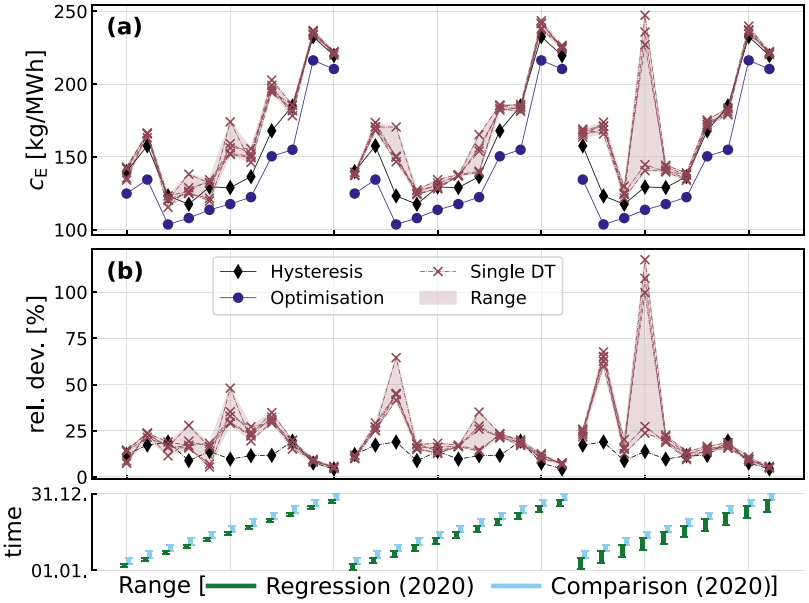

Fig. 11. Strategy of subsequent training (14, 28 or 56 days) and comparison (full month): Hourly DT operation (varied max. depth) is compared to static 1 K/29 K hysteresis and emission optimisation for 2020. Absolute values (a) and relative deviation from optimisation (b).

In principal the BT is updated at constant intervals based on the optimisation with known data (no "predictions" allowed). The (sub-)DT is then generated using the mentioned CART algorithm including the features Q_{dem}, E_{TES}, θ_{amb}, hour of the day and day-ahead-price. Latter play an important role considering the availability of renewable energies and their increased share. By including these variables, which are correlated with the specific emissions (see Section 2.5.3), a higher adaptation to that is expected. For now, only the target value c_E is used for verification; more extensive statistical analyses and introduction of scenarios without conventional energy would certainly be next insightful steps.

Thresholds T_{min} and T_{max} can be derived from that same optimisation by simple averaging of local TES extrema. Thereby the control system dynamically adjusts its operation to available data from the near past. Overall, the system is of course again more computationally intensive due to the inclusion of optimisation. Real-time performance is not affected by this, as it is not updated for every new data step. Different strategies, extracting different parameters, were investigated, as concluded in Table 5.

Table 5
Parametrisation for different strategies using the BT framework with and without DT resp. variable TES thresholds.

Strategy	T_{min}	T_{max}	DT
S1	variable	variable	x
S2	fixed (1 K)	variable	x
S3	fixed (1 K)	fixed (13 K or 29 K)	✓
S4	fixed (1 K)	variable	✓

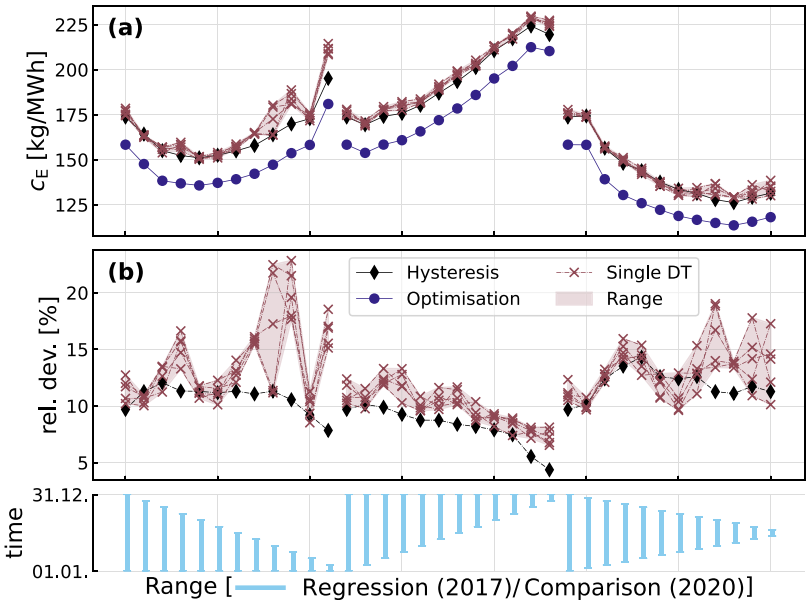

Fig. 12. Strategy of training and comparison in different years: Hourly DT operation (varied max. depth) is compared to static 1 K/29 K hysteresis and emission optimisation for 2020. Absolute values (a) and relative deviation from optimisation (b). The regression period is varied within 2017.

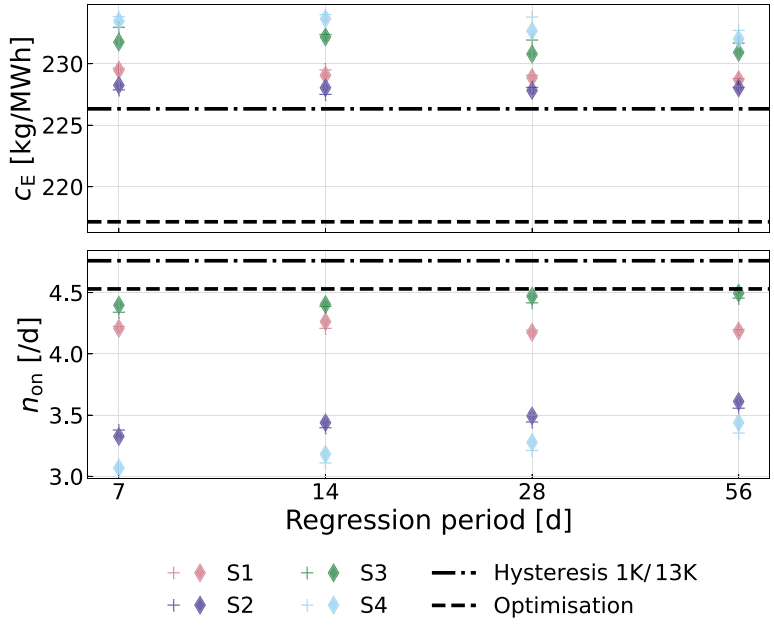

Fig. 13. Results of the different strategies from Table 5 run on 2017 data and based on emission optimisation. Although c_E of the hysteresis is not improved (top), a reduction of n_{on} is distinct (bottom). Δt_{upd} was set to 1 d (+) or 10 d (♦).

Referring to the investigation of hourly activated DT the regression period was varied between one and eight weeks (7, 14, 28, 56 d) and update frequency $1/\Delta t_{upd}$ either every day or every ten days. Δt_{ex} for the full BT was 60 s as in Section 3.1. The results regarding CO_2 emissions as well as start up behaviour based on 2017 data can be seen in Fig. 13.

First, it becomes evident that none of the tested strategies is able to undercut optimal hysteresis let alone optimisation results for c_E. An improvement, with respect to these benchmarks, using longer training periods is present but rather low. Here, also the shorter period Δt_{upd} of 1 d tends to have a small negative effect. Despite that a significant reduction of starts n_{on} was observed. This key strengths is particularly apparent in S2 and S4 as concluded in Table 6. Both possess the additional advantage, that the upper TES storage threshold does not need to be fixed in advance. In that sense also the application of S1 might be a sensible choice, although showing lower reduction in n_{on}. It appears, that the introduction of a DT decision on the actual heat output results in a trade-off between c_E and n_{on}, which can be explained by less efficient but more continuous part load operation.

Results for similar simulations based on the emission optimisation in 2020 are concluded in Fig. B.17. However, due to the overall lower heating demand in 2020 at unchanged ES dimensions n_{on} is lower from the outset, see. Table 4. A further reduction, especially of n_{on} = 1.32 for the 1 K/29 K hysteresis, does not possess reasonable advantages. Instead, the size of the HP could be questioned in this case.

The examination of strategies from Table 5 can also be applied with a focus to price, using the c_P optimisation base. Corresponding operation results for 2020 are shown in Fig. 14. c_P results are again

Table 6
Results of two strategies seen in Fig. 13 compared to that of the corresponding optimisation for emissions and hysteresis. Based on 2017 demand data, 7 d regression period.

	Strategy 2	cf. opti./best hyst.	**Strategy 4**	cf. opti./best hyst.
c_E [g kW h^{-1}]	228.25	+5.1%/+0.8%	233.46	+7.5%/+3.2%
n_{on} [d^{-1}]	3.33	−19.0%/−30.0%	3.07	−25.3%/−35.5%

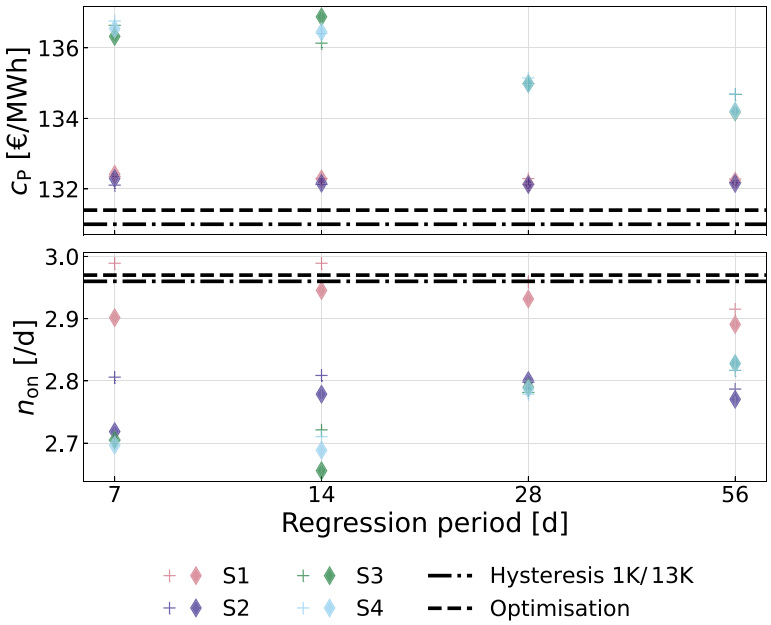

Fig. 14. Results of the different strategies from Table 5 run on 2020 data and based on price optimisation. Heating price (top) and HP start ups (bottom) compared to hysteresis and optimisation results. Δt_{upd} was set to 1 d (+) or 10 d (♦).

Table 7
Results of two strategies seen in Fig. 14 compared to that of the corresponding optimisation for price and hysteresis. Based on 2020 demand data, 14 d regression period.

	Strategy 2	cf. opti./best hyst.	**Strategy 4**	cf. opti./best hyst.
c_P [ct kW h^{-1}]	13.22	+0.9,%/+0.6%	13.65	+4.2%/+3.9%
n_{on} [d^{-1}]	2.78	−6.4%/−6.1%	2.69	−9.4%/−9.1%

very close for S1 and S2, resp. S3 and S4. Some detailed results for the latter in each case and regression period of 14 d are concluded in Table 7.

The part load capable approaches S3 and S4 using the DT subtree are not able to assert themselves significantly here. Nevertheless, the optimal c_P values are approximated again with a reduction of the static parameters and without predictions. This is confirmed by similar results for 2017 demands, shown Fig. B.18.

Actual influences on the service life of the HP can only be roughly estimated in quantitative terms. One assumption would be a linear relation to the number of compressor cycles at an estimated service life of 25 years for an optimum of three cycles per day. This results in an extension of the lifespan for the emissions focused strategies by a maximum of 8.7 a (against hysteresis) or 6.2 a (against optimisation) and for the price by approximately 2.6 a in both cases.

Although this might help to better categorise the above figures, such estimations should be treated with great caution. On the one hand, there are certainly other sources of error or indications of wear and tear, compare (Aguilera et al., 2022). On the other hand, the values used as calculation basis are difficult to verify, especially as the widespread use of heat pumps dates back no more than 25 years.

3.4. Comparison and further benchmarking

Finally a few representative results should be compared as a benchmark. First, we use the same traditional hysteresis as Schmitz et al. (2024) applied to the available data. Here temperature thresholds in the TES are not optimised, a filling level between 20% and 100% is aimed at and control steps are performed every 15 min. For the comparison to pure DT we draw on the best results of Section 3.2. Here a reasonable frequency n_{on} is achieved by using only hourly control steps. As explained in Section 2.2, the integration of renewable energies increased significantly in 2020.

With the same linear estimation from that above, service life of the HP might be included into the economic figures, i. e. extension of lifespan reduces the averaged yearly price. The total demand of the respective year was extrapolated for this purpose. Installation and purchase costs of 250 € kW^{-1} nominal heat output for a large scale HP were assumed (Vimpari, 2021).

A number of conclusions can be drawn from Table 8. From an economic perspective, proposed BT strategies can outperform the traditional hysteresis as well as DT benchmark. The effect is reduced, when including purchase costs in case of the hysteresis. A possibility for additional costs due to more sophisticated controllers cannot be

Table 8
Final benchmarking of the best BT simulation results vs. a traditional rule based strategy (hysteresis) and DT results. For the latter, results are only available for 2020, as the data from the earlier period was used for training.

		BT strategy	Hysteresis	DT (1h)
2017	c_P [ct kW h^{-1}]	12.17	13.23 (+8.7%)	–
	including purchase	13.15	13.67 (+4.0%)	–
	c_E [g kW h^{-1}]	228.25	245.88 (+7.7%)	–
2020	c_P [ct kW h^{-1}]	13.22	14.72 (+11.3%)	13.79 (+4.3%)
	including purchase	14.17	15.17 (+7.1%)	14.84 (+4.7%)
	c_E [g kW h^{-1}]	178.10	186.11 (+4.5%)	175.26 (−1.6%)

excluded here. Interestingly, relative savings when compared to the hysteresis rise for the year with higher renewable integration. While absolute emissions per kW h are lowered accordingly, advantages of using BT control are reduced or disappear completely for a comparison with pure DT control.

4. Conclusion

The development of BT as a conceptual extension of classic FSM is motivated by the need for interpretable and maintainable control structures with the ability for implementation of complex behaviour patterns (Colledanchise and Ögren, 2018). This holds true not only for traditional fields of application such as robotics or gaming, but also for the control requirements of decentralised ESs. Here, a strategy based on a superordinate BT structure was developed to obtain resilient dynamic management of a foundational heating system including thermal energy provision (HP) and storage. Our method combines the optimisation capabilities of *oemof-solph* with the tool of DT, in particular using the CART algorithm. Concrete results for monetary and environmental savings were summarised in Table 8. Weaknesses of these individual approaches besides those figures were highlighted and addressed in the BT design:

- No utilisation of predictions as in purely optimisation based approaches, i. e. MPC and initially lower computation costs,
- Feedback from the environment in contrast to pure DT, reducing sensitivity for different energy requirements and demand profiles,
- Reduction of static control parameters that must be set beforehand.

On the other hand their advantages could be exploited as follows:

- Pattern recognition of optimal operation and its integration,
- Mature and fast training algorithms for associated DT,
- Resilience through interpretable hysteresis-like envelope.

Particularly, the presented control strategy promises to increase lifetime of the HP based on reduction of compressor start ups and associated wear and tear when compared to other CO_2 or price optimised operating strategies (the “traditional” hysteresis used in the benchmark must be excluded here). At the same time it remains competitive to optimised hysteresis and MPC-like optimisation based results with respect to the introduced KPI, price and CO_2 emissions. Thus, choice of the specific objective is decisive for actual implementation. Somewhat counterintuitively, positive compromises were rather achieved for the reduction in emissions in 2017 with lower integration of renewables and a more balanced demand profile, and for the final heating price rather in 2020 with higher integration of renewable energies.

The shown results are merely an example for the expected high potential of using BT in ES for superordinate task switching, with modularity as a decisive advantage. This particularly holds true if BTs are understood as a progression of classic FSMs (attempts to quantitatively measure such advantage and aspects of expandability, interpretability or maintainability has been endeavoured elsewhere but were not the main focus of this work). Their opportunities lie, among other things, in the combination of different approaches as a connecting element between higher-level decision-making and subordinate component control. Possible applications in the energy sector, as otherwise suggested by Perger et al. (2022) or Jingsong (2023), are thereby expanded. Nevertheless structuring and development of meaningful behavioural nodes remains challenging; on the one hand, because a real system may be far more error-prone and fault tolerance can often not be represented in an idealised simulation as conducted here. On the other hand, because also BTs are relatively hierarchical, so that ultimately a prioritisation of different mechanisms must be carried out. In the context of residential buildings, customisation to individual user needs will certainly play an increasingly important role in the future.

A number of improvements might benefit the above shown performance outputs or result in better understanding of it. Especially low resolution and overall amount of demand data represents a relevant limitation. Here an extension of the data set to several years should be considered. An increase in resolution, e. g. to 15 min data, should however be viewed with caution in consideration of the computational effort as a main boundary. Further potential lies in the fine tuning of hyper parameters, here specifically also in the exploitation of extensive knowledge and tools for DT (Pre Pruning, Random Forests etc.).

Next to the need for real-time application and experiments more general further research options include direct learning of BTs using for example genetic programming as proposed by Iovino et al. (2021) and the influence of expansion or enlargement of the ES. How can BT actually be extended in that case and how could the problem of synchronisation be tackled if several components or systems use a similar or the same control implementation?

CRediT authorship contribution statement

Piet Urban: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Peter Klement:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Sunke Schlüters:** Writing – review & editing, Software, Resources, Project administration, Funding acquisition, Conceptualization. **Patrik Schönfeldt:** Writing – review & editing, Software, Resources, Methodology, Data curation, Conceptualization.

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.

Acknowledgements

This document is the results of the research project ENaQ (project number 03SBE111) funded by German Federal Ministry for Economic Affairs and Climate Action (BMWK) and the Federal Ministry of Education and Research (BMBF) and the research project WWNW (project number 03SF0624) funded by German Federal Ministry for Economic Affairs and Climate Action (BMWK). The authors thank Samuel Kees and Philipp Jaschke from *Schulz Systemtechnik* for the regular dialogue and their perspectives and suggestions on the topic.

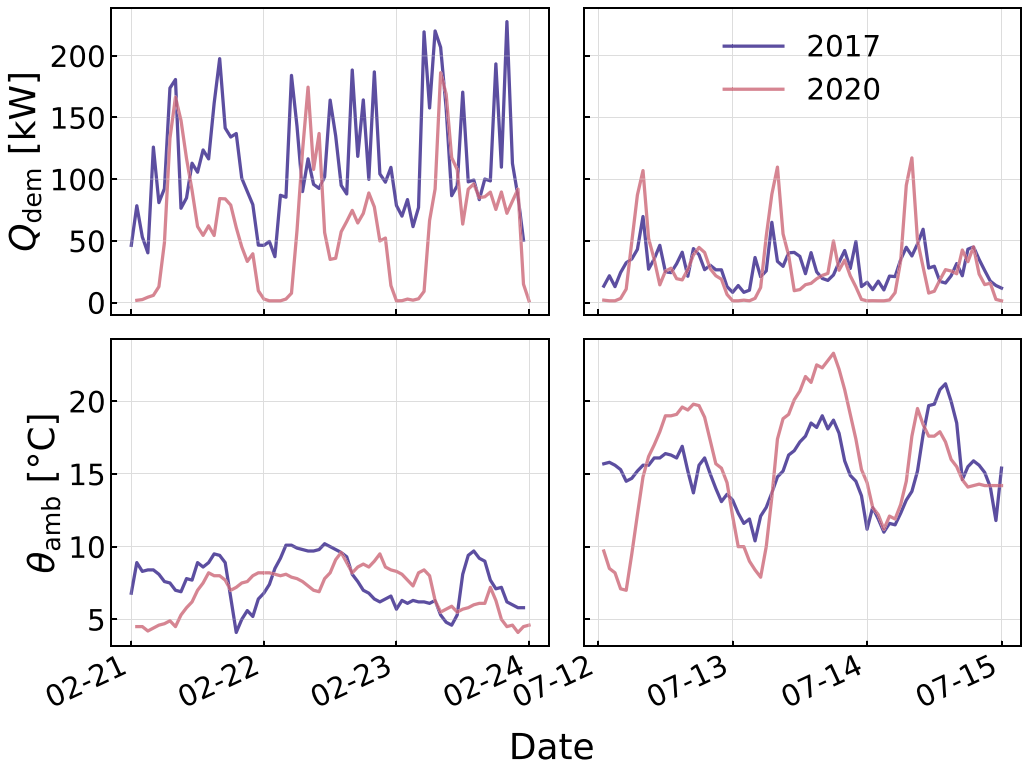

Fig. A.15. Three exemplary days for winter (left) resp. summer (right) of the simulated demand profiles. Different principal load structures despite similar ambient temperatures (plots below) become visible.

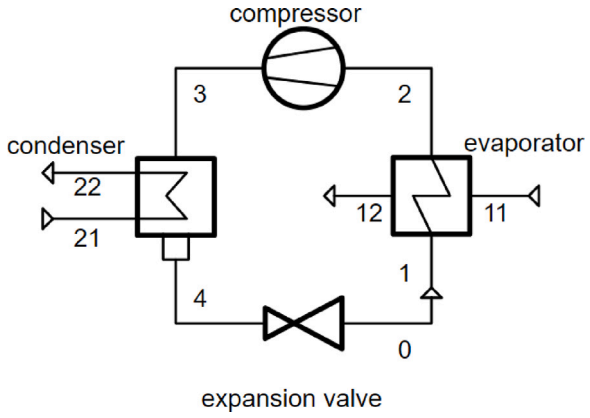

Fig. A.16. Thermodynamic model for the simulated HP. Ambient temperature is set at the evaporator, connection 11, feed and return temperature at the condenser, 22, 21 respectively. Taken from Witte (2024).

Table A.9
Node types that a BT can be build of with conditions for their respective feedback. Control flow nodes process the feedback of their children. Based on Colledanchise and Ögren (2018).

Node type		Succeeds	Fails	Running
Control flow		*feedback from children*		
Sequence	→	All succeed	One fails	One running
Fallback	?	One succeeds	All fail	One running
Parallel	⇉	$\geq M$ succeed	$> N - M$ fail	else
Decorator	◇	Custom	Custom	Custom
Execution				
Action	▭	Upon completion	If impossible	During compl.
Condition	⬭	If true	If false	–

Appendix A. Methodology

See Figs. A.15 and A.16 and Table A.9.

Appendix B. Results

See Figs. B.17 and B.18.

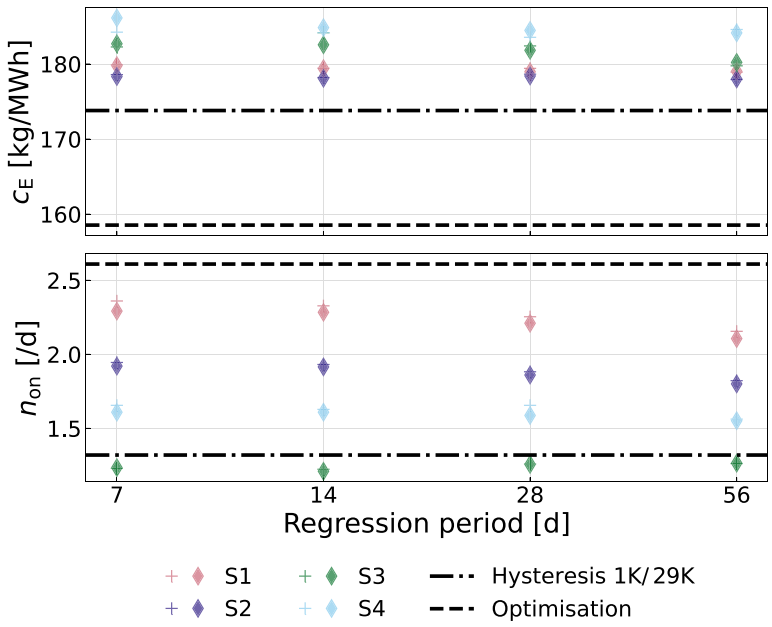

Fig. B.17. Results of the different strategies from Table 5 run on 2020 data and based on emission optimisation. Heating price (top) and HP start ups (bottom) compared to hysteresis and optimisation results. Δt_{upd} was set to 1 d (+) or 10 d (♦).

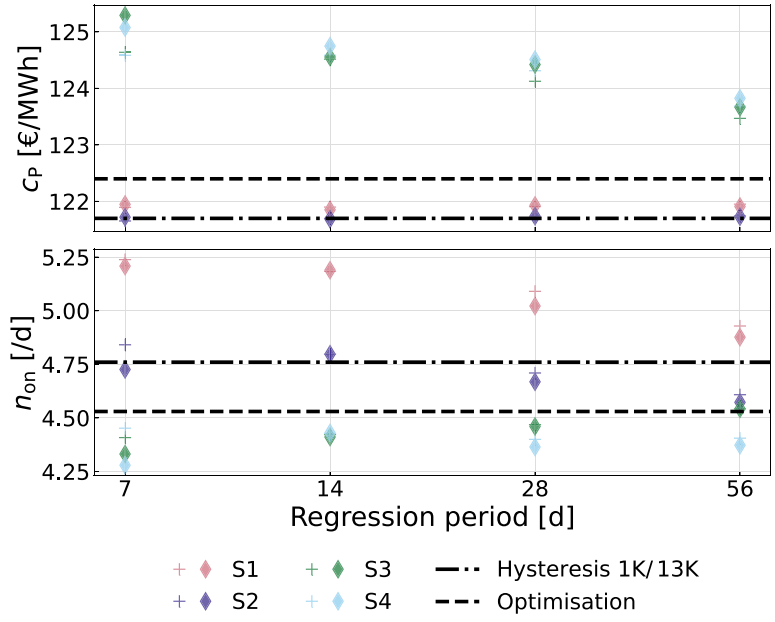

Fig. B.18. Results of the different strategies from Table 5 run on 2017 data and based on price optimisation. Heating price (top) and HP start ups (bottom) compared to hysteresis and optimisation results. Δt_{upd} was set to 1 d (+) or 10 d (♦).

Data availability

Data will be made available on request.

References

Aguilera, J.J., Meesenburg, W., Ommen, T., Markussen, W.B., Poulsen, J.L., Zühlsdorf, B., Elmegaard, B., 2022. A review of common faults in large-scale heat pumps. Renew. Sustain. Energy Rev. 168, http://dx.doi.org/10.1016/j.rser.2022.112826.

Biggar, O., Zamani, M., Shames, I., 2020a. A principled analysis of Behavior Trees and their generalisations. URL http://arxiv.org/pdf/2008.11906v2.

Biggar, O., Zamani, M., Shames, I., 2020b. On modularity in reactive control architectures, with an application to formal verification. URL http://arxiv.org/pdf/2008.12515v3.

Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J., 1984. Classification And Regression Trees. Routledge, http://dx.doi.org/10.1201/9781315139470.

Champandard, A., 2007. Behavior trees for next-gen AI. Game Developers Conference Europe.

Colledanchise, M., Natale, L., 2021. On the implementation of behavior trees in robotics. IEEE Robotics Autom. Lett. 6 (3), 5929–5936. http://dx.doi.org/10.1109/LRA.2021.3087442.

Colledanchise, M., Ögren, P., 2018. Behavior trees in robotics and AI: An introduction. http://dx.doi.org/10.1201/9780429489105.

Cruickshank, C.A., Harrison, S.J., 2010. Heat loss characteristics for a typical solar domestic hot water storage. Energy Build. 42 (10), 1703–1710. http://dx.doi.org/10.1016/j.enbuild.2010.04.013.

Deutscher Wetterdienst, 2023. [dataset] Open Data Platform of *Deutscher Wetterdienst*. availabe at: https://opendata.dwd.de/climate_environment/CDC/observations_germany/climate/hourly/. (Accessed 29 December 2023).

Deutschland, B., 2023. Erneuerbare-Energien-Gesetz vom 21. Juli 2014, Änderung Dezember 2023. https://www.bundesregierung.de/breg-de/schwerpunkte/klimaschutz/novelle-eeg-gesetz-2023-2023972. (Accessed 8 January 2024).

Grimm, A., Schönfeldt, P., Torio, H., Klement, P., Hanke, B., von Maydell, K., Agert, C., 2021. Deduction of optimal control strategies for a sector-coupled district energy system. Energies 14 (21), 7257. http://dx.doi.org/10.3390/en14217257.

Guerin, K.R., Lea, C., Paxton, C., Hager, G.D., 2015. A framework for end-user instruction of a robot assistant for manufacturing. In: 2015 IEEE International Conference on Robotics and Automation. ICRA, IEEE, pp. 6167–6174. http://dx.doi.org/10.1109/icra.2015.7140065.

Hilpert, S., Kaldemeyer, C., Krien, U., Günther, S., Wingenbach, C., Plessmann, G., 2018. The open energy modelling framework (oemof) - A new approach to facilitate open science in energy system modelling. Energy Strat. Rev. 22, 16–25. http://dx.doi.org/10.1016/j.esr.2018.07.001.

Hu, D., Gong, Y., Hannaford, B., Seibel, E.J., 2015. Semi-autonomous simulated brain tumor ablation with ravenII surgical robot using behavior tree. In: 2015 IEEE International Conference on Robotics and Automation. ICRA, IEEE, pp. 3868–3875. http://dx.doi.org/10.1109/icra.2015.7139738.

Iovino, M., Förster, J., Falco, P., Chung, J.J., Siegwart, R., Smith, C., 2022. On the programming effort required to generate Behavior Trees and Finite State Machines for robotic applications. URL http://arxiv.org/pdf/2209.07392v1.

Iovino, M., Styrud, J., Falco, P., Smith, C., 2021. Learning behavior trees with genetic programming in unpredictable environments. In: 2021 IEEE International Conference on Robotics and Automation. ICRA, IEEE, pp. 4591–4597. http://dx.doi.org/10.1109/icra48506.2021.9562088.

Isla, D., 2005. Handling complexity in the Halo 2 AI. Game Developers Conference Europe.

Jingsong, W., 2023. Microgrid real-time decision control method based on behavior trees. In: Proceedings of the 7th PURPLE MOUNTAIN FORUM on Smart Grid Protection and Control. PMF2022, Springer Nature Singapore, pp. 561–574. http://dx.doi.org/10.1007/978-981-99-0063-3_40.

Krien, U., Schönfeldt, P., Launer, J., Hilpert, S., Kaldemeyer, C., Pleßmann, G., 2020. oemof.solph — A model generator for linear and mixed-integer linear optimisation of energy systems. Softw. Impacts 6, 100028. http://dx.doi.org/10.1016/j.simpa.2020.100028.

Luo, X., Xia, J., Liu, Y., 2021. Extraction of dynamic operation strategy for standalone solar-based multi-energy systems: A method based on decision tree algorithm. Sustainable Cities Soc. 70, 102917. http://dx.doi.org/10.1016/j.scs.2021.102917.

Marzinotto, A., Colledanchise, M., Smith, C., Ogren, P., 2014. Towards a unified behavior trees framework for robot control. In: 2014 IEEE International Conference on Robotics and Automation. ICRA, IEEE, pp. 5420–5427. http://dx.doi.org/10.1109/icra.2014.6907656.

Mateas, M., Stern, A., 2002. A behavior language for story-based believable agents. IEEE Intell. Syst. 17 (4), 39–47. http://dx.doi.org/10.1109/MIS.2002.1024751.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. J. Mach. Learn. Res. 12, 2825–2830, URL https://arxiv.org/abs/1201.0490.

Perger, A.v., Gamper, P., Witzmann, R., 2022. Behavior trees for smart grid control. IFAC-PapersOnLine 55 (9), 122. http://dx.doi.org/10.1016/j.ifacol.2022.07.022.

Rabin, S., 2013. Game AI Pro: Collected Wisdom of Game AI Professionals. CRC Press, Hoboken.

Razmi, D., Lu, T., 2022. A literature review of the control challenges of distributed energy resources based on microgrids (MGs): past, present and future. Energies 15 (13), 4676. http://dx.doi.org/10.3390/en15134676.

Schmeling, L., Schönfeldt, P., Klement, P., Vorspel, L., Hanke, B., von Maydell, K., Agert, C., 2022. A generalised optimal design methodology for distributed energy systems. Renew. Energy 200, 1223–1239. http://dx.doi.org/10.1016/j.renene.2022.10.029.

Schmitz, S., Brucke, K., Kasturi, P., Ansari, E., Klement, P., 2024. Forecast-based and data-driven reinforcement learning for residential heat pump operation. Appl. Energy 371, 123688. http://dx.doi.org/10.1016/j.apenergy.2024.123688.

Schönfeldt, P., Grimm, A., Neupane, B., Torio, H., Duranp, P., Klement, P., Hanke, B., von Maydell, K., Agert, C., 2022. Simultaneous optimization of temperature and energy in linear energy system models. In: 2022 Open Source Modelling and Simulation of Energy Systems. OSMSES, IEEE, pp. 1–6. http://dx.doi.org/10.1109/OSMSES54027.2022.9768967.

Somvanshi, M., Chavan, P., Tambade, S., Shinde, S.V., 2016. A review of machine learning techniques using decision tree and support vector machine. In: 2016 International Conference on Computing Communication Control and Automation. ICCUBEA, IEEE, pp. 1–7. http://dx.doi.org/10.1109/ICCUBEA.2016.7860040.

Stonier, D., 2023. py_trees documentation. URL https://py-trees.readthedocs.io/en/devel/. (Accessed 14 August 2024).

Umweltbundesamt, 2022. Erneuerbare Energien in Deutschland. Daten zur Entwicklung im Jahr 2022. https://www.umweltbundesamt.de/sites/default/files/medien/1410/publikationen/2023-03-16_uba_hg_erneuerbareenergien_dt_bf.pdf. (Accessed 8 January 2024).

Vimpari, J., 2021. Should energy efficiency subsidies be tied into housing prices? Environ. Res. Lett. 16 (6), 064027. http://dx.doi.org/10.1088/1748-9326/abfeee.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., SciPy 1.0 Contributors, 2020. SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods 17, 261–272. http://dx.doi.org/10.1038/s41592-019-0686-2.

Wirth, H., 2023. Recent Facts about Photovoltaics in Germany. Fraunhofer ISE, https://www.ise.fraunhofer.de/en/publications/studies/recent-facts-about-pv-in-germany.html. (Accessed 28 December 2023).

Witte, F., 2024. TESPy documentation. URL https://tespy.readthedocs.io/en/main/. (Accessed 8 August 2024).

Witte, F., Tuschy, I., 2020. TESPy: Thermal engineering systems in Python. J. Open Source Softw. 5 (49), 2178. http://dx.doi.org/10.21105/joss.02178.