FPS Economy, S.M.E.s, Self-employed and Energy

ECOFLEX

With the support of the Energy Transition Fund

D5.2 Multi-energy community energy management

Version number and Date: Version 1; 15/10/2025
Author (s): Ruddick Julian (VUB)

Abstract for dissemination (PU)

This deliverable presents two studies on how advanced energy management systems (EMS) can use the flexibility of dwellings to benefit both the household and the electricity grid. The first study tests and compares several EMS approaches in a real home electric grid equipped with solar panels, a battery, and an emulated electric vehicle charger. The goal is to see which EMS best manages household flexibility to lower electricity bills under a dynamic pricing contract with hourly electricity prices and the Flemish capacity tariff. The second study uses computer simulations of 500 typical Belgian dwellings that have electrified their households with heat pumps, electric boilers and electric vehicles. It evaluates the financial and comfort benefits of installing advanced EMS in these electrified dwellings under dynamic electricity prices.

Contents

List of abbreviations

EMS	Energy management system
EV	Electric vehicle
MPC	Model predictive control
PV	Photovoltaic
SOC	State of charge
RBC	Rule-based control
RL	Reinforcement learning

1.1 Overview

This deliverable is separated into two parts. The first part is the experiment performed at the Smart Village Lab where four different EMS algorithms were compared on the real hardware of a house with photovoltaic panels, a battery and an electric vehicle. The electricity costs were defined by a dynamic contract with electricity prices changing every hour and subject to the Flemish capacity tariff. The goal of this experiment was to evaluate how well each EMS is capable of using the flexibility of the battery and electric vehicle to minimize the electricity costs of the household.

The second part of this deliverable analyses the benefits of using a state-of-the-art EMS for electrified dwellings in Belgium. This analysis is performed by simulating 500 dwellings representative of the Belgian building stock, with gas heating replaced by heat pumps, electric water boilers, and home charging of electric vehicles. This study takes into account the financial benefits as well as thermal comfort and charge completion of the electric vehicle.

The first part has been published as a peer-reviewed journal article in the Knowledge Based Systems journal [44] and the second part is currently being submitted to the peer-reviewed conference BuildSys 2026.

The following two chapters of the deliverable contain a slightly adapted version of the full text of both papers.

1.2 Publications and Abstracts

Here are the full titles and abstracts of both papers:

Real-world validation of safe reinforcement learning, model predictive control and decision tree-based home energy management systems

Abstract:

Recent advancements in machine learning based energy management approaches,

specifically reinforcement learning with a safety layer (OptLayerPolicy) and a metaheuristic algorithm generating a decision tree control policy (TreeC), have shown promise. However, their effectiveness has only been demonstrated in computer simulations. This paper presents the real-world validation of these methods, comparing them against model predictive control and simple rule-based control benchmarks. The experiments were conducted on the electrical installation of four reproductions of residential houses, each with its own battery, photovoltaic, and dynamic load system emulating a non-controllable electrical load and a controllable electric vehicle charger. The results show that the simple rules, TreeC, and model predictive control-based methods achieved similar costs, with a difference of only 0.6%. The reinforcement learning based method, still in its training phase, obtained a cost 25.5% higher than the other methods. Additional simulations show that the costs can be further reduced by using a more representative training dataset for TreeC and addressing errors in the model predictive control implementation caused by its reliance on accurate data from various sources. The OptLayerPolicy safety layer allows safe online training of a reinforcement learning agent in the real world, given an accurate constraint function formulation. The proposed safety layer method remains error-prone; nonetheless, it has been found beneficial for all investigated methods. The TreeC method, which does not require building a realistic simulation for training, exhibits the safest operational performance, exceeding the grid limit by only 27.1 Wh compared to 593.9 Wh for reinforcement learning.

Should Belgians install an advanced energy management system in their electrified dwellings?

Abstract:

As dwellings electrify and more complex electricity pricing spreads in Belgium, the question arises: What are the potential benefits of installing advanced energy management systems in an electrified Belgian dwelling?

To answer this question, we simulate 500 representative Belgian dwellings in which fossil systems are replaced by controllable electric technologies: space heating via air-to-air heat pumps, domestic hot water via electric boilers, and private transport via electric vehicles with home chargers. Photovoltaic systems are assigned according to current adoption rates.

We evaluate two advanced energy management systems (EMS): model predictive control (MPC) and TreeC, a machine learning method using decision trees to model the EMS strategy. These are compared to a rule-based controller (RBC) that mimics default operation and an MPC with perfect forecasts (MPC P) as an upper bound. Performance for these EMSs is assessed for 2023 under a real contract with hourly prices and a monthly peak tariff.

Average annual savings versus RBC are 82.5€, 116.6€, and 177.4€ for TreeC, MPC, and MPC P, respectively. All EMSs deliver similar thermal comfort (26-31 Kh). MPC saves more, while TreeC achieves higher EV charge completion (93.3% vs 88.8% for MPC, with a max charge of 95%) and a smoother aggregate load across the 500 dwellings. Savings remain modest relative to EMS purchase (280-680€) and subscription (0-360€/year) costs, which is not favourable towards widespread adoption.

2.1 Introduction

This chapter presents a comparison of four EMSs in an experimental setup with real energy assets. The four implemented EMSs are TreeC, reinforcement learning (RL), model predictive control (MPC) and rule-based control (RBC). Each EMS manages a home battery and emulated EV charger and aim to minimize household electricity costs under a dynamic tariff with prices changing every hour. Implementing EMSs in a real-world setting is important to identify potential issues that do not arise in simulations. This chapter is an adaptation of the published journal article "*Real-world validation of safe reinforcement learning, model predictive control and decision tree-based home energy management systems*" [44].

The main contribution of this chapter are:

- Implementation of TreeC in a real-world experimental setup.
- Comparison of TreeC, RL, MPC and RBC on a real experimental case with batteries, solar PV, electric loads, emulated electric vehicle charging and hourly electricity based pricing.

2.2 Method

2.2.1 Hardware setup

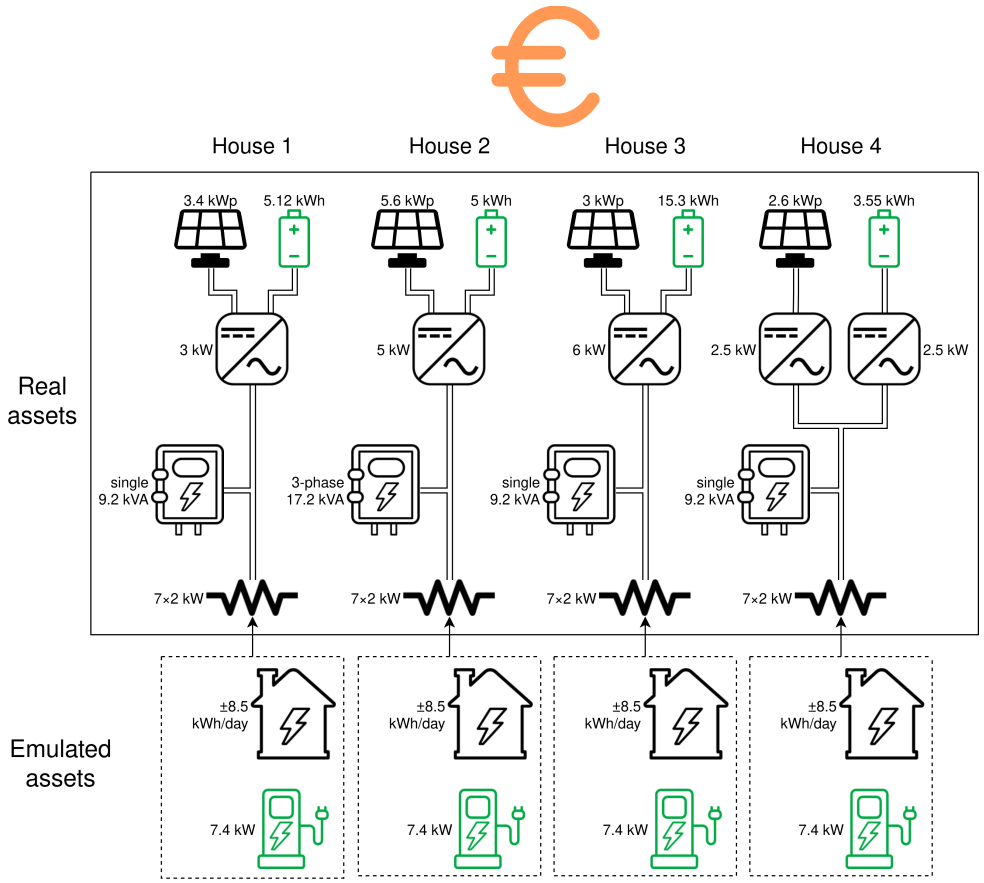

Figure 2.1: Schematic of this chapter's experimental setup for each house. From top to bottom: PV installation (with peak power), Battery (with capacity), inverter (with maximum power), grid connection (with phase specification and maximum power), dynamic loads (with maximum power), emulated house consumption (with average daily consumption), and EV charger (with maximum charge power). Assets highlighted in green are controlled by the EMS, optimization is performed on the electricity cost, indicated in orange.

The experimental setup consists of four real reproductions of a house's electrical circuit. Each house has a battery, a PV installation, an emulated electric load consumption and emulated EV charging sessions. The houses are connected to the grid and have official digital meters from the local Belgian distribution system operator that measures energy imported from and exported to the grid. The home batteries and PV installations are different for each house, whereas the emulated electric load consumption and charging sessions are identical for all houses. This setup replicates an electrical configuration typical of a modern home with any non-electrical heating system. The experiment was conducted in the Smart Village lab on the Green Energy Park site in Zellik, Belgium.

Figure 2.2: Photograph of the experimental setup: with `G` being the grid connection equipped with the official digital meter of the local distribution system operator, `B` the battery, `S` the solar PV installation inverters (solar panels equipped on the roof) and `L` the connectors to the dynamic loads (dimmers in separate cabinet and electrical loads stationed outdoors). The number represents the house numbers.

The electrical loads used for the consumption and charging sessions come from a dynamic load system composed of seven 2000W Perel®PHP2000 heaters controlled through Showtec Single DP-1 dimmers. The electrical loads are calibrated so that the setpoints match the real measured active power consumption more closely (we expect this difference to come from measurement[1] inaccuracies and manufacturing tolerances of the loads).

Measurement devices for power, voltage and current are present for the PVs, home batteries and grid connections. The power of the electrical loads and other power losses in the setup is calculated via the PV, battery and grid measurements.

A schematic and a picture of the setup are displayed in figs. 2.1 and 2.2.

The hardware specifications of the home batteries, inverters and PV installations are summarized in tables 2.1 to 2.3.

Table 2.1: Battery specifications

House	Battery model	Capacity	Max charge	Max discharge
1	BYD HVS 5.1	5.12 kWh	3.2 kW	3.2 kW
2	Huawei LUNA2000-5-EO	5 kWh	2.5 kW	2.5 kW
3	D-Centralized LF2	15.3 kWh	3.0 kW	4.0 kW
4	Pylontech US2000B	3.55 kWh	1.7 kW	2.5 kW

[1] No voltage and therefore power factor, measurement *per* circuit was operational

Table 2.2: Inverter specifications

House	Inverter model	Assets	Max AC power
1	Fronius Primo GEN24 3.0	PV + battery	3.0 kW
2	Huawei SUN2000-5KTL-M1	PV + battery	5.0 kW
3	Sermatec SMT-5K-TL-LV	PV + battery	6.0 kW
4	SMA Sunny Boy SB2.5-1VL-40	PV	2.5 kW
4	Victron Multiplus-II GX 48/3000/35-32	battery	2.5 kW

Table 2.3: Additional specifications

House	PV peak power	AC circuit	House breaker
1	3.4 kWp	mono phase	9.2 kVA
2	5.6 kWp	three phase	17.2 kVA
3	3 kWp	mono phase	9.2 kVA
4	2.6 kWp	mono phase	9.2 kVA

2.2.2 Experimental setup

2.2.2.1 Consumption profile

We utilize the real-time (10 seconds) residential electricity "household" consumption data from Schlemminger et al. [46], focusing specifically on the 2019 dataset, which predates the COVID-19 pandemic. Our experiments employ data from single-family house (SFH) 19, which does not include solar PV, battery, or EV charging. Additionally, this household dataset excludes heat pump loads. SFH 19 is selected due to its high data availability and quality, along with its typical consumption metrics, averaging 3425 kWh per year.

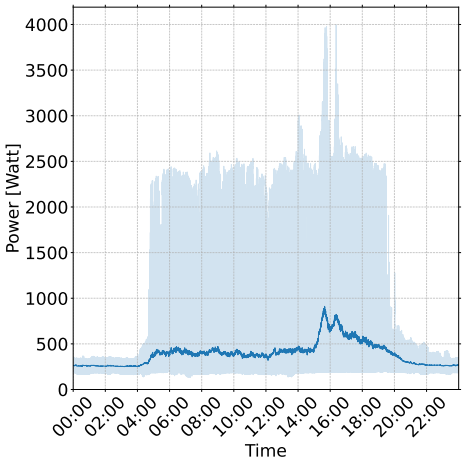

(a) Time series with 2.5th and 97.5th percentile

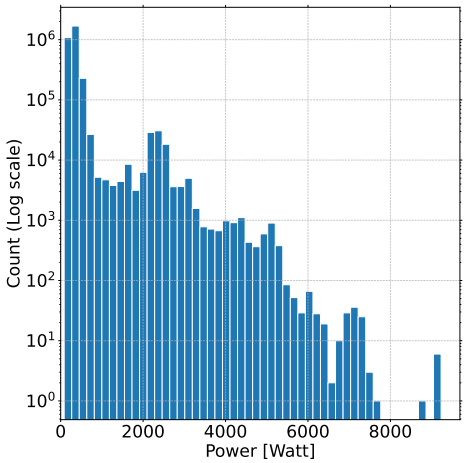

(b) Histogram

Figure 2.3: Household consumption profile: 10 seconds SFH 19 data [46] without PV, battery, EV charging or heat pump having 3425.75 kWh/year, Std: 445.85 W, max: 9229 W.

We shifted the dataset so that weekdays align (that a Monday in runtime is also a Monday in the dataset) and used this for all our hardware-in-the-loop simulated dynamic loads so that each house has the same real-time non-controllable load.

2.2.2.2 Electric vehicle and driver charging behaviour

The residential electric mobility demand consists of two parts: a) a first part specific to the EV user charging behaviour (i.e. arrival time, departure time and energy need), and b) a second part specific to the EV charging power profile.

Firstly, the EV user charging behaviour is simulated directly using an existing open-source dataset available in [48]. This dataset comprises data from 97 EV users, encompassing 6878 charging sessions at residential locations spanning from December 2018 to January 2020. Among these users, "Bl2-2" has been selected for the experiment due to its alignment with the average EV user in terms of arrival time, parking duration, and energy requirements. However, it distinguishes itself by exhibiting a notably high frequency of charging sessions, rendering it particularly interesting for the experiment as it allows for more interactions with the EMS developed in this chapter. From a statistical standpoint, this EV user has undergone 315 charging sessions, with an aggregate energy consumption of 3274.74 kWh across 386 days.

Secondly, the EV charging power profile is built based on certain variables and constraints. The EV charging power is constrained by (2.1).

$$0 \leq P_t^{ev} \leq P_{max}^{ev} \tag{2.1}$$

where:

P_t^{ev} [kW]	EV charging power at time step t
P_{max}^{ev} [kW]	Maximum charging power

P_{max}^{ev} is the maximum charging power in $[kW]$ set to 7.4 kW, following the IEC 61851 charging standards [29]. Additionally, the maximum charging power follows the Constant-Current/Constant-Voltage (CC-CV) uncoordinated charging method defined in Vagropoulos and Bakirtzis [54]. Such a method stipulates that the power is a function of the state of charge (SOC) as a piece-wise linear function summarized in (2.2).

$$P_t^{ev} = \begin{cases} P_{max}^{ev} & \text{if} \quad SOC_t^{ev} \leq SOC_{cc,cv}^{ev} \\ P_{max}^{ev} - (P_{max}^{ev} - P_{min}^{ev}) * \dfrac{SOC_t^{ev} - SOC_{cc,cv}^{ev}}{1 - SOC_{cc,cv}^{ev}} & \text{if} \quad SOC_t^{ev} > SOC_{cc,cv}^{ev} \end{cases} \tag{2.2}$$

where:

SOC_t^{ev} [%]	SOC of the EV at time step t
$SOC_{cc,cv}^{ev}$ [%]	Transition SOC between the Constant-Current and Constant-Voltage phases, set to 90%
P_{min}^{ev} [kW]	Maximum charging power at 100% SOC, set to 1 kW

Finally, both charging power and SOC variables are linked using the equality constraint (2.3).

$$SOC_{t+1}^{ev} = SOC_t^{ev} + \eta^{charge} \times \frac{P_t^{ev} \times \Delta t}{E^{ev}} \tag{2.3}$$

where:

SOC_t^{ev} [%]	SOC at time step t
E^{ev} [kWh]	capacity of the EV battery, set to 60 kWh
η^{charge} [%]	charging efficiency, set to 95%

The charging efficiency is set to 95% as it is typical for lithium-ion batteries [20].

2.2.2.3 Tariff

The real-time pricing used in this study is based on a one-year dynamic contract of a major Belgian energy supplier, as if it were taken in April 2024[2]. The total cost is composed of four costs with different pricing mechanisms: the day-ahead cost, the offtake extras cost, the peak cost and the fixed yearly cost.

The day-ahead cost depends on the prices of eqs. (2.4) and (2.5).

$$Vo_t = \begin{cases} (Vd_t + 0.011\frac{€}{kWh}) * 1.06 & \text{if } (Vd_t + 0.011\frac{€}{kWh}) \geq 0 \\ (Vd_t + 0.011\frac{€}{kWh}), & \text{otherwise} \end{cases} \tag{2.4}$$

$$Vi_t = Vd_t - 0.009\frac{€}{kWh} \tag{2.5}$$

where:

Vo_t [€/kWh]	offtake price at time step t
Vi_t [€/kWh]	injection price at time step t
Vd_t [€/kWh]	Belgian day-ahead hourly price at time step t

A value-added tax of 6% is added to the offtake price when it is positive.

[2]Full contract available on the paper's GitHub repository at https://github.com/EVERGi/real_validation_saferl_treec_paper/

Extra costs that occur for energy taken from the grid for transport, distribution and taxes are calculated together and labelled as offtake extras cost. The offtake price for all these additional offtake costs Vf is 0.114 €/kWh.

The monthly peak offtake cost Cp_m is calculated by eq. (2.6).

$$Cp_m = \frac{|B_m|}{N_m} Vp * \max_{t \in B_m}(\frac{Eo_t}{0.25h}, 2.5kW) \tag{2.6}$$

where:

B_m	set of all billed time steps in month m
N_m	total number of time steps in the month
Vp [€/kW]	peak price of 3.5€/kW
Eo_t [kWh]	energy taken out of the grid at time step t

The Cp_m cost adds a penalty to the highest offtake power averaged over a 15-minute time step that occurred for each month. In the context of this study, the monthly peak offtake cost is proportional to the billed period of each month to also include this cost for unfinished months. The offtake energy Eo_t is divided by the length of a time step to obtain the power in kW. A minimum monthly peak power of 2.5 kW is enforced in the cost, even if the measured peak power is lower.

The fixed yearly cost is calculated using the yearly price Vy of 115.84€ and is proportional to the billing period. The total electricity cost is calculated by eq. (2.7).

$$C_{tot} = \underbrace{\sum_{t \in T}(Eo_t Vo_t - Ei_t Vi_t)}_{\text{day-ahead cost}} + \underbrace{\sum_{t \in T} Eo_t Vf}_{\text{offtake extras cost}} + \underbrace{\sum_{m \in M} Cp_m}_{\text{peak cost}} + \underbrace{\frac{|T|}{N_y} Vy}_{\text{yearly cost}} \tag{2.7}$$

where:

T	Set of all billed time steps
M	Set of billed months
N_y	Total number of time steps in the year
Ei_t [kWh]	Energy injected to the grid at time step t

2.2.2.4 House switching

As described in subsection 2.2.1, the four houses have different home batteries and PV installations. Assigning a different EMS to each house and comparing the performance by running them simultaneously would not yield comparable results. One house could obtain better performances because it has a better battery or

PV installation. To avoid this bias, the EMSs are switched every 24 hours to a different house.

Since there are 4 houses and 4 EMSs to test, there are a total of 24 possible house and EMS combinations. Each combination is tested twice, resulting in 48 days of testing in total. The 48-day schedule is obtained by randomly shuffling the 24 combinations until two distinct 24-day schedules are generated. This ensures no consecutive days have the same RL or MPC EMS assigned to the same house. This rule is applied for practical reasons, as the RL and MPC EMSs behaviours depend on how they operated a house previously. Failed test days can occur, and not having an MPC or RL EMS on the same house for consecutive days provides more time to identify a failure while preventing its propagation to the following day.

The switching is conducted daily at 15:00. By this time, the batteries of all houses are charged to 100% to ensure consistent starting conditions post-switch. The EV schedule is adjusted to prevent charging sessions from overlapping with the 15:00 switching time. This is achieved by shortening the duration of overlapping sessions, either by starting or finishing them at 15:00, with the option of the longest duration being selected. The start and end SOC of the EV batteries are maintained according to the original schedule.

These switching rules have been selected because 15:00 is a time when there are fewer charging sessions compared to the morning and evening hours. At 15:00, the home battery also had the opportunity to charge using the excess energy generated by the PV installation in the preceding hours. Additionally, at this time, the day-ahead prices for the next day are available, eliminating the need to forecast them for the MPC EMS.

2.2.2.5 Enforced charging behaviour

An enforced charging behaviour is enforced in order to reach 100% SOC at the switching time for the home batteries and to reach the imposed final SOC of an EV battery at the end of a charging session. The batteries are charged at maximum power to reach their SOC goal at the latest possible time minus a buffer time $Btime$. This buffer time is used to avoid not reaching the SOC goal in time due to non-perfect battery models, safety layer activation (which could reduce the battery setpoints, see subsection 2.2.3) or other unexpected event.

The buffer time value is higher when the difference between the SOC goal and the current SOC of the battery is higher because more events could prevent the battery from reaching its SOC goal in time. This relation is described in eq. (2.8).

$$B_{time} = (SOC_g - SOC) * (B_{max} - B_{min}) + B_{min} \tag{2.8}$$

where:

B_{time} [min]	time buffer in minutes
SOC_g [%]	SOC goal of the battery
SOC [%]	current SOC of the battery
B_{max} [min]	time buffer if 100% of the battery still needs charging
B_{min} [min]	minimum time buffer

Every 5 seconds, the battery controller calculates with eqs. (2.3) and (2.9) whether the battery can reach the SOC goal at the intermediate time, defined as the target time minus the time buffer. If not, the battery is charged at maximum power for 1 minute. For both the home batteries and the EV batteries, B_{min} is set to 3 minutes. For the home batteries B_{max} is set to 1 hour and for the EV batteries B_{max} is set to 4 hours. The B_{min} and B_{max} values were obtained by testing in simulation what values would consistently meet the SOC goal in time throughout the 3 months of training data (see subsection 2.2.5.1).

2.2.3 Safety layer

A safety layer is implemented for all EMSs to ensure the maximum power limits of the main circuit breaker is not exceeded. The safety layer is run every 5 seconds and can override the setpoints of the EMSs if the power limits would be exceeded. The method used in this study is the `OptLayerPolicy` safety layer from Ceusters et al. [8] - as it is considered one of the state-of-the-art methods.

2.2.4 Simulation

A simulation of the homes is implemented in order to train the TreeC EMS and compare the real experiment with the simulation. A simulation time step of 15 minutes was used. The simulation models the behaviour of the home battery, EV charging, PV, electrical load and grid connection and how they react to different control strategies from EMSs.

The simulation model of the home battery is described in eq. (2.9).

$$SOC_{t+1}^{batt} = SOC_t^{batt} + \frac{P_t^{ch} * \eta^{batt} - P_t^{di} / \eta^{batt}}{E^{batt}} \Delta t \tag{2.9}$$

where

where :

SOC_t^{batt} [%]	SOC of the home battery at time step t
P_t^{ch} [kW]	Charge power at time step t
P_t^{di} [kW]	Discharge power at time step t
E^{batt} [kWh]	Energy capacity of the home battery
η^{batt} [%]	Charge and discharge efficiency of the home battery
Δt [h]	Time step duration in hours

The efficiency of the home battery η^{batt} is set to 95% [20] for home batteries 1,2 and 4 and 96% for battery 3 as stated in the documentation of the battery. The energy capacity of each home battery is stated in table 2.1.

The EV is modelled using eqs. (2.1) to (2.3). The PV and electric loads come from the measurements and are included in the simulation. The grid power is calculated with eq. (2.10).

$$P_t^{grid} = -P_t^{load} - P_t^{ev} + P_t^{pv} + P_t^{batt} \forall t \tag{2.10}$$

To simplify the modelling of the simulations, the limit imposed by the hybrid inverter on the maximum combined power from the PV installation and the home battery is not enforced (see table 2.2). This simplification is not expected to significantly impact the results of the simulations because the hybrid inverter limits are very rarely reached in practice for the setups used in this study.

The enforced charging behaviour described in subsection 2.2.2.5 is closely approximated in the simulation. The time buffer is calculated using eq. (2.8) at each time step. If following the battery charge or discharge power given by the EMS does not allow reaching the SOC goal at the intermediate time, the battery is charged at the minimum power necessary to meet that goal.

2.2.5 Energy management systems

This section presents the four EMS methods used and the data used to train the MPC, RL and TreeC EMSs. All EMSs send a setpoint to the home battery and the EV charger every 15 minutes; how they calculate this setpoint is described in this section.

2.2.5.1 Training data

The MPC, RL, and TreeC EMSs required training data and used three months of data spanning from 2024/01/01 at 00:00 to 2024/04/01 at 00:00. The dataset included real PV power production, household power consumption described in subsection 2.2.2.1 and EV charging sessions from subsection 2.2.2.2. The MPC utilized this data to train its forecasting models, while the RL models were pre-trained on the training data with additional power profiles for the home batteries

and EV chargers. The power profiles of the home batteries and EV chargers were generated as if an RBC EMS was controlling these assets. The TreeC EMS used this data in the simulator, described in subsection 2.2.4, to construct its decision tree.

2.2.5.2 Rule-based control

The RBC EMS implemented in this setup consists of default rules commonly found in many devices. The primary rule for the home battery focuses on self-consumption "behind-the-meter". It charges when excess power is available from the PV installation and discharges to meet load demands. The rule for the EV charger is to always charge the EV immediately at maximum allowed power. This EMS serves as a benchmark to evaluate the performance of the other EMSs.

2.2.5.3 Model predictive control

The MPC EMS in this setup is implemented as a Mixed-Integer Linear Program, which is solved at each (15-minute) time step. The horizon of the MPC is variable and ends at the next switching time. The MPC model is solved every 15 minutes. The variables of the MPC are the power profiles of the home battery and the EV charger, which are the assets it needs to control.

The objective function of the MPC is described in eq. (2.11).

$$O = \underbrace{\sum_{t \in H} (Eo_t Vo_t - Ei_t Vi_t)}_{\text{day-ahead cost}} + \underbrace{\sum_{t \in H} Eo_t Vf}_{\text{offtake extras cost}} \underbrace{+ Vp * \max_{t \in H}(\frac{Eo_t}{0.25h}, Pp)}_{\text{peak cost}} \tag{2.11}$$

where:

H	Set of time steps in the MPC horizon
Pp [kW]	Previous grid peak power for that house when running MPC

At the beginning of the experiment, Pp is set to 2.5 kW. This objective function penalizes heavily exceeding the previous grid peak power Pp.

The MPC system models take as inputs the time of the day, the SOC of the home battery and EV battery, the Belgian day-ahead hourly prices Vd_t until the next switching time, the forecasted power profiles until the next switching time of the PV installation and the electric load and the forecasted final SOC and departure time of the connected EV.

The MPC constraints for the home battery, grid power, and EV are consistent with the simulation models described in eqs. (2.1), (2.3), (2.9) and (2.10). However, the EV charging behaviour in the MPC formulation differs from the simulator model in eq. (2.2), as detailed in eq. (2.12).

$$P_t^{ev} \leq \begin{cases} P_{max}^{ev} & \text{if} \quad SOC_{t+1}^{ev} \leq SOC_{cc,cv}^{ev} \\ P_{max}^{ev} - (P_{max}^{ev} - P_{min}^{ev}) * \dfrac{SOC_{t+1}^{ev} - SOC_{cc,cv}^{ev}}{1 - SOC_{cc,cv}^{ev}} & \text{if} \quad SOC_{t+1}^{ev} > SOC_{cc,cv}^{ev} \end{cases} \tag{2.12}$$

Calculating the average maximum power of the EV charger over the next time step when the $SOC_{cc,cv}^{ev}$ is exceeded would introduce an exponential relationship to the MPC model. The linearity of the MPC model is kept by limiting the maximum power of the EV based on the SOC of the EV at the following time step $t + 1$ instead.

A constraint requires the modelled SOC of the home battery to be charged to 100% at the next switching time. A similar constraint is imposed to ensure the EV battery reaches the forecasted final SOC at the forecasted departure time. This forecasted SOC and departure time can be underestimated, leading to the activation of the enforced charging behaviour (see subsection 2.2.2.5). For this reason, the approximated enforced EV charging power for the next time step (see subsection 2.2.4) is set as the minimum EV power setpoint in the MPC model. Finally, a constraint requires the modelled grid power not to exceed the maximum grid power of 9.2 kW.

Forecasting of electrical load and PV generation are accomplished using a Light Gradient Boosting Machine (LightGBM) model [30]. The model predicts the next 96 quarter-hourly values (totalling 24 hours) using the previous 96 quarter-hourly values, the time of the day and the day of the week. An optimal combination of hyperparameters is found using the Tree-structured Parzen Estimator [3]. The ranges of hyperparameters used for the optimization are shown in Table 2.4. We refrain from using the weather data provided by external services as an input feature for the PV forecast to avoid possible errors, such as the external server being down for maintenance. For electrical load prediction, the model is trained using the original data (see subsection 2.2.2.1) and therefore the same for the four houses while for PV generation the model is trained on a building-by-building basis. The forecasting model for EV charging duration and demands employs k-Nearest Neighbors (k-NN) [10] regression. The k-NN models utilize the initial SOC and the hour of the day as features, leveraging cosine similarity to find the most similar historical patterns. This model has previously shown superiority in data-scarce applications, such as the prediction of EV charging session parameters [24].

Table 2.4: Hyperparameter Ranges for LightGBM Model

Parameter	Range
Learning Rate	0.001, 0.01, 0.1, 0.2, 0.3
Num Leaves	32, 64, 128, 256
Max Depth	2, 4, 6, 8, 10
ColSample ByTree	0.6, 0.8, 1.0
Min Child Samples	1, 2, 3, 4, 5

2.2.5.4 Reinforcement learning

The specific reinforcement learning algorithm used in this case is the twin delayed deep deterministic policy gradient (TD3) [22], with its implementation taken from the stable baseline [42] library.

Every 15 minutes, the algorithm receives the following inputs from the system (instantaneous values at the start of each time step):

- Non-controllable electrical load [kW]
- PV power [kW]
- State of charge of the home battery [%]
- State of charge of the EV [%]
- Belgian day-ahead hourly price (Vd_t) [€/kWh]
- Hour of the day [h]
- Day of the week [-]

At each time step, the RL algorithm selects an action for both the home battery and the EV charger. This action is a continuous value representing the power setpoint in kW, constrained by the respective maximum charge and discharge limits of the home battery and the maximum charge limit of the EV charger.

At the end of each time step, the RL algorithm receives a reward R_t based on the previous time step t. The reward is computed as follows:

$$R_t = -(\underbrace{Eo_t Vo_t - Ei_t Vi_t}_{\text{day-ahead cost}} + \underbrace{Eo_t Vf}_{\text{offtake extras cost}}) - z_t \tag{2.13}$$

Where $z_t = 1e$ if the safety layer corrected the setpoint of either the home battery or the EV charger, and $z_t = 0e$ otherwise. This gives a significant penalty to unsafe behaviours, encouraging the RL model to avoid them.

The RL agent is initially pre-trained using Behavioral Cloning [2], with the simulated data of the RBC EMS over the three-month training period to reproduce the RBC's behaviour. This approach reflects a practical scenario where only historical rule-based EMS data is available, and no simulator is present, highlighting the RL algorithm's ability to learn directly from real-world data without requiring a system model. The pre-training was performed using an 80%-20% train-test split and trained over a period of 7 months and 6 days of data (i.e., three times the training data without the test data).

The TF3 algorithm was implemented with the hyperparameters described in table 2.5.

Table 2.5: TD3 hyperparameters. The parameters are the result of an optimization study from Ceusters et al. [7], yet with an increased `train_freq` and reduced `learning_starts`.

Parameters	**TD3**
gamma	0.7
learning_rate	0.000583
batch_size	16
buffer_size	1e6
train_freq	(4, "step")
noise_type	normal
noise_std	0.183
learning_starts	96

2.2.5.5 TreeC

Using the TreeC method, one decision tree is generated to control the home battery and another one to control the EV charger for each house. The decision trees used in the experiment are generated using the training data of subsection 2.2.5.1, the simulator of subsection 2.2.4 and the objective function described in eq. (2.7).

Two changes have been made to the original TreeC method described in **??**:

- The metaheuristic algorithm Covariance matrix adaptation evolution strategy seemed to end up in local minima during different training runs regularly. The covariance matrix adaptation evolution strategy was therefore replaced by the metaheuristic algorithm particle swarm optimization [31], which obtained better training results. This study used the particle swarm optimization algorithm implemented in the Python library pygmo [4] with default parameters which corresponds to the canonical particle swarm optimization algorithm described in [40].
- At the end of the training, the best decision tree is pruned by removing all leaves that do not make the objective function score worse than 1%

compared to the unpruned decision tree. The least used leaves are tested for removal first.

The splitting features of the decision trees are the following:

- Non-controllable electrical load (P_{t-1}^{load}) [kW]
- PV power (P_{t-1}^{pv}) [kW]
- Home battery power (P_{t-1}^{batt}) [kW]
- EV charging power (P_{t-1}^{ev}) [kW]
- State of charge of the home battery (SOC_t^{batt}) [%]
- State of charge of the EV (SOC_t^{ev}) [%]
- Belgian day-ahead hourly price (Vd_t) [€/kWh]
- Hour of the day [h]
- Day of the week [-]
- Difference with minimum day-ahead price [€/kWh]

All power values are the average power over the previous 15-minute time step (i.e., from $t-1$ to t). The SOC, price and time values are instantaneous values at the start of the time step t. The "difference with minimum day-ahead price" feature represents the day-ahead price shifted to have the minimum price of the ongoing experiment day equal to 0. This additional feature is calculated by subtracting to the current day-ahead price the minimum price of the experiment day period ranging from the previous switching time to the next switching time.

The actions of the TreeC EMS include charging or discharging the home battery at different power levels and also the possibility to do self-consumption. The action space is normalized between 0 and 1. If the action value is less than 0.1, then it does self-consumption; otherwise, it charges or discharges the home battery at the power level corresponding to the denormalized action value.

The TreeC EMS is trained in simulation with the particle swarm optimization algorithm over 1000 generations and a population of 1000 individuals. This means 1 000 000 decision tree based EMSs are evaluated in simulation on the three months of training data and the best one is then selected. The population size of 1000 was chosen as it obtained better results for the real-world and non-unimodal benchmarks of the IEEE Competitions in Evolutionary Computation compared to the more historically implemented population sizes of 20-50 individuals [39]. As in classical TreeC, five different EMSs are generated for each house

through training then pruning. The best-performing EMS of the five is then used in the experiment.

The TreeC EMS obtained after training for house 1 is presented in fig. 2.4.

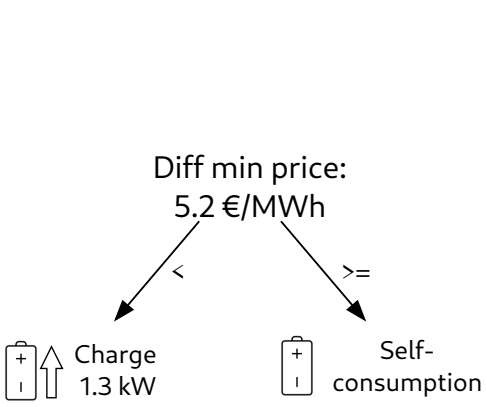

(a) Decision tree for the home battery

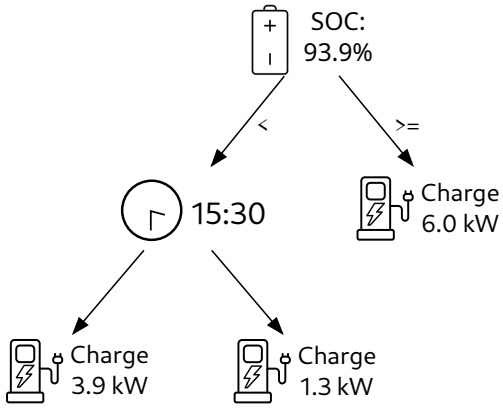

(b) Decision tree for the EV charger

Figure 2.4: The TreeC EMS obtained after training and pruning for house 1.

The obtained EMS is interpretable. The home battery charges when the electricity price is close to the minimum price for the experiment day and otherwise does self-consumption. The EV charger sets a high charging power when the home battery is almost full; otherwise, it sets a low charging power after 15:30 and an intermediate charging power before 15:30.

Here below are the EMSs obtained for the other houses.

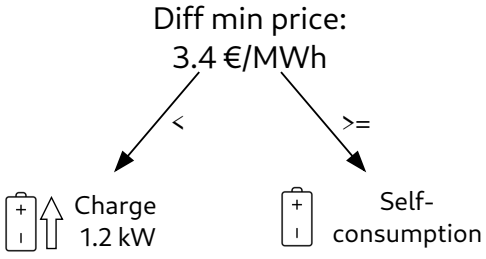

(a) Decision tree for the home battery

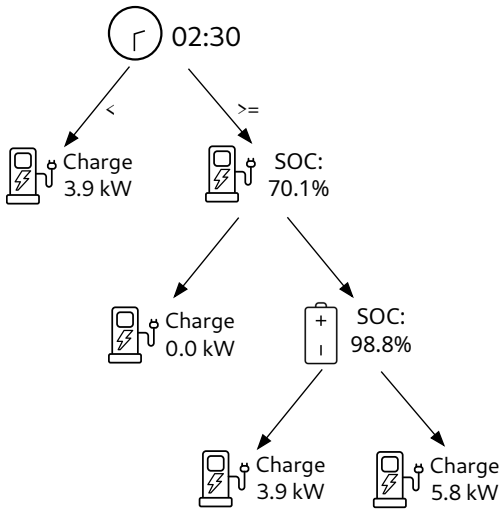

(b) Decision tree for the EV charger

Figure 2.5: The TreeC EMS of house 2.

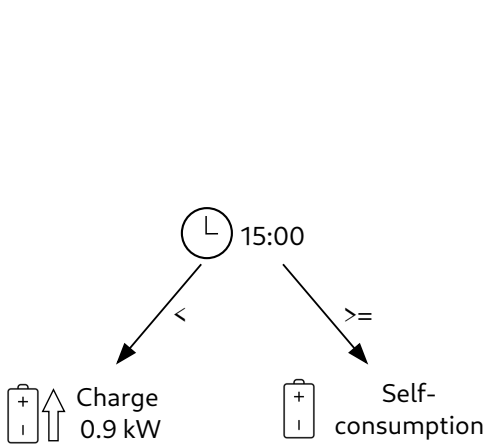

(a) Decision tree for the home battery

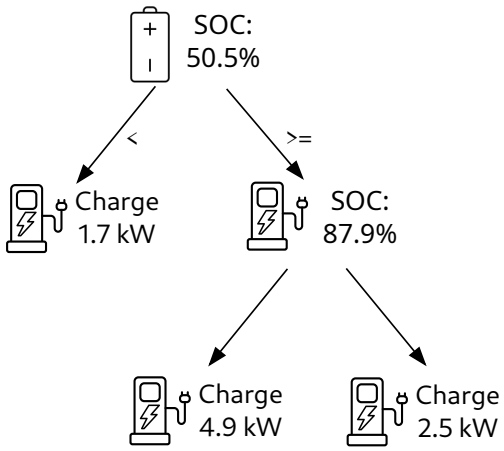

(b) Decision tree for the EV charger

Figure 2.6: The TreeC EMS of house 3.

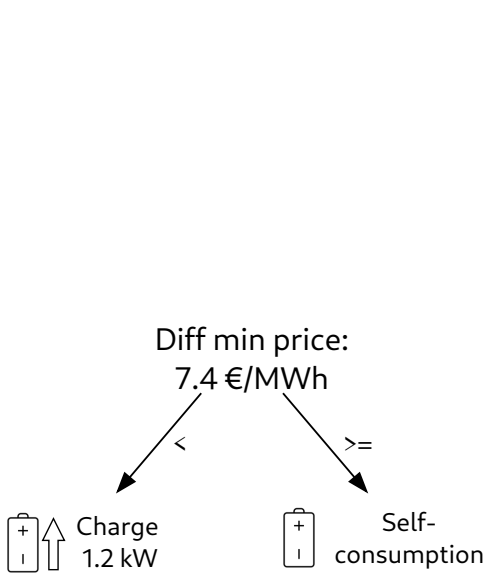

(a) Decision tree for the home battery

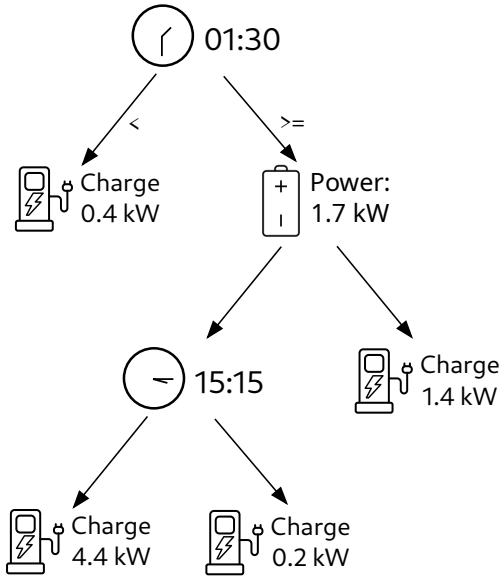

(b) Decision tree for the EV charger

Figure 2.7: The TreeC EMS of house 4.

2.2.6 Experiment execution

The experiment was conducted from 2024/04/11 at 15:00 to 2024/06/17 at 15:00 to obtain 48 experimental days fit for comparison. In total, 19 experimental days were not used due to home batteries not reaching 100% at the switching time, missing data or circuit breaker trips. The reasons for these failures were the following:

- Server restart interrupting the execution program (2 days)
- Battery of house 1 stopped responding until manual intervention (2 days)

- Battery of house 4 stopped responding when fully charged, a bug fix was applied to limit the SOC to 95% (1 day)
- The bug fix for house 4 did not fix the issue. A new fix was applied to perform self-consumption above 95% SOC (2 days)
- A circuit breaker from the dynamic loads circuit of house 2 tripped because the breaking capacity was too low for high loads; a circuit breaker with a higher breaking capacity was used instead (6 days)
- A previously uncaught error caused the experiment program to stop (2 days)
- The main circuit breaker of house 3 blew because the apparent power of the house was too high for too long (2 days)
- Missing grid measurements due to a malfunction of the grid meter of house 1 (2 days).

Some changes have been made to the experiment during its execution. As listed above, we encountered an unexpectedly high amount of reactive power from the dynamic loads within our hardware setup (fig. 2.1), causing a circuit breaker trip. This was because the safety layer was formulated using active power rates, essentially assuming a perfect power factor of 1. Therefore, halfway through the experiment, we incorporated the reactive power in the constraints - imposing limits on the grid's apparent power instead to limit the maximum main circuit breaker current correctly. Extra measurements were taken on the main grid breaker, and the grid limit of 9.2 kVA apparent power corresponds to approximately 8.7 kW active power. The maximum grid power in simulation and the MPC constraints was therefore changed from 9.2 kW to 8.7 kW. The correction was made on the 2024/5/21 at 15:00.

Two other minor changes were made during the experiment. The EV SOC input used by the TreeC EMS had a value 100 times too low. This was corrected and had an impact on only one charging session without affecting the performance much. The error was also reproduced in the simulation (Corrected 2024/4/15 15:00:00). The maximum SOC for the battery of house 4 was limited to 95% as described in the bullet points above (First correction 2024/04/22 15:00:00, second correction 2024/04/30 15:00:00).

Figure 2.8 shows the RBC EMS controlling house 3 for an experiment day. The figure shows the power profiles of the different assets for this experiment day, the SOC of the EV and home battery, the enforced charging behaviour at the end of the day for the home battery and two periods when the safety layer corrected the EV and home battery powers.

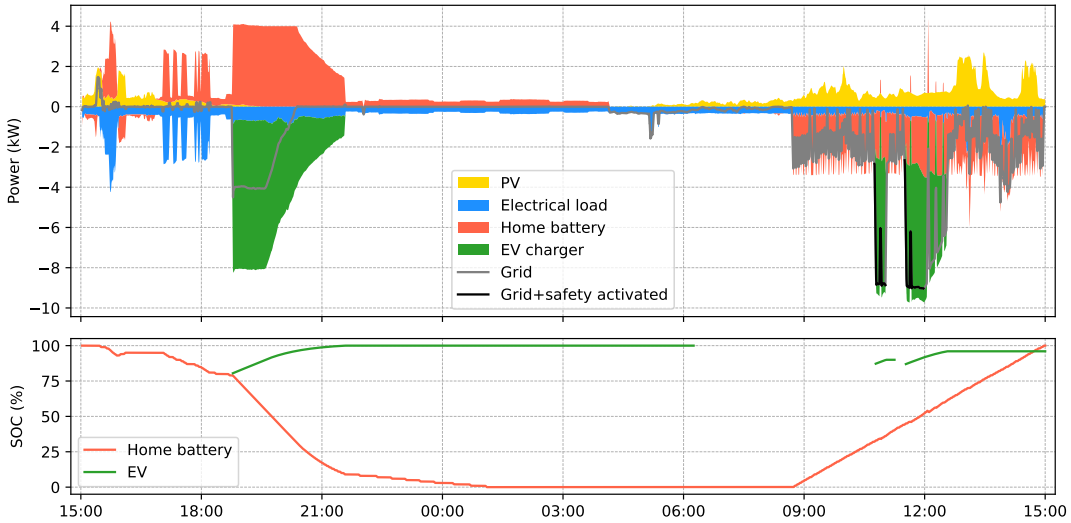

Figure 2.8: Representation of the experiment day starting on the 30th of May at 15:00. The upper plot shows the power profiles of the different assets of house 3 controlled by the RBC EMS. The grid power is represented in grey when the safety layer is not activated and in black when it is. The bottom plot shows the SOC of the home battery and EV. The enforced charging behaviour of the home battery is executed at the end of the experiment day to reach the 100% SOC goal at 15:00.

2.3 Results and discussion

2.3.1 EMS PERFORMANCE COMPARISON

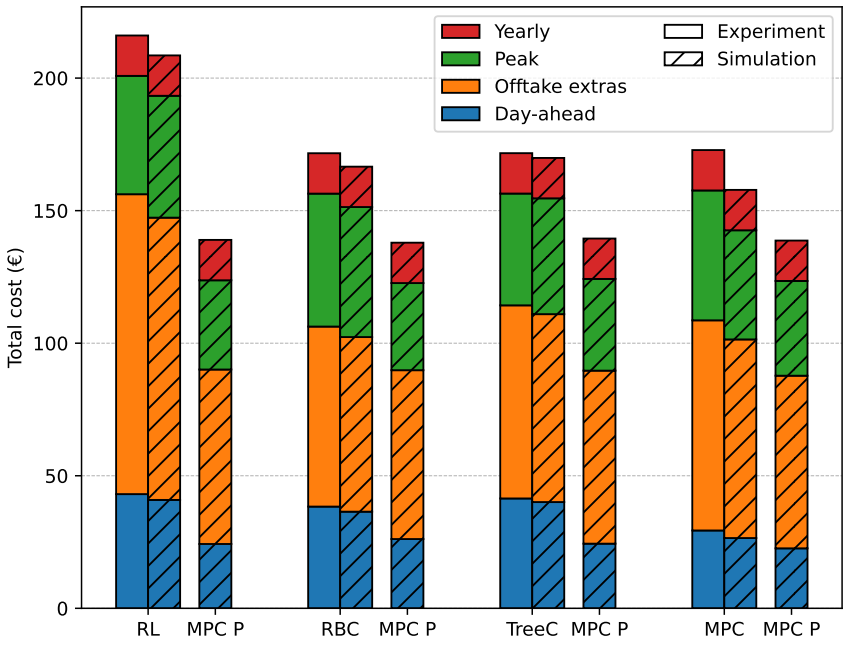

Figure 2.9: Costs of electricity for each EMS over the 48-day experiment from the real measurements and simulation. The total cost of electricity is divided into the four different costs from eq. (2.7). In addition to EMSs costs, the simulated costs of an MPC EMS with perfect forecast is also shown with the label MPC P. The MPC P costs are simulated using the same set of 48 days and house combinations as the EMS on its left and serve as benchmarks.

The RBC, TreeC and MPC EMSs obtained similar performances in the experiment as shown in fig. 2.9. The RL EMS obtained a worse performance than the other EMSs, which is expected as the RL performed more exploratory actions than the other EMSs. For the three other EMSs, them being close is unexpected but is explained in the following sections.

No EMS is close in performance to the MPC P benchmarks, indicating better implementations or methods could still be found. The MPC P EMSs obtained scores between 137.9€ and 139.5€ even though they were evaluated on different day house combinations (i.e. the same day house combinations than the EMS for which they serve as a benchmark). The score difference is small, which indicates that the comparison procedure is fair.

Table 2.6: Electrical energy metrics for each EMS over the 48-day experiment from the real measurements. Net consumption is defined as the difference between grid-imported and grid-exported energy.

Result	RL	RBC	TreeC	MPC
Grid imported [kWh]	991.7	595.3	638.6	694.9
Grid exported [kWh]	458.3	91.3	142.2	190.5
Net consumption [kWh]	533.4	504.0	496.3	504.4

The results of table 2.6 show that the RL EMS imported more electricity from the grid compared to the other EMSs, and it also has a higher net consumption. A higher net consumption is most probably caused by a higher use of the battery, which would result in more efficiency energy losses.

2.3.1.1 Rule-based control

The strategy of the RBC is to do as much self-consumption as possible, which means trying not to take electricity off the grid as shown in table 2.6. This strategy is reflected in a comparatively low offtake extras cost, which is proportionally the highest cost category for all EMSs.

2.3.1.2 Reinforcement learning

The RL EMS obtained the worst score, mainly due to a higher offtake extras cost (distribution and taxes) as can be seen in fig. 2.9 - which is not unexpected as we observed early-stage learning despite the pre-training, given the new (real) experience tuples. Moreover, even though the RL agents were pre-trained offline with a 3-month-long historical dataset, each agent-house combination only had an experimental runtime of 1.7 weeks (12 days). When comparing to relevant simulated case studies, more required training time is also observed. For example, Ceusters et al. [9] reported 95% MPC-like performance of the RL EMS after $\sim$ 18.6 weeks and only matched the MPC EMS after $\sim$ 2 years - also using 15-minute time steps.

2.3.1.3 TreeC

The TreeC EMS obtained the lowest peak cost but couldn't obtain as low offtake extras and day-ahead costs as the other MPC and RBC EMSs. Looking at the decision trees obtained in figs. 2.4 to 2.7, the low peak cost is explainable by lower charging setpoints of the EVs.

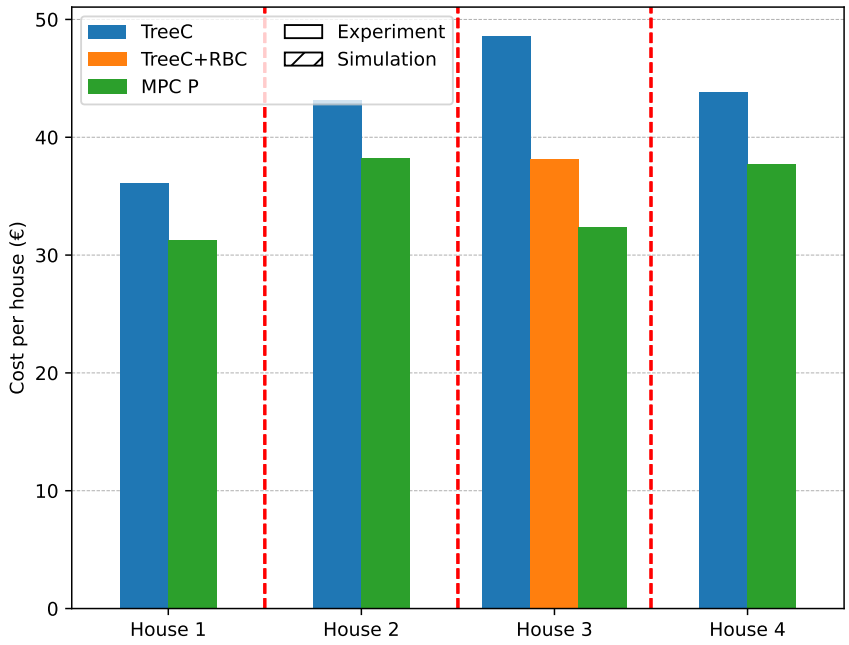

Figure 2.10: Comparison for each house of the total cost obtained with the TreeC EMS and the benchmark MPC P EMS. For house 3, the TreeC+RBC EMS is added to the comparison and represents the TreeC EMS using the self-consumption control strategy of the RBC EMS to control the home battery.

Intuitively, it is surprising that the TreeC EMS did not perform better than the RBC EMS. Figure 2.10 shows that the benchmark costs of the MPC P EMS are 10-15% lower than the TreeC EMS costs for houses 1,2 and 4. For house 3, though, the MPC P EMS cost is 33% lower than the TreeC EMS cost, showing a much larger difference than for the other houses. Looking at the decision trees obtained from the TreeC training (see figs. 2.4 to 2.7), the main difference between houses 1,2 and 4 and house 3 is the home battery strategy. Houses 1,2 and 4 charge their battery when the electricity price is very low and otherwise do self-consumption. House 3, on the other hand, does self-consumption until midnight and then charges the battery until the switching time at 15:00. This strategy could well be valid in the training months of January to March when there is less PV production but during the experiment months of April to June, the strategy would lead to more offtake from the grid and, therefore, offtake extras costs compared to the self-consumption strategy of the RBC EMS. When replacing the TreeC home battery strategy of house 3 with the self-consumption strategy of the RBC EMS in simulation, the MPC P cost is now only 15% lower for that house as shown in fig. 2.10. The TreeC+RBC policy for house 3 would have reduced the total cost of the TreeC EMS by 10.5€, making it better than the RBC EMS over the total cost of the 4 houses (see fig. 2.9 for original total costs).

Overall, this shows the importance of using data from a training period that is representative of the experiment or test period. In this case, the training data was not representative enough of the experiment, which led to the TreeC EMS obtaining a bad control strategy for the home battery of house 3.

2.3.1.4 Model predictive control

The MPC EMS obtained the lowest day-ahead costs but couldn't obtain as low of a peak and offtake extras costs as the RBC and TreeC EMSs. The low day-ahead cost is expected as the MPC knows the future Belgian day-ahead prices of electricity and can optimize accordingly. Intuitively, it is surprising that the MPC did not optimize the other costs efficiently.

Figure 2.9 shows a bigger cost difference between simulation and reality than the other EMSs and obtained lower peak, offtake extras and day-ahead costs. Certain days had cost differences exceeding 1€, as seen in fig. 2.11. A deeper analysis was performed to understand these differences on three days that had this cost difference for houses 1 and 3. House 4 also had two days when the simulation results were worse than the experiment by more than 1€. However, on one day, the simulation outperformed the experiment by approximately 2€, which mitigates the impact of the two worse days.

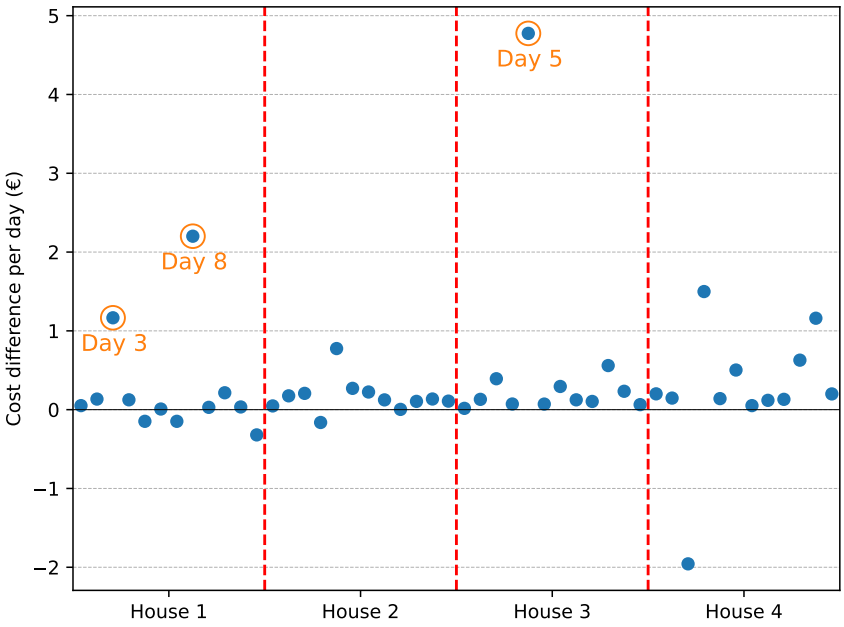

Figure 2.11: Difference in costs between the experiment and the simulation for the MPC EMS for each day.

In all three instances, the main driver behind the cost differences is the peak cost.

- Day 5 of house 3, a difference of 4.8€: On the first time step of the day, a new short charging session forced the EV charger to charge at maximum power. In the simulation, this peak was avoided by discharging the home battery at the same time, but this did not occur in the experiment. The MPC EMS in the experiment decided on the setpoint for the home battery and EV charger without receiving the information that an EV session would start at that time step and, therefore, did not set an appropriate home battery setpoint. This is a coding error that would happen when an EV session starts at the first time step of the day.

- Day 8 of house 1, a difference of 2.2€: The peak of the day was higher in the experiment than in the simulation because the MPC set a higher charging power for the EV charger than in the simulation at midnight. The MPC EMS forecasted that there would be some PV production at that time because it received the PV production of the previous day in a wrong order. This wrong order was due to the device failing to take the grid measurements. The device failed to collect data from the 29th of May at 10:45 to the 30th of May at 06:45, and the investigated day started on the 30th of May at 15:00. The forecaster of the MPC requested the data from the previous 24 hours, but due to the missing data and the formatting functions not sorting by default on time; the data was received starting from 06:45 until the time of request and then from 24 hours before the time of request until 06:45. This mixed order caused the forecaster to predict a PV production between 300W and 400W during nighttime and at a time when the MPC EMS decided to charge the EV at power close to the previous peak, thus obtaining a higher peak in the experiment than in the simulation for the day.

- Day 3 of house 1, difference of 1.7€: Close to the end switch time of the investigated experiment day, the MPC EMS managed in simulation to charge the home battery at 100% SOC and complete an EV session without increasing the peak cost. In contrast, in the experiment, the peak was increased due to bad management of the EV session and home battery 100% SOC switching time constraint. The error came from a difference in the experiment and simulation MPC. Whenever the forecaster would predict a departure time later than 15:00, the simulation MPC would set the departure time to 15:00 instead. This correction was mistakenly not applied in the experiment MPC.

These errors were reproduced in simulation, and the total cost went from 157.8€ to 165€. This brings the simulation cost closer to the experimental total cost of 172.8€ by half of the original difference. A better implementation of the MPC EMS in the experiment would improve its performance. On the other hand,

this shows that the peak cost can penalize the performance of the MPC if there are some errors in the implementation. EMSs based on less complex models and inputs, such as the RBC and TreeC EMS, did not experience similar errors. Also, it would have been difficult to notice these errors without a simulation reproducing the MPC behaviour.

2.3.2 Difference experiment vs simulations

The following section analyses differences between simulation and experiment in addition to the ones already pointed out specifically for the MPC EMS in subsection 2.3.1.4.

The experiment runs in real-time while the simulation approximates this behaviour on a 15-minute time step. In the experiment, the grid sensor measures both imported and exported energy for the same time step, while in simulation, there is only one net consumption value per time step. The net consumption is the difference between the imported and exported energy from the grid. Calculating the cost for one time step using both the imported and exported energy instead of the net consumption is always more expensive because offtake extras costs get removed when doing the energy difference for net consumption. Table 2.7 shows the score obtained in the experiment for each EMS when using the import and export energy values as well as the score obtained when using the net consumption to understand if the approximation of the simulation has a big impact on the results.

Table 2.7: Costs of electricity obtained for each EMS in the experiment when taking the export and import energy values and when calculating net consumption.

EMS	Experiment (import/export)	Experiment (net consumption)	Simulation (net consumption)
RL	216.1€	207.2€	208.6€
RBC	171.7€	171.2€	166.6€
TreeC	171.7€	170.7€	169.9€
MPC	172.8€	170.1€	157.8€

The cost difference between import/export and net consumption is larger for the RL and MPC EMSs than for the TreeC and RBC EMSs. A major behavioural difference between the MPC and RL EMSs with the TreeC and RBC EMSs is that the former control the home battery's setpoint while the latter often or exclusively do self-consumption. Doing self-consumption allows the battery to adapt to real-time changes and has no or very little imported and exported energy when the battery is not full or empty. Giving a fixed setpoint to the home battery for a whole time step, on the other hand, does not allow the battery to adapt to

real-time changes and can lead to both imported and exported energy during the same time step.

A solution to this problem could be to require the home battery to maintain the grid power at a certain setpoint instead of keeping a fixed battery power for the whole time step. Further simulation and experimentation would be needed to confirm if this is a good solution.

Both the MPC differences from subsection 2.3.1.4 and the grid energy exchange calculation difference from this section explain some differences between the experiment and the simulation. Other differences observed between experiment and simulation but more difficult to quantify in terms of cost include:

- Battery capacities and efficiencies not being exactly the same as on the specification sheet.
- Batteries 1 and 2 having a clear minimum soc of 5%.
- Battery 4 having a variable minimum SOC, an upped soc of 95% to avoid the battery not responding and a recurrent jump from 85% to 95% SOC due to internal recalibration.
- The real-time safety layer of the experiment is approximated on a 15-minute time step in the simulation.
- The hybrid inverter's power limit is not modelled in the simulation.

With the MPC errors of subsection 2.3.1.4 reproduced in simulation, the difference in total cost between experiment and simulation ranges from 1% to 4.5% for all EMSs. When calculating the experiment's total cost based on the net consumption, this range goes down to between 0.5% to 3% difference. The remaining difference between the experiment and simulation can be explained by the points listed above and some other unnoticed differences but, in total, they do not account for a large difference in costs. As stated in subsection 2.2.1, the electrical loads and possible power losses of the system are not separately measured; therefore, the simulation uses data combining the two. The possible power losses could depend on the home battery and EV usage, which would not be reflected in the simulation and needs to be further investigated in future work. Nonetheless, the simple battery model (see eq. (2.9)) and the capacity and efficiency parameters provided by the battery manufacturers (see table 2.1) were sufficiently accurate to perform meaningful simulations.

2.4 Conclusion

This chapter presented the results of the real-world experimental validation of a novel machine learning method using a metaheuristic algorithm to generate an

interpretable control policy modelled as a decision tree (i.e., TreeC [45]). This, together with the comparison between an MPC, a RL, an RBC and an explainable Tree-based EMS on any case study, simulated or otherwise, using a novel evaluation procedure, aimed to be as fair as possible. We come to the following conclusions:

- The RBC, TreeC and MPC EMSs obtained a similar economic performance with only a 0.6% difference in cost between the three. The results highlight the importance of representative training data for the TreeC EMS and expert implementation for the MPC EMS to obtain better performances in real implementations;
- A pre-trained, yet model-free, RL agent still has a considerably long training time (even with hyperparameters chosen for a higher learning rate) and, therefore, monetary training cost that would need to be set off against the extra (if any) cost of other methods (e.g. modelling time);
- With the MPC experiment errors reproduced in simulation, the difference in costs between experiment and simulation ranges from 1% to 4.5% for all EMSs. This difference is in part explained by the simulation using a single net consumption value per time step for the energy exchanged with the grid, while the experiment measures both imported and exported energy per time step. Overall, the simulations rather accurately represented reality in terms of cost and were essential to finding the MPC errors and TreeC training data issues;

3.1 Introduction

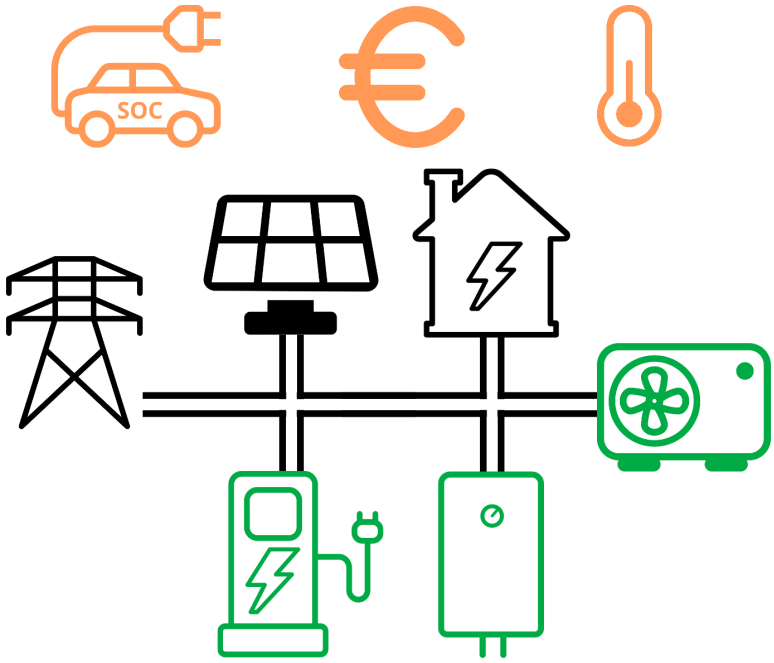

Figure 3.1: Representation of the energy grid analysed in this chapter. The optimized metrics indicated in orange are EV charge completion, electricity cost, and thermal comfort. The controlled assets indicated in green are the EV charger, water heater, and heat pump.

An effective management of energy assets in residential buildings is becoming increasingly relevant in Belgium. This is driven by the increasing electrification of our energy systems, the adoption of dynamic electricity pricing, and the implementation of grid peak tariffs in the residential sector [44]. As these systems and regulations become more complex, the question arises: What are the benefits of implementing an advanced energy management system in electrified dwellings ?

This chapter aims to answer this question by simulating a set of representative Belgian dwellings which have been electrified and are subject to dynamic pricing and grid peak tariffs. The electrification of the dwellings is achieved by replacing the traditional heating systems with air-to-air heat pumps, adding domestic

electric vehicle chargers due to the replacement of petrol cars with electric vehicles, and replacing gas water heaters with electric water heaters. Other dwelling characteristics such as the presence of photovoltaic panels, insulation level, and construction type are kept close to current Belgium situation according to the most recent statistics we could find. Real-world datasets are used for base electricity consumption, domestic water usage, electric vehicle charging sessions, photovoltaic production, weather data, and electricity tariffs. Two years of data (2022 for training, 2023 for evaluation) are used, reflecting a realistic scenario where a household adopts an EMS in January 2023 based on data from the previous year. Home batteries are excluded from the analysis, as their adoption remains limited in Belgium. According to available statistics, about 3% of dwellings had a home battery in 2024 [23]. Replacing gas boilers with heat pumps and petrol cars with electric vehicles is essential for reducing fossil fuel use, while the case for home batteries is less evident.

This study compares three distinct energy management system (EMS) strategies: rule-based control (RBC), which mimics the default operation of residential energy assets; model predictive control (MPC), a state-of-the-art optimization-based approach; and TreeC, a machine learning-based method that models EMS strategy as decision trees. In addition to these realistic EMSs, a model predictive control EMS with perfect forecast (MPC P) is also compared to show a maximum performance an EMS could technically achieve.

Buying a home EMS through a professional is advertised as costing between 280 € and 680 € with some additional subscription costs ranging between 0 € and 360 € per year [34]. An EMS can be implemented cheaper using for example Home Assistant and the additional EMHASS module [26, 13] for which the basic hardware costs close to 100 € but usually requires some technical knowledge to set up and configure. To these basic costs can be added the maintenance, more expensive compatible equipment and additional necessary sensors.

This chapter addresses the following research questions for a representative set of electrified Belgian dwellings:

- What are the benefits of EMSs like TreeC and MPC compared to a simple RBC in terms of cost savings, thermal comfort and EV charging satisfaction ?
- How much can algorithmic improvements realistically improve these benefits ?
- What is the influence of different dwelling characteristics on cost savings ?

3.2 Building a representative set of Belgian dwellings

This section describes the methodology for constructing a set of 500 representative Belgian dwellings. Each dwelling is first assigned a set of characteristics (such as region, number of rooms, and insulation level) based on the latest available Belgian statistics (see subsection 3.2.1). Using these characteristics, a simulation model is created for each dwelling, incorporating real-world data sources (see subsections 3.2.2 to 3.2.6). This process yields 500 unique simulation models, each reflecting its own set of characteristics and data inputs. The overall methodology is illustrated by the flow chart in fig. 3.2.

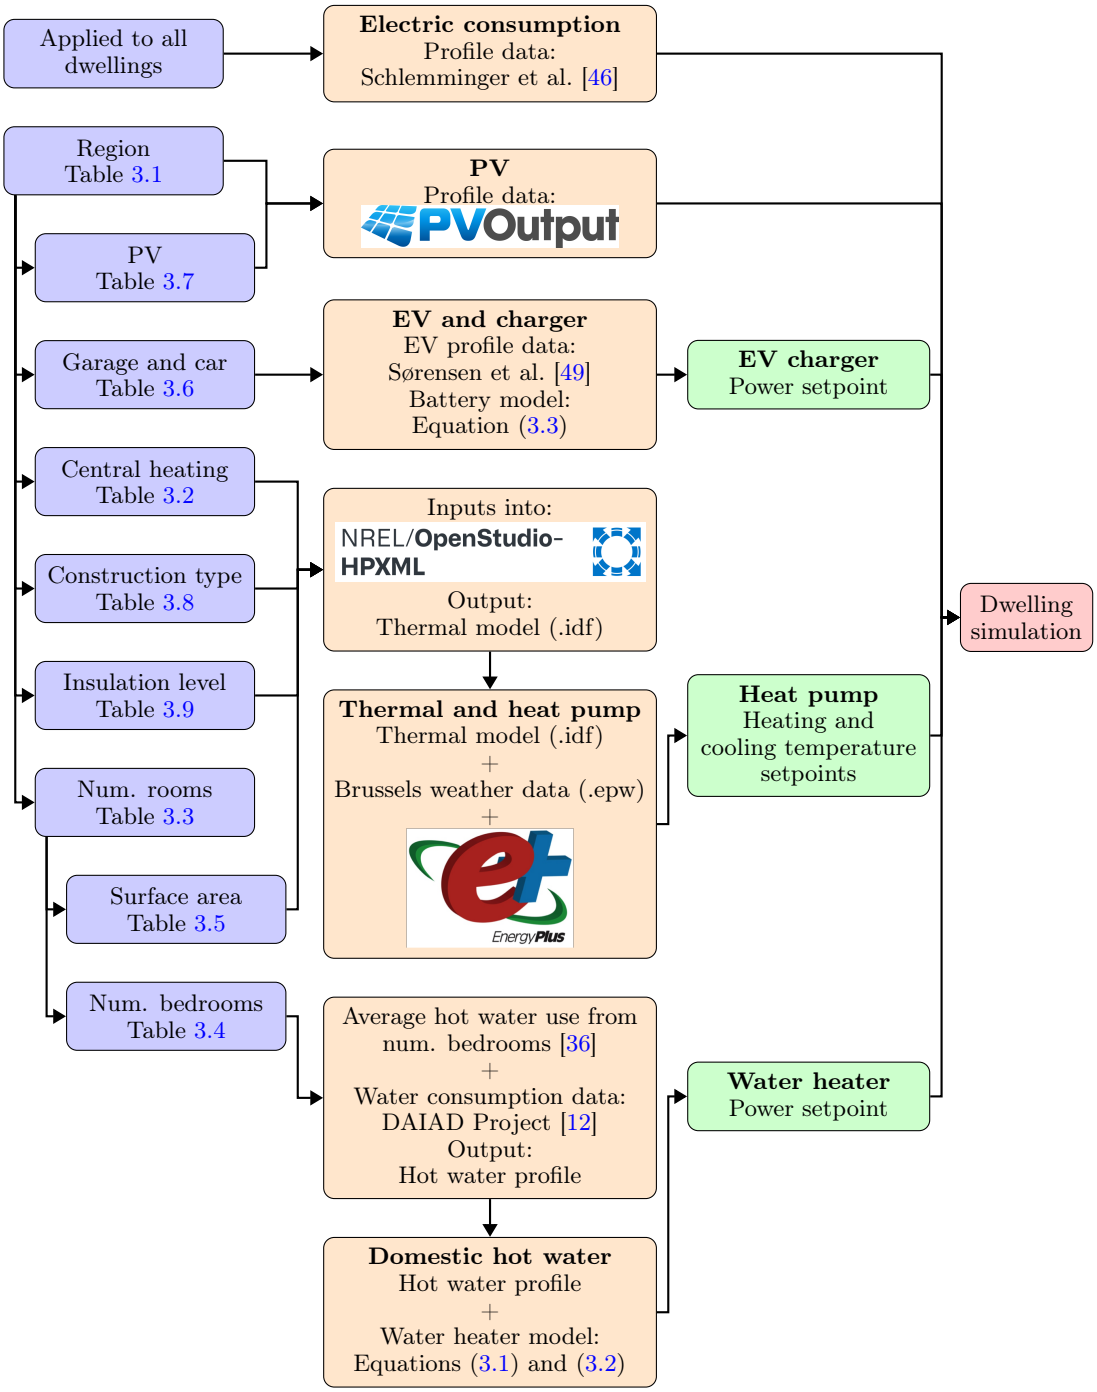

Figure 3.2: Flow chart of this study's methodology to simulate a dwelling. Blue nodes represent the sampled characteristics of the dwelling (subsection 3.2.1), orange nodes represent the data and simulation models used for each part of the dwelling (subsections 3.2.2 to 3.2.6), green nodes represent the possible control actions of the assets (section 3.3), and the red node is the dwelling simulation that combines everything together.

3.2.1 Dwelling characteristics

Different statistics were collected to build the set of 500 representative Belgian dwellings. This section goes through each of them giving the source and explanation when necessary. In the following tables, each cell shows the percentage of all dwellings that have both the characteristic specified by the row and the characteristic specified by the column.

To assign characteristics to the 500 representative dwellings, the number of dwellings with each characteristic is calculated by multiplying the relevant percentage by 500. For example, according to table 3.1, 10.6% of dwellings are in Brussels, 57.4% in Flanders, and 32.0% in Wallonia, resulting in 53 dwellings in Brussels, 287 in Flanders, and 160 in Wallonia. This establishes the initial pool of unassigned characteristics.

Assignment proceeds sequentially, starting with the most independent characteristic (region), and moving towards the most dependent (number of bedrooms and surface area) as illustrated by fig. 3.2. For each dwelling, characteristics are sampled randomly from the remaining pool, using any already assigned characteristics as conditional constraints. This process continues until all dwellings have a complete set of characteristics, and the overall distribution matches the target statistics.

The percentage for the number of dwellings per region (table 3.1), the presence of central heating (table 3.2) and the number of rooms (table 3.3) are all obtained from the 2021 Belgian Census which gives the total number of dwellings fitting each characteristic without need for extrapolation [52, 51].

The number of bedrooms data from table 3.4 comes from an extrapolation of Meireles et al. [36]. The source gives the number of bedrooms of 9 206 Belgian apartments. To fit better to the global Belgian situation with 66% houses (see table 3.8), the total proportion of 0 and 1 bedroom dwellings was reduced by 25% and the total proportion of 3 to 6 bedroom dwellings was increased by 47.3% to match the reduction of the 0 and 1 bedroom dwellings. The correlation between the number of bedrooms and the number of rooms is established using common sense with the total bedroom and total room proportions as a constraint (see table 3.3). Meireles et al. [36] indicates that the number of bedrooms in a dwelling is a good indicator of average water consumption. The number of bedrooms is used solely to determine the appropriate hot water consumption profile to assign to the dwelling, as detailed in subsection 3.2.3.

The surface area data from table 3.5 comes from tables IV.24 and IV.28 of Vanneste et al. [55]. Table IV.24 gives the number of dwellings with a certain surface area per region. Table IV.28 gives the number of dwellings with a certain number of rooms per region. The correlation between the number of rooms and

the surface area is determined to ensure that the final proportions per region closely match those presented in tables IV.24 and IV.28 of Vanneste et al. [55].

The garage and car ownership data from table 3.6 comes from page 68 of Hubert and Toint [28]. The source gives the proportion of dwellings that have a garage in each region and the proportion of people with a garage that have at least one car. Table 3.6 shows the combination of these two proportions, it is used to determine the proportion of dwellings that would have a EV charger in each region.

The photovoltaic data from table 3.7 comes from different sources for the different regions. The number of photovoltaic installations in Brussels and Flanders comes from Brugel [5] and Vlaams Energie- en Klimaatagentschap [56] respectively. For Wallonia, the proportion of dwellings with a photovoltaic installation is given by Lallemand [32].

The construction type data in table 3.8 combines two sources: Statbel [50], which provides the proportion of houses and apartments in each region, and Figure IV.13 from Vanneste et al. [55], which details the breakdown of terraced, semi-detached, and detached constructions by region.

The insulation data in table 3.9 is based on Recticel Insulation [43], which reports the proportions of dwellings with roof, wall, and ground insulation in each region. However, the source does not specify the exact overlap between these insulation levels (e.g., how many dwellings have both wall and roof insulation, or all three types). To construct the final table, the following assumptions were made: dwellings with ground insulation also have wall and roof insulation, and dwellings with wall insulation also have roof insulation. This approach ensures that the total proportions match the source data while providing a reasonable estimate of the distribution of combined insulation levels.

Table 3.1: Distribution of dwellings per region [52]

Brussels	Flanders	Wallonia
10.6%	57.4%	32.0%

Table 3.2: Distribution of dwellings with central heating [52]

	Brussels	Flanders	Wallonia	Total
With central heating	10.1%	51.9%	27.4%	89.4%
Without central heating	0.5%	5.5%	4.6%	10.6%

Table 3.3: Distribution of dwellings per number of rooms [51]

	Brussels	Flanders	Wallonia	Total
1 room	0.3%	0.5%	0.4%	1.2%
2 rooms	0.8%	1.6%	1.0%	3.4%
3 rooms	2.1%	4.6%	2.5%	9.2%
4 rooms	3.2%	10.0%	5.1%	18.3%
5 rooms	1.8%	14.0%	7.4%	23.2%
6 rooms	0.9%	12.6%	7.1%	20.6%
7 rooms	0.5%	7.3%	4.1%	11.9%
8 rooms	0.4%	3.6%	2.3%	6.3%
9+ rooms	0.6%	3.2%	2.1%	5.9%

Table 3.4: Distribution of dwellings per number of bedrooms [36]

Number of rooms	1	2	3	4	5	6	7	8	9+	Total
0 bedroom	0.4%	1.2%	2.5%	2.5%	2.0%	0.3%	-	-	-	8.9%
1 bedroom	0.8%	2.2%	5.6%	5.6%	4.6%	1.2%	-	-	-	20.0%
2 bedrooms	-	-	1.1%	8.2%	12.3%	10.3%	5.1%	2.1%	2.0%	41.1%
3 bedrooms	-	-	-	2.0%	4.0%	6.5%	5.1%	2.2%	2.1%	21.9%
4 bedrooms	-	-	-	-	0.3%	2.2%	1.4%	1.3%	1.0%	6.2%
5 bedrooms	-	-	-	-	-	0.1%	0.3%	0.6%	0.5%	1.5%
6 bedrooms	-	-	-	-	-	-	-	0.1%	0.3%	0.4%

Table 3.5: Distribution of dwellings per surface area [55]

Number of rooms	1	2	3	4	5	6	7	8	9+	Total
$<35m^2$	1.1%	1.7%	0.9%	-	-	-	-	-	-	3.7%
35-54m^2	0.1%	1.0%	2.8%	5.5%	4.6%	3.1%	-	-	-	17.1%
55-84m^2	-	0.7%	3.7%	7.3%	7.0%	5.1%	1.2%	-	-	25.0%
85-104m^2	-	-	1.8%	3.7%	7.0%	6.2%	3.6%	0.5%	0.5%	23.3%
105-124m^2	-	-	-	1.8%	2.3%	4.1%	3.6%	1.6%	1.5%	14.9%
$>124m^2$	-	-	-	-	2.3%	2.1%	3.5%	4.2%	3.9%	16.0%

Table 3.6: Distribution of dwellings with a garage and a car [28]

	Brussels	Flanders	Wallonia	Total
With garage and car	3.8%	42.7%	22.4%	68.9%
Without garage or car	6.8%	14.7%	9.6%	31.1%

Table 3.7: Distribution of dwellings with photovoltaic installation [5, 56, 32]

	Brussels	Flanders	Wallonia	Total
With photovoltaic	0.2%	19.2%	6.4%	25.8%
Without photovoltaic	10.4%	38.2%	25.6%	74.2%

Table 3.8: Distribution of dwellings per construction type [55]

	Brussels	Flanders	Wallonia	Total
Terraced house	1.3%	11.7%	8.4%	21.4%
Semi-detached house	0.2%	10.6%	5.2%	16.0%
Detached house	0.1%	18.6%	10.1%	28.8%
Apartment unit	9.0%	16.5%	8.3%	33.8%

Table 3.9: Distribution of dwellings by insulation level [43]

	Brussels	Flanders	Wallonia	Total
No insulation	1.3%	7.8%	11.5%	20.6%
Roof insulation	5.6%	16.5%	15.5%	37.6%
Wall and roof insulation	-	10.7%	0.2%	10.9%
Ground and roof insulation	0.1%	-	-	0.1%
Ground, wall and roof insulation	3.6%	22.4%	4.8%	30.8%

3.2.2 Dwelling thermal and heat pump models

The thermal models are created using the Openstudio-HPXML workflow [27]. This workflow takes as inputs the various characteristics of the dwelling, in this case construction type, insulation level and surface area, and returns and .idf file containing a detailed thermal model of the dwelling with a properly sized heat pump that can then be used for Energyplus [11] simulations (see fig. 3.2).

The roof and ground are insulated with an R-value of 6.4 m^2K/W and the walls with a value of 7.3 m^2K/W following the values from Recticel Insulation [43]. The uninsulated R-values are averages of typical Belgian building constructions: 0.5 m^2K/W for the walls, 0.3 m^2K/W for the roof and 0.4 m^2K/W for the ground, as reported by Énergie+ [16].

The windows U-values correspond to single glazing (5.8 W/m^2K), double glazing (2.8 W/m^2K) and triple glazing (0.6 W/m^2K) values from Energuide.be [17].

The air tightness values are taken from Laverge et al. [33] and correspond to the median air tightness values for Belgian houses:

- Constructed between 1990 and 1995 (9.3 ACH_{50})
- Constructed between 2007 and 2010 and randomly sampled (5.5 ACH_{50})
- Constructed between 2006 and 2013 and sampled from engaged owners that reported their air tightness levels (2.9 ACH_{50})
- With high energy performance (1.4 ACH_{50})

The link between the insulation level, the window U-values and the air tightness is shown in table 3.10.

Table 3.10: Window U-values and air tightness values

Insulation level	Window U-value	Air tightness
No insulation	5.8 W/m^{2}K	9.3 ACH_{50}
Roof insulation	2.8 W/m^{2}K	5.5 ACH_{50}
Double insulation	2.8 W/m^{2}K	2.9 ACH_{50}
Ground, wall and roof insulation	0.6 W/m^{2}K	1.4 ACH_{50}

Since only ranges of surface areas are given in the statistics from table 3.5, each range is assigned a specific value as input to the Openstudio-HPXML workflow. The surface area is set to 25m^2 for the dwellings of <35m^2, 150m^2 for the dwellings of >124m^2 and the average of the upper and lower bounds for the other surface area categories (e.g. 70m^2 for the 55-84m^2 range).

The dwelling's heating and cooling are managed by an air-to-air heat pump. This type of heat pump was selected due to its higher sales volume compared to air-to-water and geothermal heat pumps in Belgium [18]. The heat pump is modelled with a primary heating coil and a primary cooling coil, both operating via a compressor. Additionally, it includes a backup electric resistance heating coil that activates only when the primary heating coil cannot meet the heating demand. The heat pump is also modelled with a resistive defrost heater to prevent ice formation. The heat pump is controllable by setting a heating and cooling set point temperature.

The historical weather data for Brussels is used in all simulations and was downloaded using the Weather Data Downloader from Oikolab [37]. The obtained weather data file is in .epw format.

The simulation is conducted in EnergyPlus using the generated .idf thermal model file and the .epw Brussels weather data file as inputs. Subsection 3.3.1 details how the heat pump is controlled during the simulation.

3.2.3 Domestic hot water

The domestic hot water system is modelled as a storage tank following the equations from Dolan et al. [14] with a hot water consumption profile extrapolated from the DAIAD Project [12] dataset.

The hot water tank is modelled using the equations presented in eqs. (3.1)

and (3.2):

$$
\begin{aligned}
C &= \rho^w c^w V \\
G &= \frac{SA}{R} \\
B_t &= \rho^w c^w F_t \\
R'_t &= \frac{1}{G + B_t} \\
D_t &= \exp\left(-\frac{\Delta t}{R'_t C}\right)
\end{aligned}
\tag{3.1}
$$

where:

C [J/K]	tank heat capacity
ρ^w [kg/m^3]	water density (1000 kg/m^3)
c^w [J/kg/K]	specific heat capacity of water (4186 J/kg/K)
V [m^3]	tank volume
G, B_t, R'_t, D_t	intermediate parameters and variables
SA [m^2]	tank surface area
R [m^{2}K/W]	tank insulation R-value (used: 10 m^{2}K/W)
F_t [m^3/s]	water flow rate at time t
Δt [s]	time step duration

$$
T^{tank}_{t+1} = D_t \underbrace{T^{tank}_t}_{\text{init temp}} + \left(\underbrace{G R'_t T^{room}}_{\text{loss to room}} + \underbrace{B_t R'_t T^{inlet}_t}_{\text{loss inlet water}} + \underbrace{\eta^{heater} P_t R'_t}_{\text{gain heater}} \right) (1 - D_t) \tag{3.2}
$$

where:

T^{tank}_t [°C]	mean tank temperature at time t
T^{room} [°C]	ambient room temperature (used: 15 °C)
T^{inlet}_t [°C]	inlet (mains) water temperature at time t
η^{heater} [–]	heater efficiency (used: 1)
P_t [W]	heater power at time t

The DAIAD Project [12] dataset comprises two and a half years of hourly water consumption data from 1 007 households in Alicante. Meireles et al. [36] studied the water consumption of 9 206 Belgian apartments and gave a month dependant ratio between total water consumption and hot water consumption. The mean of the monthly ratios is 32.3% with a standard deviation of 2.5% between the months. The water consumption profiles from the DAIAD Project [12] dataset are converted to hot water consumption profiles by multiplying each month's data by the corresponding monthly hot-to-total water ratio.

The study also gives statistical links between water consumption and the number of bedrooms in the apartments. A probability distribution is created linking the

number of bedrooms to an estimated average hot water consumption in order to link a dwelling with a hot water consumption profile.

To sample the daily hot water consumption given a number of bedrooms in a dwelling, we use an inverse Gaussian distribution. An inverse Gaussian distribution is defined by its mean, shape and location parameters. Each bedroom count gets its own distribution, with the mean value taken directly from Meireles et al. [36]. The study also provides overall statistics (median, standard deviation, skewness, kurtosis) for hot water use across all apartments.

We assume the same shape and location parameters for all distributions. These parameters are fitted so the combination of distributions matches the median value reported in the original dataset. The obtained shape and location parameters are 67.68 and 0 respectively. The resulting probability density functions for each bedroom count are shown in fig. 3.3.

The mean and median of the combination of distributions is therefore the same as the original dataset but the standard deviation (77.2 for dataset vs 65.5 for distribution), skewness (61.3 for dataset vs 5.4 for distribution) and kurtosis (11.3 for dataset vs 66.6 for distribution) are still different. Knowing the exact distribution for each bedroom count would require the original data which is not available.

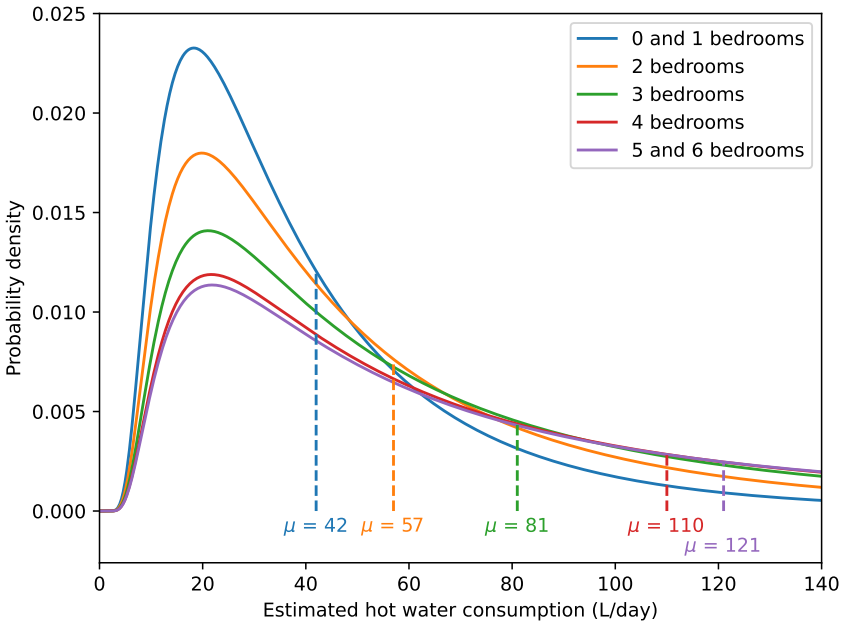

Figure 3.3: Probability density function of the estimated average daily hot water consumption per number of bedrooms in a dwelling. The mean value of the distribution is indicated by the vertical dashed line.

The hot water consumption profile for a dwelling is determined by sampling

from the inverse Gaussian distribution that matches the number of bedrooms in the dwelling. The profile selected is the one with the average daily hot water consumption closest to the sampled value.

In addition to the hot water consumption profile, the model from Dolan et al. [14] need a temperature for the incoming cold water from the mains. Fuentes et al. [21] established that a sinusoidal function, which provides the mains water temperature as a function of the month is a good approximation. The upper and lower bounds of the sinusoidal function are set to 11 °C and 16 °C respectively [53].

The water heater is modelled as an electric resistance heater with a heater efficiency (η^{heater}) of 1. The volume of the storage tank is three times the expected average daily hot water consumption determined by the number of bedrooms to account for peak consumption periods (so 146 litres for a dwelling with 1 bedroom and 363 litres for a dwelling with 6 bedrooms).

3.2.4 EV CHARGER

The dwellings sampled to have a garage and a car (see table 3.6) are equipped with an EV charger for electric vehicle. Each simulation uses a real user charging profile, which is a record of every home charging session information that occurred over the year. These profiles are sourced from Sørensen et al. [49]. This dataset includes the arrival time, departure time and charged energy of each charging session of electric vehicles from 267 users charging at 12 residential locations in Norway. The dataset also provides assumed arrival and departure state of charge, battery sizes and maximum charging powers for each charging session [47]. All charging profiles containing more than 600 days of data were included in this study, resulting in a total of 16 profiles. Each dwelling with an EV charger is randomly assigned one of these 16 charging profiles.

Table 3.11: Summary statistics of the 16 EV charging profiles used in this study.

Statistic	Mean	Std. dev.
Total energy charged per profile per year	2 414 kWh	1 464 kWh
Number of sessions in a year per profile	193.6	131.2
Session duration across all profiles	12h16min	11h12min
Session SOC charged across all profiles	31.7%	20.0%
Std. dev. of session duration per profile	9h46min	5h32min
Std. dev. of session SOC charged per profile	17.3%	2.6%

Table 3.11 summarizes the variability in the EV charging profiles.

On average, each EV charges by 2 414 kWh per year at home. Assuming an EV efficiency of 0.191 kWh/km [19], this equates to an annual driving distance of

12 638 km. This figure is close to the average yearly travelled distance by Belgian cars of 14 420 km [6], especially considering that not all users charge their EV exclusively at home.

On average, each profile has about 194 charging sessions per year, but this number varies widely between users (standard deviation: 131). The typical charging session lasts 12 hours and 16 minutes, with a standard deviation of 11 hours and 12 minutes, showing that some sessions are much shorter or longer than others. The average state of charge (SOC) charged per session is 31.7%, but again, this varies a lot (standard deviation: 20%). Even when looking at individual profiles, the average standard deviation within a user's sessions remains high at 9 hours and 46 minutes for session duration, and at 17.3% for charged SOC.

In summary, the average charged energy per year across the users is typical for Belgian drivers, and EV charging behaviour is highly variable both between users and from session to session for the same user. This highlights potential challenges in predicting this behaviour for energy management systems.

The model of the charged electric vehicle is defined by eq. (3.3):

$$\mathrm{SOC}_{t+1} = \mathrm{SOC}_t + \frac{\eta^{charge} \times \mathrm{P}_t^{\mathrm{charge}} \times \Delta t}{\mathrm{E}^{\mathrm{batt}}} \tag{3.3}$$

where:

SOC_t [–]	State of charge of the EV battery at time t
$\mathrm{P}_t^{\mathrm{charge}}$ [kW]	Power of the charger at time t
Δt [h]	Time step duration
$\mathrm{E}^{\mathrm{batt}}$ [kWh]	Energy capacity of the battery
η^{charge} [–]	Charging efficiency

The maximum charging power provided in the dataset for the 16 charging profiles is 3.7 kW. We are aware that home chargers in Belgium generally have a higher maximum power (7.4 kW or 11 kW) but we chose to keep the dataset's maximum charging power because the maximum power of the charger could influence the user's charging profile unpredictably (e.g. shorter connection times).

The charging efficiency η^{charge} is set to 88% and the departure SOC is always 95% which are the assumptions made by Sørensen et al. [47] for the dataset.

3.2.5 Photovoltaic panels

The PV panels are simulated using real Belgian PV production data available on a 5-minute granularity from the PVoutput.org website [41]. A total of 380 PV profiles were downloaded from the website, representing all installations with a capacity below 10 kW that recorded measurements for at least 95% of the days

in both 2022 and 2023. Each dwelling with PV panels is randomly assigned a PV profiles from its corresponding region out of the 380 profiles.

3.2.6 Electricity consumption

The electricity consumptions are simulated using real house electricity consumption profiles available on a 2-minute granularity for houses located in Hamelin, Germany [46]. The heating and cooling electricity consumption is measured separately from the rest of the electricity consumption, this study only uses the rest of the electricity consumption.

Out of the 38 profiles available in the dataset, this study used the 22 profiles with data available for at least 98% of the time steps for the years 2019 and 2020. Since all missing data must be imputed for the simulation, this threshold ensures that only a minimal amount of data requires imputation.

The 22 consumption profiles have an average annual electricity consumption of 2 833 kWh, with a standard deviation of 862 kWh. For comparison, the average Flemish household consumed 2 534 kWh according to VREG, while the European average reported by Eurostat is 3 500 kWh [57]. These profiles are therefore reasonably representative of typical household yearly consumption, with the added benefit of being based on real, high-resolution measurements.

3.2.7 Electricity prices

The real-time pricing used in this study is based on a one-year dynamic contract of a major Belgian energy supplier, as if it were taken in January 2023[1]. The total cost is composed of four costs with different pricing mechanisms: the day-ahead cost, the offtake extras cost, the peak cost and the fixed yearly cost.

The day-ahead cost depends on the prices of eqs. (3.4) and (3.5).

$$Vo_t = \begin{cases} (Vd_t + 0.002\frac{€}{kWh}) \times 1.06 & \text{if } (Vd_t + 0.002\frac{€}{kWh}) \geq 0 \\ (Vd_t + 0.002\frac{€}{kWh}), & \text{otherwise} \end{cases} \tag{3.4}$$

$$Vi_t = Vd_t \frac{€}{kWh} \tag{3.5}$$

where:

Vo_t [€/kWh]	Offtake price at time step t
Vi_t [€/kWh]	Injection price at time step t
Vd_t [€/kWh]	Belgian day-ahead hourly price

[1] Full contract available on the paper's GitHub repository at https://github.com/EVERGi/ems_adoption_electrified_belgium_paper

Table 3.12 summarizes the Belgian day-ahead electricity prices (Vd_t) from 2020 to 2024. The table highlights significant year-to-year variation, with 2022 standing out as an exceptional year due to unusually high prices. In contrast, 2023 falls within the typical range observed over recent years.

Table 3.12: Belgian day-ahead electricity prices for 2020–2024. All values are in €/kWh.

Year	Mean	Std. dev.	Min	Max
2020	0.032	0.017	-0.115	0.200
2021	0.104	0.079	-0.070	0.620
2022	0.245	0.135	-0.100	0.871
2023	0.097	0.046	-0.120	0.330
2024	0.070	0.0430	-0.140	0.566

A value-added tax of 6% is added to the offtake price when it is positive. Extra costs that occur for energy taken from the grid for transport, distribution and taxes are calculated together and labelled as offtake extras cost. The offtake price for all these additional offtake costs Vf is 0.079 €/kWh.

The monthly peak offtake cost Cp_m is calculated by eq. (3.6).

$$Cp_m = Vp \times \max_{t \in B_m}\left(\frac{Eo_t}{\Delta t}, 2.5kW\right) \tag{3.6}$$

where:

Vp [€/kW]	Peak price (used: 3.66 €/kW)
Eo_t [kWh]	Energy taken from the grid at time step t
B_m	Set of time steps in month m
Δt [h]	Length of time step

This cost adds a penalty to the highest offtake power measured over any 15 minute time step within each month. The offtake energy is divided by the length of a time step to obtain the power in kW. A minimum monthly peak power of 2.5 kW is enforced in the cost, even if the measured peak power is lower.

The monthly peak offtake cost, otherwise called capacity tariff, is a pricing mechanism applied to all residential consumers equipped with a digital meter in the Belgian region of Flanders. The goal is to incentivize consumers to lower their peak offtake power from the grid.

Finally, part of the electric bill is a fixed yearly cost Cy of €115.23 per year.

The total electricity cost $Ctot$ is calculated by eq. (3.7).

$$Ctot = \underbrace{\sum_{t \in T} (Eo_t Vo_t - Ei_t Vi_t)}_{\text{day-ahead cost}} + \underbrace{\sum_{t \in T} Eo_t Vf}_{\text{offtake extras cost}} + \underbrace{\sum_{m \in M} Cp_m}_{\text{peak cost}} + \underbrace{Cy}_{\text{yearly cost}} \tag{3.7}$$

where:

T Set of all time steps
M Set of all months

3.3 Asset management

The three controllable assets, when present in the dwelling, are the air-to-air heat pump, the water heater and the electric vehicle charger.

Every simulation is run for a period of one year with a time step of 15 minutes as this is the time period used to calculate the peak electricity cost in Belgium (see eq. (3.6)). The EMS can change at each time step the cooling and heating temperature set points of the heat pump, the electric power of the water heater and the charge power of the electric vehicle charger.

3.3.1 Heat pump control

The heat pump operates within a default comfort range. If the indoor temperature is outside the comfort range, the heat pump will bring the indoor temperature back to the comfort range. The EMS can set a heating or cooling temperature set point within this comfort range.

Based on different studies, Peeters et al. [38] established usable comfort ranges for the different rooms of a residential building. The comfort range depends on the reference external temperature $T_d^{e,ref}$ which is based on the minimum and maximum temperatures of the current and the last three days as shown in eq. (3.8).

$$T_d^{e,ref} = \frac{T_d^{avg} + 0.8 \times T_{d-1}^{avg} + 0.4 \times T_{d-2}^{avg} + 0.2 \times T_{d-3}^{avg}}{2.4} \tag{3.8}$$

where:

T_d^{avg} Average of the maximum and minimum temperature on day d

To avoid relying on weather forecasts for T_d^{avg}, this study uses T_{d-1}^{avg} as a substitute in the calculation. The mean absolute difference in the resulting comfort range is only 0.18 °C (standard deviation 0.24 °C), showing that this simplification has minimal impact while reducing forecast dependency.

For each room a neutral temperature T_d^n is calculated based on the reference external temperature $T_d^{e,ref}$. The neutral temperature represents the temperature at which most people would feel comfortable in the room. From T_d^n a comfort range is calculated using the width of the comfort band w and the band asymmetry parameter α. This study uses a width w of 5 °C and an asymmetry parameter α of 0.7 which corresponds to a 10% predicted percentage of dissatisfied (PPD) users [38].

For the bedroom T_d^n is calculated as shown in eq. (3.9).

$$T_d^n = \begin{cases} 16°C & \text{if } T_d^{e,ref} < 0°C \\ 0.23 \times T_d^{e,ref} + 16°C & \text{if } 0°C \leq T_d^{e,ref} < 12.6°C \\ 0.77 \times T_d^{e,ref} + 9.18°C & \text{if } 12.6°C \leq T_d^{e,ref} < 21.8°C \\ 26°C & \text{if } T_d^{e,ref} \geq 21.8°C \end{cases} \tag{3.9}$$

The comfort range for the bedroom is calculated as shown in eq. (3.10).

$$\begin{aligned} T_t^{low} &= \max(16°C, T_d^n - w \times (1-\alpha)) \\ T_t^{up} &= \min(26°C, T_d^n + w \times \alpha) \end{aligned} \tag{3.10}$$

where:

T_t^{low}	Lower bound of the comfort range at time step t
T_t^{up}	Upper bound of the comfort range at time step t

For the rooms of the house other than the bedroom and the bathroom, T_d^n is calculated as shown in eq. (3.11).

$$T_d^n = \begin{cases} 20.4°C + 0.06 \times T_d^{e,ref} & \text{if } T_d^{e,ref} < 12.5°C \\ 16.63°C + 0.36 \times T_d^{e,ref} & \text{if } T_d^{e,ref} \geq 12.5°C \end{cases} \tag{3.11}$$

The comfort range for the rooms of the house other than the bedroom and the bathroom are calculated as shown in eq. (3.12).

$$\begin{aligned} T_t^{low} &= \max(18°C, T_d^n - w \times (1-\alpha)) \\ T_t^{up} &= T_d^n + w \times \alpha \end{aligned} \tag{3.12}$$

Figure 3.4 illustrates how the comfort temperature bounds T_t^{low} and T_t^{up} vary with the reference external temperature $T_d^{e,ref}$ for both sleeping and home occupancy

profiles. For an example of how the comfort range evolve throughout a weekend, check fig. 3.6.

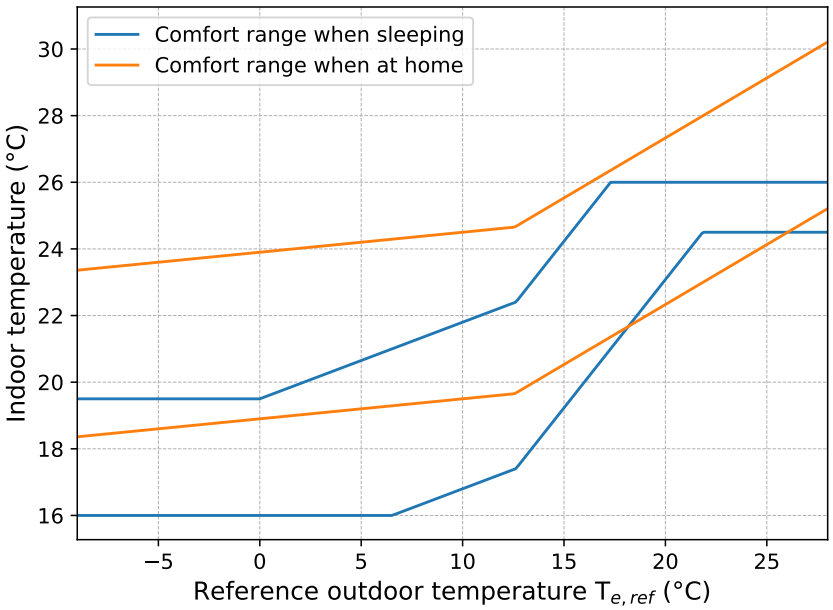

Figure 3.4: Illustration of the comfort temperature ranges for sleeping and home occupancy profiles as a function of the reference external temperature $T_d^{e,ref}$. The plot covers $T_d^{e,ref}$ values from −9 °C to 28 °C, representing the coldest and warmest reference temperatures observed in Belgium during the 21st century.

Aerts et al. [1] defines a set of occupancy profiles out of 3 474 Belgian households. The paper finds 7 daily occupancy profiles representative of people's behaviour which cover 79% of the evaluated households. The occurrence of each profile depends on whether the day is a weekday, Saturday, or Sunday. In these profiles, the household is either at home, absent or sleeping.

Each dwelling is assigned one of the 7 representative occupancy profiles for the weekdays, Saturdays, and Sundays based on the distribution provided by Aerts et al. [1]. The occupancy profiles are used to determine the comfort range of the dwelling at each time of the day. When at home, the comfort range is defined by eq. (3.12), when sleeping by eq. (3.10) and when absent the range is set to $T_t^{low} = 5°C$ and $T_t^{up} = 40°C$.

3.3.2 Water heater control

The hot water tank has a temperature range from 55 °C to 65 °C that it should maintain. These values are the minimum and maximum domestic hot water temperatures from the European, Australian, Canada, Japan and Unites States standards [21]. If the water temperature goes below 55 °C, the water heater will heat the water at maximum power for one 15 minute time step. Whenever

the water temperature is above 65 °C, the water heater will not heat the water. Otherwise, the water can be heated between 0 kW and the maximum power of the water heater.

3.3.3 Electric vehicle charger control

The electric vehicle charger is controlled by setting a power between 0 kW and 3.7 kW. Once the required departure state of charge is reached, charging stops. Currently, most EVs do not provide information about their battery's state of charge [47]. Therefore, in this work, the EMSs are assumed to operate without knowledge of the battery's state of charge.

For dwellings with an EV charger, two types of simulations are run. In the first, the EMS fully controls EV charging, allowing to assess how well each EMS charges the EV (see figs. 3.7d and 3.8). In the second, charging is enforced so the EV always reaches full charge before departure, ensuring a fair cost comparison across EMSs (see figs. 3.7a, 3.7b, 3.9 and 3.10).

In the enforced scenario, if the EV is not fully charged near departure, the charger switches to 3.7 kW to guarantee a full charge, simulating a driver's request for a complete charge before leaving.

3.4 Energy management systems

3.4.1 Rule-based control

The rule-based control EMS uses simple, predefined strategies for each controllable asset, reflecting typical default operation modes of these assets:

- **Heat pump:** The heat pump setpoint is determined by obtaining the comfort range for the next 1 hour and 30 minutes. The controller gets the highest lower bound (T_t^{low}) and the lowest upper bound (T_t^{up}) during this period and sets the heating and cooling setpoints to these values.
- **Domestic water heater:** The water heater operates at maximum power for 15 minutes whenever the tank temperature drops below 55 °C.
- **Electric vehicle charger:** The EV is always charged at the maximum power.

The 1 hour and 30 minute time horizon was selected based on the final results to ensure that the rule-based controller achieves a thermal discomfort level similar to the other EMSs (see fig. 3.7c).

3.4.2 Model predictive control

Model predictive control (MPC) is a control strategy that optimizes the operation of controllable assets over a prediction horizon by solving an optimization problem at each time step. The MPC uses a mathematical model of the assets and predicts all necessary parameters such as weather, electricity prices, and energy consumption of non-controllable assets to define the set points for the controllable assets.

This study implements two MPC approaches: a realistic MPC (MPC) and a perfect forecast MPC (MPC P). The realistic MPC is a state-of-the-art control algorithm for EMSs. The perfect forecast MPC serves as an unrealistically performant EMS to indicate an upper EMS performance limit.

The realistic MPC uses realistic forecasts for most parameters and uses a prediction horizon that extends up to when day-ahead prices (Vd_t) are available, that is until midnight of the current day if before 14:00 Central European Time or until midnight of the following day if after 14:00, resulting in a horizon ranging from 10 hours and 15 minutes to 34 hours. MPC P uses a perfect forecast for all parameters and has a fixed prediction horizon of 48 hours.

Both MPC approaches are implemented using the Gurobi solver [25].

3.4.2.1 Mathematical model

The MPC model is formulated as a mixed-integer linear programming problem.

The dwelling's thermal model is approximated using the resistor-capacitor analogy for thermal modelling [35, 15]. Equation (3.13) shows the implementation of this model in the MPC. EnergyPlus does not allow to control the power of the heat pump directly so the control setpoint used is the temperature of the model at the next time step (T^{in}_{t+1}).

$$C^{dwell}\frac{T^{in}_{t+1}-T^{in}_{t}}{\Delta t}=\frac{T^{out}_{t}-T^{in}_{t}}{TR}+Q^{pump}_{t}+gAQ^{sun}_{t} \tag{3.13}$$

where:

C^{dwell} [J/K]	Thermal capacitance of the dwelling
T^{in}_t [K]	Indoor temperature at time step t
T^{out}_t [K]	Outdoor temperature at time step t
TR [K/W]	Thermal resistance of the dwelling
Q^{pump}_t [W]	Thermal power of the heat pump at time step t
gA [m^2]	Solar aperture
Q^{sun}_t [W/m]	Solar irradiance at time step t

The primary heating and cooling coils and the backup heating coil are modelled linearly using a fixed efficiency conversion between the electric power the thermal power as shown in eq. (3.14).

$$Q_t^{pump} = \eta^{heat} P_t^{heat} - \eta^{cool} P_t^{cool} + \eta^{backup} P_t^{backup} \tag{3.14}$$

where:

η^{coil} [–]	Efficiency of the heat, cool or backup coil
P_t^{coil} [W]	Electric power of the heat, cool or backup coil at time step t

The total electric power of the heat pump P_t^{pump} is modelled as the sum of the electric power of the primary heating and cooling coils and the backup heating coil as shown in eq. (3.15).

$$P_t^{pump} = P_t^{heat} + P_t^{cool} + P_t^{backup} \tag{3.15}$$

The parameters for the thermal capacitance (C^{dwell}), thermal resistance (TR), solar aperture (gA), and coil efficiencies (η^{heat}, η^{cool}, η^{backup}) are calibrated by fitting the MPC model to the detailed simulation results for each dwelling over the year 2022.

When using the accurate simulation values for T_t^{out}, T_t^{in}, Q_t^{pump} and Q_t^{sun}, the calibrated model from eq. (3.13) predicts the indoor temperature for the next time step T_{t+1}^{in} with a mean absolute error of 0.12 °C (standard deviation 0.14 °C) across all dwellings. This outperforms a naive baseline that assumes the indoor temperature remains constant, which yields a mean absolute error of 0.16 °C (standard deviation 0.31 °C). These results indicate that the simulation's thermal dynamics are more complex than those captured by the simplified RC model used in the MPC.

The maximum electric power of each heat pump coil is constrained by the maximum capacities defined by the Openstudio-HPXML workflow for the respective dwelling. The MPC model does not include the ventilation consumption and the defrosting coil consumption present in the simulation model.

The indoor temperature T_t^{in} is constrained to be within the comfort range defined in subsection 3.3.1 for the predicted horizon.

The water heater's MPC model closely mirrors the simulation model by implementing eq. (3.2) and enforcing heating at maximum power for 15 minutes whenever the tank temperature falls below 55 °C. The main difference is that, in the MPC model, the tank temperature is strictly limited to a maximum of 65 °C,

rather than allowing it to exceed this value and simply disabling heating in the following time step.

The electric vehicle charger is modelled using eq. (3.3) for the electric vehicle charging and the charging session is constrained to charge a given amount of energy in a given time period. This exactly matches the simulation model of the charger.

The total electric power of the dwelling is the sum of the different assets as shown in eq. (3.16).

$$P_t^{total} = P_t^{pump} + P_t^{water} + P_t^{EV} - P_t^{PV} + P_t^{load} \tag{3.16}$$

Overall, the MPC model provides a close approximation of the simulation model.

3.4.2.2 Objective function

To calculate the objective function, the energy offtake Eo_t and injection Ei_t are calculated using the total electric power P_t^{total} following eqs. (3.17) and (3.18).

$$Eo_t = \begin{cases} P_t^{total}\Delta t, & \text{if } P_t^{total} \geq 0 \\ 0, & \text{otherwise} \end{cases} \tag{3.17}$$

$$Ei_t = \begin{cases} -P_t^{total}\Delta t, & \text{if } P_t^{total} < 0 \\ 0, & \text{otherwise} \end{cases} \tag{3.18}$$

Model predictive control uses an adapted version from eq. (3.7) as its objective function without the yearly cost, which cannot be optimized, and where the peak cost increase is directly reflected on the prediction horizon (see eq. (3.19)).

$$O^M PC = \underbrace{\sum_{t\in H}(Eo_t Vo_t - Ei_t Vi_t)}_{\text{day-ahead cost}} + \underbrace{\sum_{t\in H} Eo_t Vf}_{\text{offtake extras cost}} + \underbrace{Vp \times \max_{t\in H}(\frac{Eo_t}{\Delta t}, Pp)}_{\text{peak cost}} \tag{3.19}$$

where:

H	Set of time steps in the prediction horizon
Pp [kW]	Previous grid peak power for that month (minimum 2.5 kW)

3.4.2.3 Forecast

The forecasted variables used by the MPC's model are:

- The outdoor temperature (T_t^{out})
- The solar irradiance (Q_t^{sun})
- The hot water flow (F_t)
- The arrival time of the EV
- The departure time of the EV
- The charged energy of the EV
- The PV production (P_t^{PV})
- The base load consumption (P_t^{load})

MPC P uses the actual future values of all relevant variables over its 48-hour prediction horizon.

For the realistic MPC, all non EV related variables in the next hour are predicted using the most recent value, while subsequent time steps are predicted using values from the preceding 24 hours at the same time of the day.

Predicting the arrival time of an EV is particularly challenging due to the high variability in user behaviour. Therefore, the MPC only assumes that an EV arrives when a new charging session starts. For each new charging session, the departure time and charged energy are predicted using k-Nearest Neighbor (k-NN) [10] regression as it showed good performance in Genov et al. [24] to predict departure time and charged energy of EVs. The forecasting features used are the arrival time and week day, the charged energy from the previous charging session, the total connection time of the previous charging session, and the time since the start of the previous charging session. The model is trained on the 2022 charging profile data of the dwelling the MPC is controlling and the k-NN model is retrained with one additional charging session data every time a new charging session is finished during simulation.

To account for prediction errors, if the MPC incorrectly predicts that the EV is fully charged but it remains connected and not fully charged, the MPC will schedule charging at 3.7 kW for the next 2 hours to ensure additional charging. Conversely, if the MPC overestimates the departure time and the enforced charging behaviour was triggered at the previous time step (see subsection 3.3.3), then the MPC assumes the EV will be charged at maximum power for the next 2 hours.

3.4.3 TreeC

3.4.3.1 Training process

The training of TreeC is performed by evaluating the performance of different decision tree based EMSs on the simulation of the year 2022. The enforced charging mechanism described in subsection 3.3.3 is not used during the TreeC training simulations.

Two changes have been made to the original TreeC method described in [45]:

- The covariance matrix adaptation evolution strategy was replaced by the metaheuristic algorithm particle swarm optimization [31] as it performed better in [44]. This study used the particle swarm optimization algorithm implemented in the Python library pygmo [4] with default parameters which corresponds to the canonical particle swarm optimization algorithm described in [40].
- After training, the best decision tree is pruned by iteratively removing leaves whose removal does not degrade the objective function score by more than 5 € compared to the unpruned tree. This threshold is selected to balance performance and interpretability: it is low enough to avoid sacrificing meaningful cost savings, yet high enough to eliminate leaves that contribute minimally to overall performance.

  The least used leaves are tested for removal first.

The TreeC EMS is trained in simulation with the particle swarm optimization algorithm over 300 generations and a population of 200 individuals. This means 60 000 decision tree based EMSs are evaluated in simulation on the 2022 training data and the best one is then selected. The population size of 200 was chosen as it obtained better results for the real-world and non-unimodal benchmarks of the IEEE Competitions in Evolutionary Computation compared to the more historically implemented population sizes of 20-50 individuals [39]. To reduce the risk of converging to a local optimum, five independent EMSs are trained and pruned for each dwelling. The EMS with the best performance among the five is then selected for use in the experiment.

The TreeC training was performed on a high performance computer composed of Intel® Xeon® Gold, AMD EPYC™ 9384X and AMD EPYC™ 9535 processors. Each training used 100 cores and evaluated 300 000 EMSs (60 000 metaheuristic optimization × 5 to avoid local optimums). On average a training took 5 days and 20 hours to complete with each EMS evaluation taking on average 34 seconds on a single core corresponding to the simulation of the dwelling for one year.

3.4.3.2 Decision tree model

The training can select from up to 23 different inputs to use as splitting features in the decision trees. The inputs related to specific assets are not used in the dwellings where the asset is not present.

The inputs are:

- The time of the day
- The day of the week
- The cosine of the day of the year, normalized so that -1 corresponds to December 21st (winter solstice) and $+1$ to June 21st (summer solstice)
- P_{t-1}^{asset}, the average power of each asset (heat pump, water heater, the EV charger, base load, PV) and total power for the previous time step (6 inputs)
- T_t^{in}, the indoor temperature
- T_t^{low}, the lower bound of the comfort range for the next time step, in 30 minutes and in 2 hours (3 inputs)
- T_t^{tank}, the water tank temperature
- T_t^{out}, the outdoor temperature
- Q_t^{sun}, the solar irradiance
- Vd_t, the day-ahead price
- $\max_{h \in H}(Vd_h) - Vd_t$ the difference between the maximum and the current price for the next 6, 12 and 24 hours with H being the set of future time steps (3 inputs)
- $Vd_t - \min_{h \in H}(Vd_h)$, the difference between the current price and the minimum price for the next 6, 12 and 24 hours with H being the set of future time steps (3 inputs)

The TreeC EMS can control the heat pump through three types of setpoints: comfort rule-based, temperature setpoint and adjustment to indoor temperature. The comfort rule-based setpoint in TreeC follows the same logic as in the rule-based controller: it sets the heat pump setpoints for cooling and heating to match the comfort range over a specified time window (see subsection 3.4.1). Unlike the rule-based controller, however, TreeC select the length of this time window anywhere from 0 to 12 hours. The temperature setpoint sets the heat pump setpoints for cooling and heating to a fixed temperature between 5, °C and 40 °C. The adjustment to indoor temperature setpoint sets the heat pump's

heating and cooling setpoints relative to the current indoor temperature T_t^{in}. For heating, the setpoint is defined as T_t^{in} plus a fixed offset; for cooling, as T_t^{in} minus a fixed offset. This fixed offset can be selected within the range of $-2\,°C$ to $5\,°C$.

The TreeC EMS control the water heater and EV charger by setting directly the power of the asset.

3.4.3.3 Objective function

The TreeC EMS's goal is to minimize the total electricity cost, keep the indoor temperature within the comfort range, and charge the EV sufficiently before departure. The total electricity cost $Ctot$ is calculated using eq. (3.7).

Discomfort is measured in Kh (Kelvin-hours), which represents how long and by how much the indoor temperature deviates from the comfort range. For example, being 2 °C above the comfort range for 1 hour results in 2 Kh of discomfort.

The discomfort above the upper bound of the comfort range δ_t^{up} and below the lower bound δ_t^{low} are calculated using eqs. (3.20) and (3.21).

$$\delta_t^{up} = \int_0^{\Delta t} \max\left(\frac{T_{t+1} - T_t}{\Delta t}x + (T_t - T_t^{up}),\, 0\right) dx \tag{3.20}$$

$$\delta_t^{low} = \int_0^{\Delta t} \max\left(\frac{T_t - T_{t+1}}{\Delta t}x + (T_t^{low} - T_t),\, 0K\right) dx \tag{3.21}$$

where:

T_t [K]	Indoor temperature at time step t
Δt [h]	Length of the time step
T_t^{up} [K]	Upper bound of the comfort range at time step t
T_t^{low} [K]	Lower bound of the comfort range at time step t

The total discomfort cost $C\delta$ is calculated by eq. (3.22).

$$C\delta = max(0€, w1 \times (\sum_{t \in T}(\delta_t^{up} + \delta_t^{low}) - 30Kh)) \tag{3.22}$$

where:

$C\delta$ [€]	Discomfort cost
w_1 [€/Kh]	Penalty weight for discomfort cost (5 €/Kh)

A discomfort penalty is only applied if the total annual discomfort exceeds 30 Kh, which corresponds to a low level of discomfort over the year. For any discomfort

above this threshold, a penalty of 5 €/Kh is imposed, strongly incentivizing the EMS to maintain thermal comfort.

The cost for missed EV charging energy C_{charg} is calculated using eq. (3.23).

$$Ccharg = \sum_{s \in S} \begin{cases} w_2 \times E_s^{miss}, & \text{if } E_s^{miss} \leq E^{5\%} \\ w_2 \times E^{5\%} + w_3 \times (E_s^{miss} - E^{5\%}), & \text{otherwise} \end{cases} \tag{3.23}$$

where:

C_{charg} [€]	Penalty cost for missed EV charging energy
S	Set of all charging sessions
E_s^{miss} [kWh]	Missed energy in session s
$E^{5\%}$ [kWh]	Energy corresponding to charging the EV by 5% SOC
w_2 [€/kWh]	Penalty weight for missed energy up to $E^{5\%}$ (0.5 €/kWh)
w_3 [€/kWh]	Penalty weight for missed energy exceeding $E^{5\%}$ (2 €/kWh)

A missed energy event occurs when the simulated EMS charges less energy than what was originally delivered in the dataset for a given charging session. In other words, it is the shortfall between the original and the actually delivered energy by the EMS.

Missing up to 5% of the EV's SOC in a charging session is undesirable but not critical, so a moderate penalty of 0.5 €/kWh is applied reflecting an exceptionally high price of electricity (see table 3.12). However, missing more than 5% SOC is considered a significant problem, and a much higher penalty of 2 €/kWh is imposed to strongly discourage such outcomes.

The total objective function O of the TreeC EMS is calculated by summing the different costs from eqs. (3.7), (3.22) and (3.23) as shown in eq. (3.24).

$$O^{TreeC} = Ctot + C\delta + Ccharg \tag{3.24}$$

3.4.3.4 Interpretable decision tree

As an example, the TreeC EMS obtained after training for dwelling 1 is presented in fig. 3.5 as it is the first dwelling with all possible assets present (PV, heat pump, water heater and EV charger).

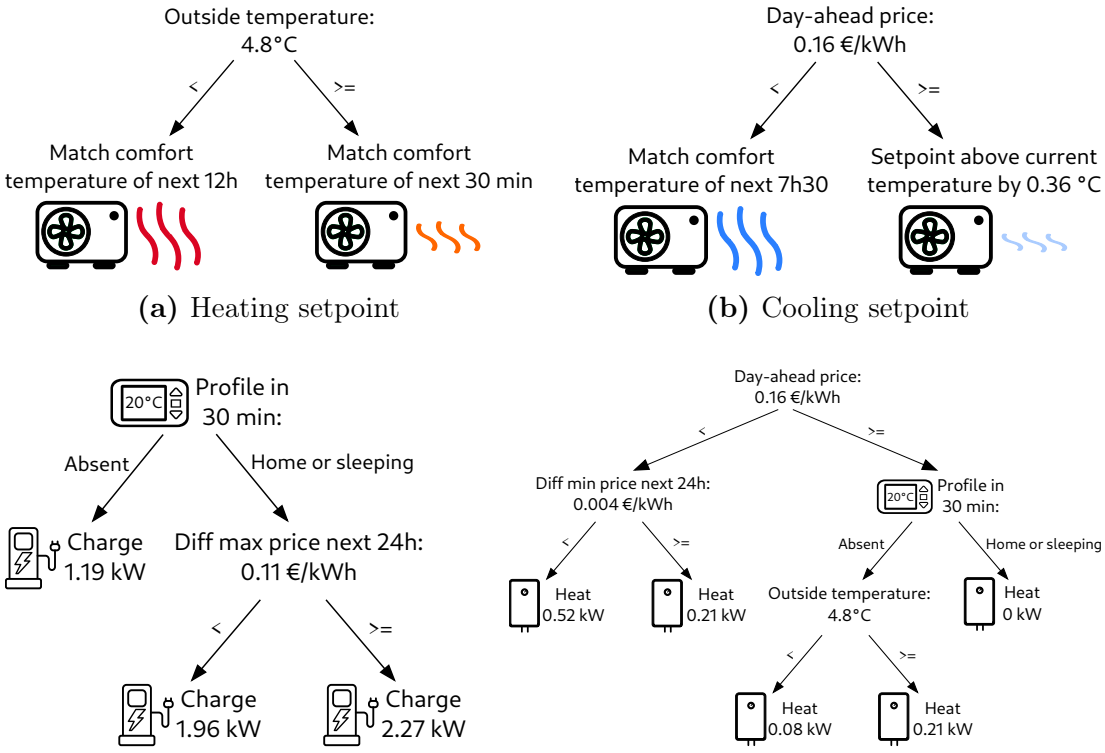

(a) Heating setpoint

(b) Cooling setpoint

(c) EV charger power setpoint

(d) Water heater power setpoint

Figure 3.5: Decision trees for each controllable asset in the TreeC EMS for dwelling 1.

The decision tree is interpretable for each of the set points it controls. For the heating setpoint in fig. 3.5a, when the outdoor temperature is cold (below 4.8 °C), the EMS is precautious and matches the comfort range for the next 12 hours to not incur discomfort, otherwise it matches it for the next 30 minutes as this is more energy efficient. For the cooling setpoint in fig. 3.5b, when the day-ahead price is considered low enough (below 0.16 €/kWh), the EMS matches the comfort range 7h30 minutes ahead, otherwise it just avoids that the indoor temperature increases too fast. For the EV charger in fig. 3.5c, the EMS charges between 1.96 kW and 2.27 kW when the occupancy profile is marked as home or sleeping in 30 minutes and at a lower power of 1.19 kW then occupancy profile is absent. The EV charger never charges at full power to avoid peaks. For the water heater in fig. 3.5d, the EMS heats the water more when the day-ahead price is considered low enough and doesn't heat at all when the occupancy profile is marked as home or sleeping in 30 minutes.

The figures and explanation highlight the interpretability of the TreeC EMS and also illustrate how it can coordinate the operation of different assets. For example, the EV charger tends to charge more when the water heater is not heating, helping to avoid simultaneous power peaks. This is achieved by both controllers using the same occupancy profile feature.

3.5 Results

Figure 3.6 shows how the different EMSs operate during the last weekend of winter for dwelling 1. This dwelling is a detached house in Flanders with all assets present: heat pump, water heater, EV charger, and PV system. The water heater has a 171 litre tank, the house is partially insulated (roof only), and the PV system has a capacity of 5 kWp. This dwelling was chosen because it is the first with all assets present, and the selected weekend highlights an unsuccessful EV charging session under the MPC EMS. In this example, MPC incorrectly forecasted that the EV would remain connected until 12:30 on the 19$^{\text{th}}$ of March, whereas in reality, the EV departed at 18:00 on the 18$^{\text{th}}$ of March. Because the day-ahead prices, and therefore injection prices, before 18:00 on the 18$^{\text{th}}$ of March were higher than those before 12:30 on the 19$^{\text{th}}$ of March, MPC planned to inject excess PV generation on the 18$^{\text{th}}$ March and use the excess PV on 19 March to charge the EV. Mathematically this all makes sense, but since the prediction was incorrect, the EV was not fully charged by the time of departure contrary to the other EMSs. RBC charges to the maximum power, TreeC charges at a lower power to avoid peaks but still high enough to ensure the EV is fully charged and MPC P charges the EV using all excess PV production and cheap day ahead prices.

For the domestic hot water, TreeC heats the water tank over almost the entire weekend because the price of electricity is mostly below 0.16 €/kWh (see fig. 3.5d), MPC and MPC P make more use of the excess PV production to heat the water tank, and RBC keeps the water tank around the lower 55 °C bound.

For heating, all EMSs generally maintain the indoor temperature near the lower bound of the comfort range except for moments where the sun and outdoor temperatures heat the house naturally. MPC and MPC P take advantage of excess PV production to provide additional heating at certain times. RBC starts heating earlier in the morning than the other EMSs, which results in slightly lower energy efficiency.

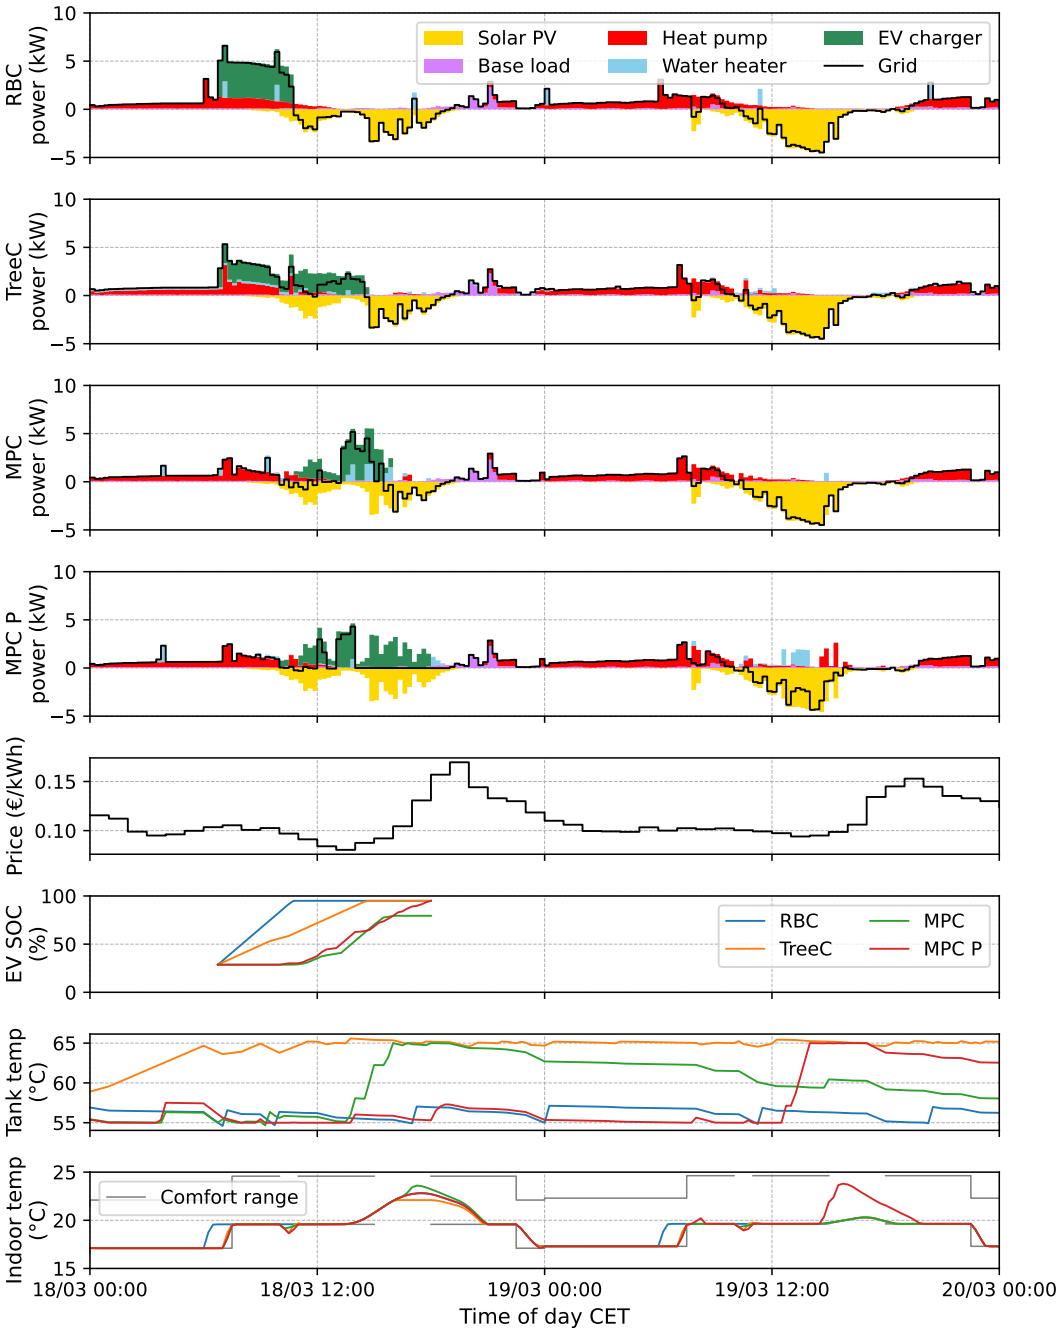

Figure 3.6: Illustration of the operation of the different EMSs over the weekend of 18-19 March 2023 for dwelling 1. The top four plots show the stacked power consumption and production of each asset, with the total grid connection power indicated in black. The four lower plots show the day-ahead electricity prices, the state of charge of the connected EV, the water heater's tank temperature, and the indoor temperature for each EMS.

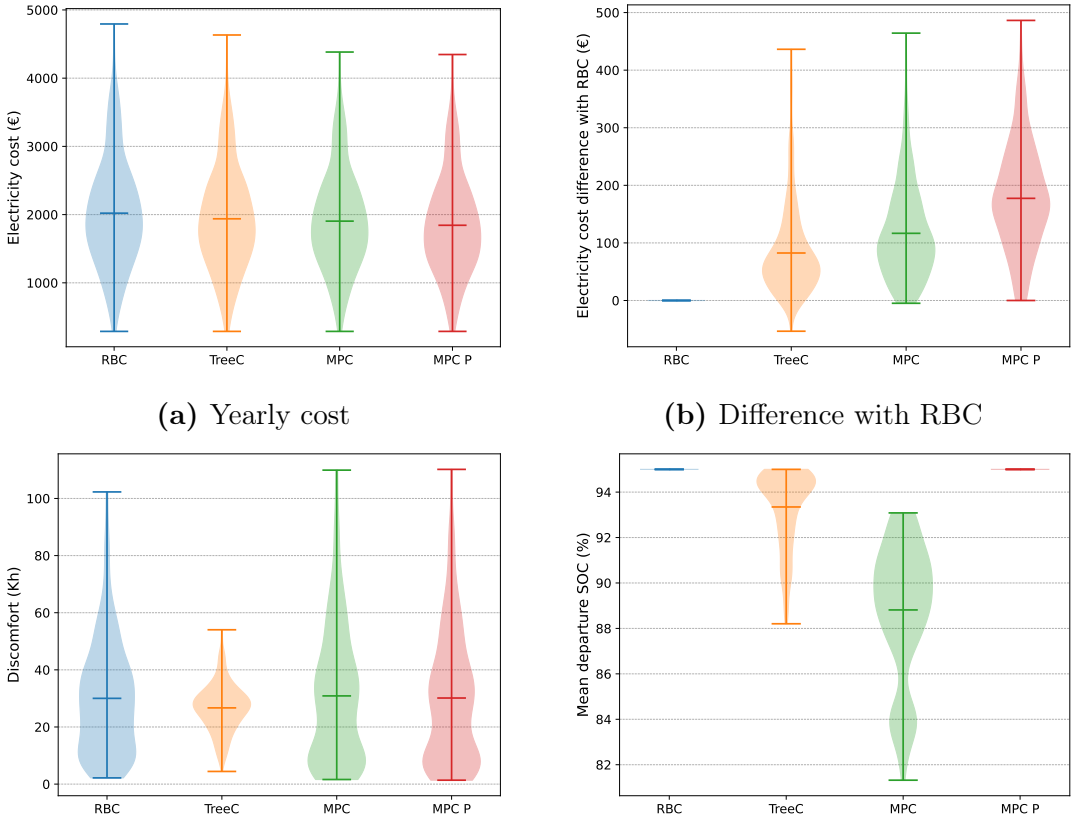

(a) Yearly cost

(b) Difference with RBC

(c) Thermal discomfort

(d) EV charge completion

Figure 3.7: Comparison of key performance metrics for the different EMSs: (a) total yearly cost, (b) difference in electricity cost with the RBC EMS, (c) thermal discomfort, and (d) mean departure state of charge for EV charging sessions. For each EMS, violin plots (a) and (b) include 500 dwellings, (c) includes the 447 dwellings with a heat pump, and (d) includes the 344 dwellings with an EV charger. Each violin plot shows the distribution of results across all dwellings. The width of the violin at each value reflects the relative frequency of dwellings with that result. The horizontal dashes within each violin indicate the minimum, mean, and maximum values for each EMS.

Results in fig. 3.7 show different metrics to compare EMSs across the 500 dwellings. Figure 3.7a shows the total yearly cost for each dwelling. The results indicate a hierarchy in performance: the RBC EMS has the highest costs, followed by TreeC, then MPC, with MPC P achieving the lowest costs overall. There is considerable variation in total costs across dwellings, reflecting the impact of factors such as insulation quality, and the presence of EV chargers or PV systems on overall energy consumption.

To provide a clearer comparison of the EMS performance relative to RBC, fig. 3.7b presents the reduction in total yearly cost for each EMS compared to the RBC EMS. The results show that TreeC achieves a mean cost reduction of 82.5 €, MPC achieves 116.6 €, and the upper bound MPC P achieves 177.4 €. TreeC

and MPC rarely achieve a worse cost than RBC, which is likely counteracted by an increase in thermal comfort.

Figure 3.7c shows the thermal discomfort for each EMS with all EMSs achieving similar mean discomfort levels between 26 and 31 Kh. TreeC manages to keep all dwellings below 55 Kh discomfort while the other EMSs have some dwellings with discomforts up to 111 Kh. Overall even 111 Kh can be considered acceptable especially since some dwellings are not insulated at all and this only represents 111 hours in the year when the indoor temperature is 1 °C above or below the comfort range with the rest of the year being within the comfort range.

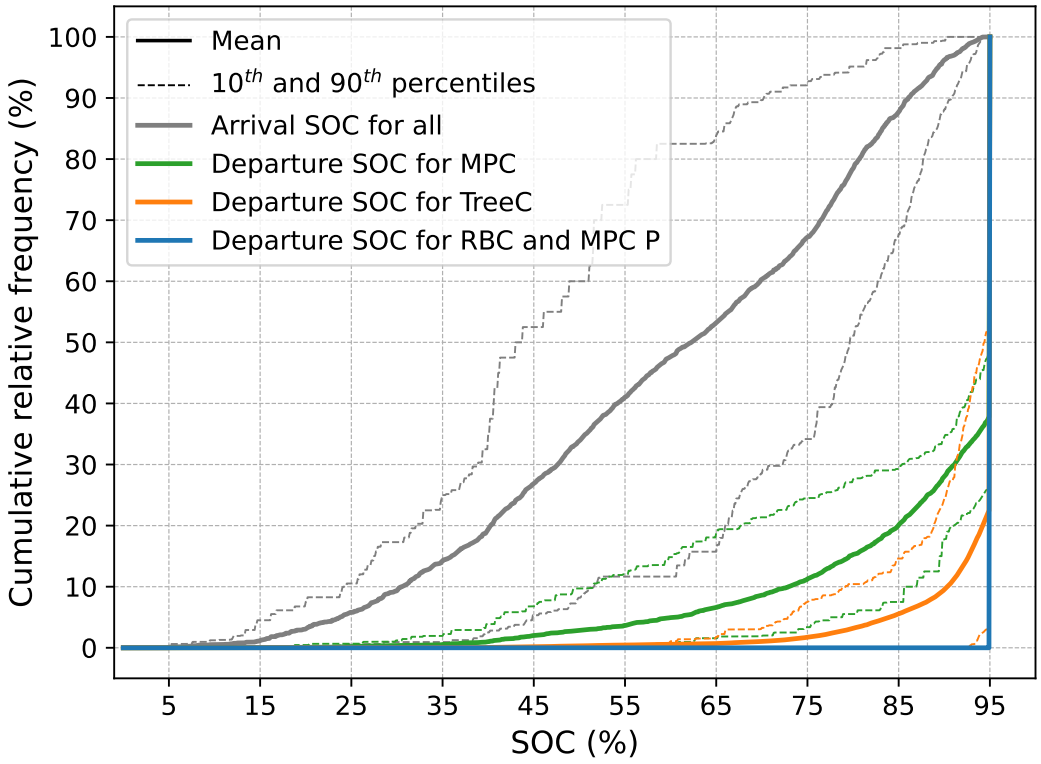

Figure 3.8: Cumulative distribution of EV charging session SOC at arrival (grey) and departure (coloured) for each EMS. Thick lines show the mean cumulative relative frequency of SOC across all dwellings with an EV charger, representing the percentage of sessions at or below the given SOC. Thin dashed lines indicate the 80% interpercentile range (10^{th} to 90^{th} percentiles), highlighting variability between dwellings. This figure demonstrates how well each EMS ensures EVs are sufficiently charged before departure.

Figure 3.7d shows the mean departure SOC for the EV charging sessions across all dwellings with an EV charger. The original dataset assumed the departure charging SOC to be 95% for all sessions so it is normal that RBC and MPC P always achieve this value. For this particular metric TreeC performs better than MPC achieving a mean of 93.3% departure SOC compared to 88.8% for MPC.

This shows that TreeC is able to charge the EV more effectively than MPC but the average departure SOC does not always give a complete representation of the charging performance.

To give a better representation, fig. 3.8 presents the cumulative distribution of EV SOC at departure for each EMS. The results show that, with MPC, the SOC falls on average below 85% in 20% of charging sessions and below 65% in 6.6% of sessions. In contrast, TreeC only drops below these thresholds in 5.5% and 0.7% of sessions, respectively. We think that TreeC generally maintains an acceptable charging level for most users while the MPC EMS may leave a significant share of drivers dissatisfied. Developing more refined metrics for driver charging satisfaction would be valuable for future studies to better quantify this aspect.

Figure 3.9 shows the impact of different dwelling features on the EMS performance. Most results are as expected: fig. 3.9a shows that an EMS makes better improvements when the dwelling has a PV system, fig. 3.9b shows the same for EV chargers, figs. 3.9c and 3.9d shows the same for central heating and fig. 3.9e shows that generally the more assets a dwelling has, the better the EMS performs.

Interestingly fig. 3.9c shows that more insulation has negative impact on the EMSs improvement compared to RBC. This is certainly because dwellings with more insulation have a lower energy consumption and therefore the EMSs have less opportunity to reduce costs. For the water tank size, fig. 3.9d shows a positive correlation between the tank size and the EMS performance.

To understand better how the EMSs achieve cost reductions, fig. 3.10 breaks down the savings by cost component.

Only minimal savings are achieved in the offtake extras costs related to energy efficiency. This is expected, as the RBC is already quite energy efficient for water and space heating. Efficiency gains can be realized by directly consuming electricity from PV, but 74.2% of dwellings do not have a PV system.

MPC and MPC P achieve significant cost saving in the day-ahead costs, this is the main cost component that gives these methods their edge over TreeC. This can be explained by MPC based techniques having perfect hourly knowledge of the future day-ahead prices and being able to schedule assets based on these prices more effectively than TreeC which relies on incomplete knowledge of the future prices and a less sophisticated decision tree control model.

For the peak cost component, TreeC slightly outperforms MPC. During training, TreeC learns to minimize peak costs by coordinating asset operation and selecting lower power setpoints, thereby avoiding large simultaneous peaks. While MPC also attempts to coordinate assets to reduce peaks, prediction errors can occur due to forecast or model inaccuracies and result in unexpected peaks, which may

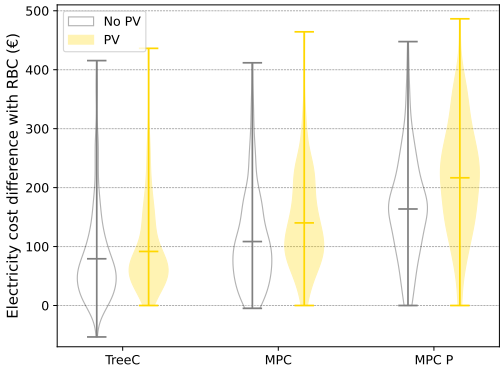

(a) PV

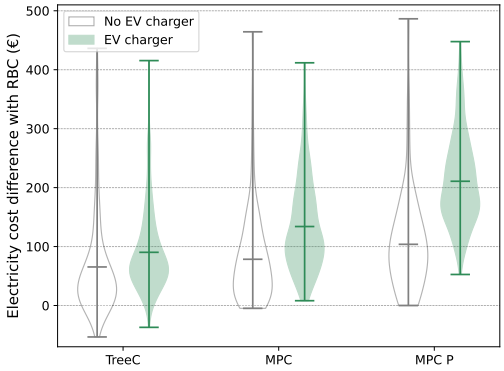

(b) EV charger

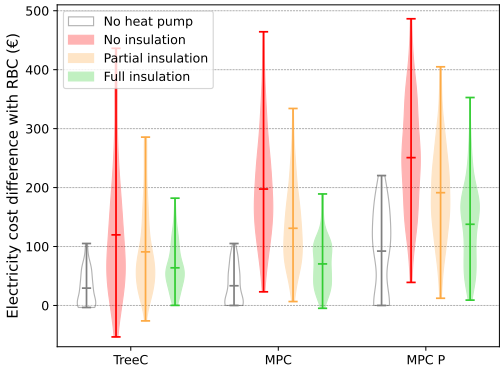

(c) Insulation

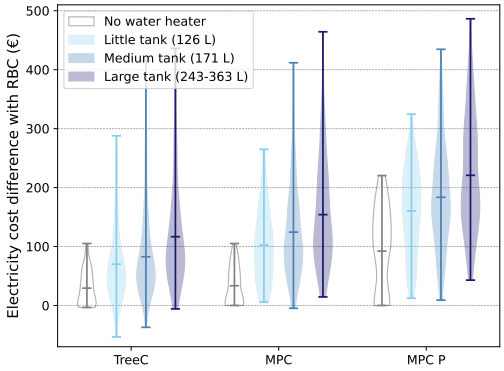

(d) Tank volume

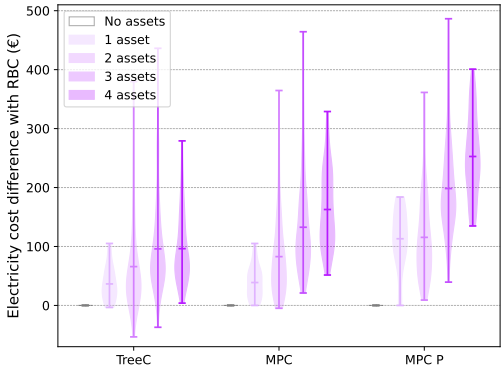

(e) Asset number

Figure 3.9: Comparison of EMS performance for different dwelling features: (a) PV presence, (b) EV charger presence, (c) insulation level, and (d) water heater's tank volume (e) number of assets.

explain its lower performance compared to TreeC in this area.

Another relevant aspect to consider is how the peak and hourly day-ahead costs influence the total grid power consumption of the 500 dwellings. Figure 3.11 shows on 24 hours the mean total grid power for each EMS. The MPC and MPC P EMSs show almost identical average behaviour. Overall TreeC and the MPC EMSs consume more than RBC during hours with lower day ahead prices, and

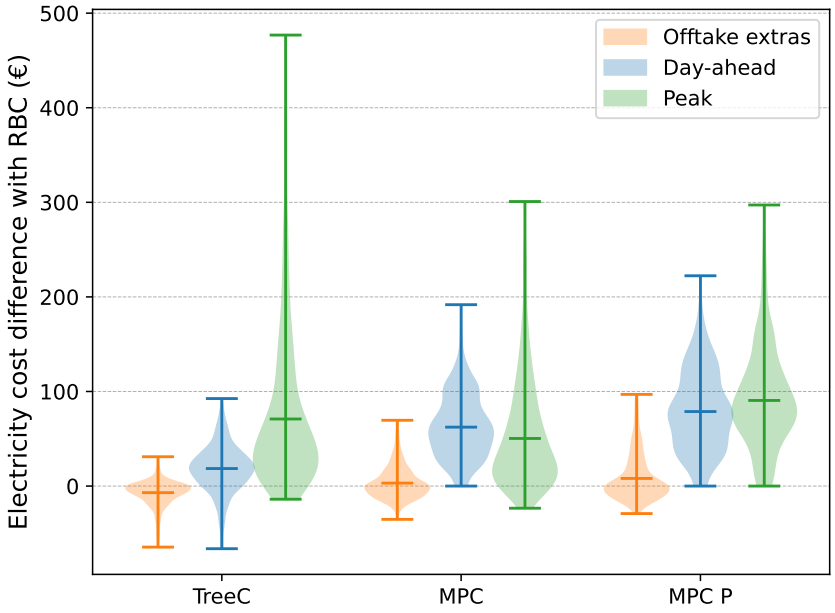

Figure 3.10: Breakdown of cost savings relative to the RBC EMS for each cost component defined in eq. (3.7). Offtake extras costs improvement reflects better energy efficiency, day-ahead costs improvement reflects better hourly price optimization, and peak costs improvement reflects reductions in peak grid power. Each box plot summarizes the results across the 500 dwellings.

all manage reduce the evening RBC total peak. However, the MPC EMSs show a clear peak for the last time step of the hours between 04:00 and 08:00 and between 14:00 and 18:00 which are the times the dwelling heat up to reach a higher comfort temperature. This is because the MPC EMSs calculates that it is better to heat the dwelling at the last time step of the hour to gain advantage of the low day-ahead price of that hour and at the same time loose as little as possible of the heat bought at that price. Buying at an earlier time step of that hour would result in more of the bought heat being lost to the exterior. The TreeC EMSs didn't learn this behaviour and therefore has a much smoother total grid power consumption profile.

The peaks caused by the MPC EMSs are problematic for the grid because: 1) they can cause local congestion, requiring costly grid reinforcement, and 2) the sharp, short-lived spikes in aggregate demand increase intra-hour variability, making grid balancing more challenging for operators. This undermines the intent of hourly pricing signals, which are designed to encourage demand shifting across entire hours rather than concentrating consumption at specific moments within those hours.

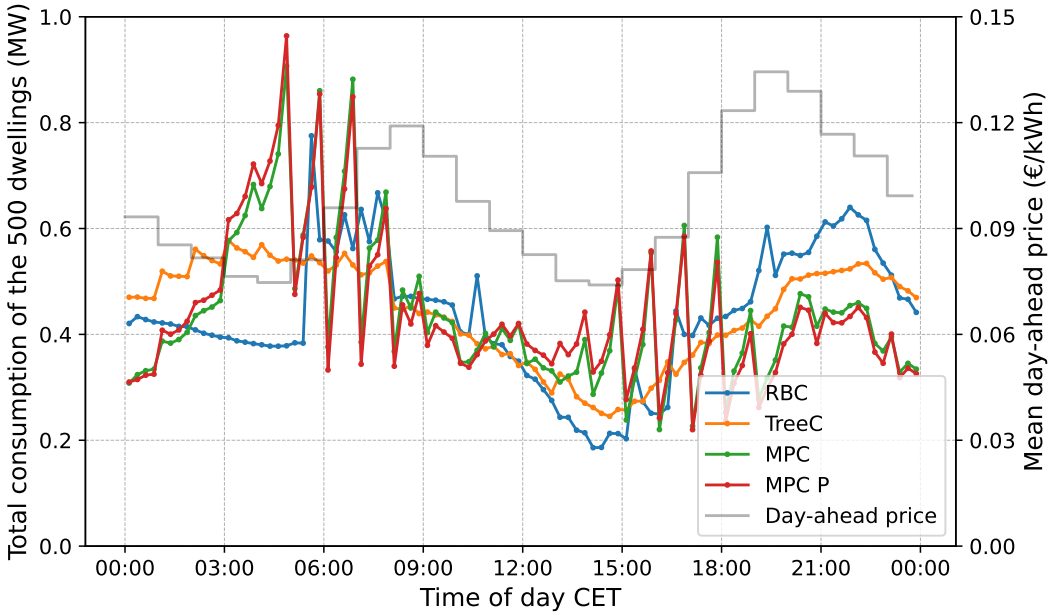

Figure 3.11: Yearly mean of the total grid power consumption of the 500 dwellings for each time step of the day. Each EMS is represented as well as the mean day-ahead prices for the year 2023.

3.6 Discussion

When analysing all the results, TreeC and MPC have their strengths and weaknesses. TreeC is more interpretable (see fig. 3.5), achieves higher EV charge completion rates, and results in a smoother total grid power consumption profile when all dwellings are equipped with the same EMS. However, MPC obtains more cost reductions, which is arguably the most important metric for evaluating EMS performance.

To improve the cost reduction of TreeC, future work should focus on improving its ability to reduce the day-ahead costs. This can be done by retraining every month the EMS to adapt to the changing day-ahead prices, adding more price related inputs that can be selected by the decision trees, or using more sophisticated action on the leaf nodes of the decision trees. Improving the training time of TreeC through reduced simulation speed, improved training methodology, or transfer learning is also necessary to make it more practical for real-world applications.

Table 3.13: Mean absolute error (MAE) of two forecasts for the EV charging sessions. Values show the mean and standard deviation of MAE across all dwellings with an EV charger for 2023.

Forecast method	Predicted value	MAE (Mean ± Standard deviation)
K-NN regression	Connection time	7h14min ± 4h33min
Repeat last	Connection time	9h14min ± 5h7min
K-NN regression	SOC demand	12.8% ± 3.4%
Repeat last	SOC demand	16.4% ± 5.4%

The results for MPC P indicate that improved forecasting can enhance both cost savings and EV charge completion. However, accurately predicting EV user behaviour and charging sessions remains challenging due to the high variability observed in table 3.11 and the substantial mean absolute errors (MAE) for both K-NN regression and "repeat last" forecasts in table 3.13 exceeding 7 hours for connection time and 12% for SOC demand. To address this, future work should consider multi-scenario or stochastic MPC approaches that explicitly account for the uncertainty in EV charging behaviour, as conventional forecasting methods struggle with this aspect.

The mean annual savings compared to the RBC EMS are 82.5€ for TreeC, 116.6€ for MPC, and 177.4€ for MPC P. These savings are relatively modest compared to the cost of buying, installing, maintaining, and subscription fees of current home EMS solutions on the market (between 280€ and 680€ to buy and 0€ and 360€ per year for subscription fees [34]). Specific cases can justify an EMS implementation as the results do show some cost reduction going up to 464€ are possible, but these are not the norm.

Different factors could increase the commercial attractiveness of advanced EMSs:

- More volatility in electricity prices. Having more extreme low and high prices would increase the potential savings of shifting consumption towards low price periods.
- Allowing participation in balancing markets. This would create additional revenue streams for households providing flexibility to the grid.
- A wider adoption of home batteries. The additional flexibility provided by batteries could increase the potential savings from an advanced EMS.
- Wider adoption of standardized protocols in large electrical appliances. Standardization would lower EMS integration costs by making it easier to connect and control devices, potentially allowing simple plug-and-play EMS installation.

The code and data for this study are open-source, enabling other researchers to benchmark their EMS solutions in the same context. All materials are available at: https://github.com/EVERGi/ems_adoption_electrified_belgium_paper However, further algorithmic improvements alone are unlikely to deliver greater benefits than those demonstrated by MPC P.

3.7 Conclusion

This chapter presented a comparison of three EMSs for electrified residential buildings in Belgium: a rule-based controller (RBC), a model predictive control (MPC) EMS, and a decision tree-based EMS called TreeC. The study evaluated the performance across one year simulations of 500 dwellings with characteristics representative of current Belgian dwellings but where the fossil fuel based assets for heating, hot water and car transportation are replaced by electric equivalents.

On average, TreeC and MPC achieved cost reductions of 82.5 € and 116.6 €, respectively, compared to the RBC EMS, which had a mean total cost of 2021 €. MPC obtained better cost reduction but worse EV charge completion with an average of 88.8% SOC at departure compared to 93.3% for TreeC with a full charge set at 95% SOC. The results show that TreeC and MPC performed better than RBC in terms of cost for dwellings, with PV systems, with EV chargers, larger water tanks, and more assets in general which is expected, it also worked better for dwellings with less insulation which is a less intuitive result. The total power consumption of the 500 dwellings together obtained large volatility from one time step to the next when all dwellings use MPC, while TreeC obtained a smoother total power consumption profile.

The main recommended improvement for TreeC is to enhance its ability to reduce hourly day-ahead costs. This could be achieved by retraining the EMS more frequently to adapt to changing electricity over time or by making the decision tree model more responsive to hourly price signals. For MPC, future work should focus on addressing the uncertainty in EV charging sessions, for example by adopting stochastic or multi-scenario MPC approaches.

In most cases, the cost savings achieved by advanced EMSs are modest relative to the investment and effort required for their implementation, when compared to the default operation of assets (RBC EMS). Further algorithmic improvements are unlikely to deliver significantly greater benefits than a mean cost reduction of 177.4 € achieved by the MPC with perfect forecast (MPC P) EMS. Increasing the potential cost savings, and thus the attractiveness of advanced EMSs, could be achieved by modifying electricity pricing structures, enabling participation in flexibility markets, reducing EMS investment and maintenance costs, or integrating home batteries.

[1] D. Aerts, J. Minnen, I. Glorieux, I. Wouters, and F. Descamps. "Discrete Occupancy Profiles From Time-use Data For User Behaviour Modelling In Homes". In: Building Simulation 2013. Vol. 13. Building Simulation. IBPSA, 2013, pp. 2421–2427. ISBN: 978-2-7466-6294-0. DOI: 10.26868/25222708.2013.1273.

[2] M. Bain and C. Sammut. "A Framework for Behavioural Cloning". In: *Machine Intelligence 15, Intelligent Agents [St. Catherine's College, Oxford, July 1995]*. GBR: Oxford University, Jan. 1, 1999, pp. 103–129. ISBN: 978-0-19-853867-7.

[3] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl. "Algorithms for Hyper-Parameter Optimization". In: *Advances in Neural Information Processing Systems*. Vol. 24. Curran Associates, Inc., 2011.

[4] F. Biscani and D. Izzo. "A Parallel Global Multiobjective Framework for Optimization: Pagmo". In: *Journal of Open Source Software* 5.53 (Sept. 13, 2020), p. 2338. ISSN: 2475-9066. DOI: 10.21105/joss.02338.

[5] Brugel. *Énergies Renouvelables En Région de Bruxelles-Capitale*. 2024. URL: https://app.powerbi.com/view?r=eyJrIjoiNDkyNWIyNDgtNWNkNi00MWY2LTgxY2QtZTZlZWI2MDM1YmRhIiwidCI6ImMwYjg2YzA3LWRhZGUtNDkyMC1hYzEzLWIwZWNhZDNiMmM5NSIsImMiOjh9 (visited on 02/28/2025).

[6] Car Pass. *Rapport annuel 2024*. Mar. 21, 2025. URL: https://www.car-pass.be/fr/blog/car-pass-jaarverslag-2024 (visited on 10/03/2025).

[7] G. Ceusters, L. R. Camargo, R. Franke, A. Nowé, and M. Messagie. "Safe Reinforcement Learning for Multi-Energy Management Systems with Known Constraint Functions". In: *Energy and AI* 12 (Apr. 1, 2023), p. 100227. ISSN: 2666-5468. DOI: 10.1016/j.egyai.2022.100227.

[8] G. Ceusters, M. A. Putratama, R. Franke, A. Nowé, and M. Messagie. "An Adaptive Safety Layer with Hard Constraints for Safe Reinforcement Learning in Multi-Energy Management Systems". In: *Sustainable Energy, Grids and Networks* 36 (Dec. 1, 2023), p. 101202. ISSN: 2352-4677. DOI: 10.1016/j.segan.2023.101202.

[9] G. Ceusters, R. C. Rodríguez, A. B. García, R. Franke, G. Deconinck, L. Helsen, A. Nowé, M. Messagie, and L. R. Camargo. "Model-Predictive Control and Reinforcement Learning in Multi-Energy System Case Studies". In: *Applied Energy* 303 (Dec. 1, 2021), p. 117634. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2021.117634.

[10] T. Cover and P. Hart. "Nearest Neighbor Pattern Classification". In: *IEEE Transactions on Information Theory* 13.1 (Jan. 1967), pp. 21–27. ISSN: 1557-9654. DOI: 10.1109/TIT.1967.1053964.

[11] D. B. Crawley et al. "EnergyPlus: Creating a New-Generation Building Energy Simulation Program". In: *Energy and Buildings*. Special Issue: BUILDING SIMULATION'99 33.4 (Apr. 1, 2001), pp. 319–331. ISSN: 0378-7788. DOI: 10.1016/S0378-7788(00)00114-6.

[12] DAIAD Project. *Smart Water Meter Consumption Time Series*. May 5, 2020. URL: https://data.hellenicdataservice.gr/dataset/78776f38-a58b-4a2a-a8f9-85b964fe5c95 (visited on 10/23/2024).

[13] davidusb-geek. *Energy Management for Home Assistant*. June 21, 2025. URL: https://github.com/davidusb-geek/emhass (visited on 06/24/2025).

[14] P. S. Dolan, M. H. Nehrir, and V. Gerez. "Development of a Monte Carlo Based Aggregate Model for Residential Electric Water Heater Loads". In: *Electric Power Systems Research* 36.1 (Jan. 1, 1996), pp. 29–35. ISSN: 0378-7796. DOI: 10.1016/0378-7796(95)01011-4.

[15] J. Drgoňa et al. "All You Need to Know about Model Predictive Control for Buildings". In: *Annual Reviews in Control* 50 (Jan. 1, 2020), pp. 190–232. ISSN: 1367-5788. DOI: 10.1016/j.arcontrol.2020.09.001.

[16] Énergie+. *Coefficient de transmission thermique de parois types peu ou non isolées*. Energie Plus Le Site. Oct. 26, 2007. URL: https://energieplus-lesite.be/theories/enveloppe9/coefficient-de-transmission-thermique/valeurs-de-coefficients-de-transmission-thermique-u-de-parois-types/ (visited on 03/03/2025).

[17] Energuide.be. *Double Glazing, Triple Glazing or Super-Insulating Glazing – Which Should You Choose?* Energuide. URL: https://www.energuide.be/en/questions-answers/double-glazing-triple-glazing-or-super-insulating-glazing-which-should-you-choose/215/ (visited on 03/03/2025).

[18] EurObserv'ER. *22nd Report on The State of Renewable Energies in Europe*. 2023. URL: https://www.eurobserv-er.org/22nd-annual-overview-barometer/ (visited on 09/02/2024).

[19] EV Database. *Energy Consumption of Full Electric Vehicles*. EV Database. 2025. URL: https://ev-database.org/cheatsheet/energy-consumption-electric-car (visited on 10/03/2025).

[20] S. Farhad and A. Nazari. "Introducing the Energy Efficiency Map of Lithium-Ion Batteries". In: *International Journal of Energy Research* 43.2 (2019), pp. 931–944. ISSN: 1099-114X. DOI: 10.1002/er.4332.

[21] E. Fuentes, L. Arce, and J. Salom. "A Review of Domestic Hot Water Consumption Profiles for Application in Systems and Buildings Energy Performance Analysis". In: *Renewable and Sustainable Energy Reviews* 81 (Jan. 1, 2018), pp. 1530–1547. ISSN: 1364-0321. DOI: 10.1016/j.rser.2017.05.229.

[22] S. Fujimoto, H. Hoof, and D. Meger. "Addressing Function Approximation Error in Actor-Critic Methods". In: *Proceedings of the 35th International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 3, 2018, pp. 1587–1596.

[23] E. van Gastel. "De harde cijfers | Vorig jaar ruim 19.000 thuisbatterijen geïnstalleerd in Vlaanderen". In: *Solar Magazine* (2025).

[24] E. Genov, C. D. Cauwer, G. V. Kriekinge, T. Coosemans, and M. Messagie. "Forecasting Flexibility of Charging of Electric Vehicles: Tree and Cluster-Based Methods". In: *Applied Energy* 353 (Jan. 1, 2024), p. 121969. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2023.121969.

[25] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2024. URL: https://www.gurobi.com.

[26] Home assistant. *Home Assistant*. Home Assistant, June 24, 2025. URL: https://github.com/home-assistant/core (visited on 06/24/2025).

[27] S. Horowitz, J. Robertson, Y. Zhou, B. Park, N. Merket, and A. Speake. *OpenStudio® HPXML Workflow [SWR-25-13]*. National Renewable Energy Laboratory (NREL), Golden, CO (United States), 2024. DOI: 10.11578/DC.20241125.4. URL: https://www.osti.gov/doecode/biblio/147986 (visited on 02/28/2025).

[28] J.-P. Hubert and P. Toint. *La Mobilité quotidienne des Belges*. Presses universitaires de Namur, 2003. 358 pp. ISBN: 978-2-87037-388-0. Google Books: wUH8rUNo3qMC.

[29] International Electrotechnical Commission. *IEC 61851-1: 2017 Electric Vehicle Conductive Charging System-Part 1: General Requirements*. Geneva, 2017.

[30] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu. "LightGBM: A Highly Efficient Gradient Boosting Decision Tree". In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017.

[31] J. Kennedy and R. Eberhart. “Particle Swarm Optimization”. In: *Proceedings of ICNN’95 - International Conference on Neural Networks*. Proceedings of ICNN’95 - International Conference on Neural Networks. Vol. 4. Nov. 1995, 1942–1948 vol.4. DOI: 10.1109/ICNN.1995.488968.

[32] C. Lallemand. “"16% des ménages wallons sont équipés de panneaux solaires, mais les autres technologies vertes sont désespérément à la traîne"”. In: *Trends-Tendances. a-la-une* (Nov. 9, 2023).

[33] J. Laverge, M. Delghust, N. V. D. Bossche, and A. Janssens. “Airtightness Assessment of Single Family Houses in Belgium”. In: *International Journal of Ventilation* 12.4 (Mar. 1, 2014), pp. 379–390. ISSN: 1473-3315. DOI: 10.1080/14733315.2014.11684031.

[34] Maak Je Meter Slim. *Alle EMS- Producten*. June 24, 2025. URL: https://maakjemeterslim.be/ems-producten (visited on 06/24/2025).

[35] H. Madsen and J. Holst. “Estimation of Continuous-Time Models for the Heat Dynamics of a Building”. In: *Energy and Buildings* 22.1 (Mar. 1, 1995), pp. 67–79. ISSN: 0378-7788. DOI: 10.1016/0378-7788(94)00904-X.

[36] I. Meireles, V. Sousa, B. Bleys, and B. Poncelet. “Domestic Hot Water Consumption Pattern: Relation with Total Water Consumption and Air Temperature”. In: *Renewable and Sustainable Energy Reviews* 157 (Apr. 1, 2022), p. 112035. ISSN: 1364-0321. DOI: 10.1016/j.rser.2021.112035.

[37] Oikolab. *Weather Data Downloader*. URL: https://weatherdownloader.oikolab.com/app (visited on 03/03/2025).

[38] L. Peeters, R. de Dear, J. Hensen, and W. D’haeseleer. “Thermal Comfort in Residential Buildings: Comfort Values and Scales for Building Energy Simulation”. In: *Applied Energy* 86.5 (May 1, 2009), pp. 772–780. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2008.07.011.

[39] A. P. Piotrowski, J. J. Napiorkowski, and A. E. Piotrowska. “Population Size in Particle Swarm Optimization”. In: *Swarm and Evolutionary Computation* 58 (Nov. 1, 2020), p. 100718. ISSN: 2210-6502. DOI: 10.1016/j.swevo.2020.100718.

[40] R. Poli, J. Kennedy, and T. Blackwell. “Particle Swarm Optimization”. In: *Swarm Intelligence* 1.1 (June 1, 2007), pp. 33–57. ISSN: 1935-3820. DOI: 10.1007/s11721-007-0002-0.

[41] PVOutput.org. *PVOutput.Org*. URL: https://pvoutput.org/ (visited on 03/05/2025).

[42] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann. “Stable-Baselines3: Reliable Reinforcement Learning Implementations”. In: *Journal of Machine Learning Research* 22.268 (2021), pp. 1–8. ISSN: 1533-7928.

[43] Recticel Insulation. *Baromètre de l'isolation 2020*. 2020. URL: https://www.recticelinsulation.com/be-fr/barometre-de-lisolation-2020 (visited on 02/28/2025).

[44] J. Ruddick, G. Ceusters, G. Van Kriekinge, E. Genov, C. De Cauwer, T. Coosemans, and M. Messagie. "Real-World Validation of Safe Reinforcement Learning, Model Predictive Control and Decision Tree-Based Home Energy Management Systems". In: *Energy and AI* 18 (Dec. 1, 2024), p. 100448. ISSN: 2666-5468. DOI: 10.1016/j.egyai.2024.100448.

[45] J. Ruddick, L. Ramirez Camargo, M. A. Putratama, M. Messagie, and T. Coosemans. "TreeC: A Method to Generate Interpretable Energy Management Systems Using a Metaheuristic Algorithm". In: *Knowledge-Based Systems* 309 (Jan. 30, 2025), p. 112756. ISSN: 0950-7051. DOI: 10.1016/j.knosys.2024.112756.

[46] M. Schlemminger, T. Ohrdes, E. Schneider, and M. Knoop. "Dataset on Electrical Single-Family House and Heat Pump Load Profiles in Germany". In: *Scientific Data* 9.1 (Feb. 15, 2022), p. 56. ISSN: 2052-4463. DOI: 10.1038/s41597-022-01156-1.

[47] Å. L. Sørensen, I. Sartori, K. B. Lindberg, and I. Andresen. "A Method for Generating Complete EV Charging Datasets and Analysis of Residential Charging Behaviour in a Large Norwegian Case Study". In: *Sustainable Energy, Grids and Networks* 36 (Dec. 1, 2023), p. 101195. ISSN: 2352-4677. DOI: 10.1016/j.segan.2023.101195.

[48] Å. L. Sørensen, K. B. Lindberg, I. Sartori, and I. Andresen. "Residential Electric Vehicle Charging Datasets from Apartment Buildings". In: *Data in Brief* 36 (June 1, 2021), p. 107105. ISSN: 2352-3409. DOI: 10.1016/j.dib.2021.107105.

[49] Å. L. Sørensen, I. Sartori, K. B. Lindberg, and I. Andresen. "Electric Vehicle Charging Dataset with 35,000 Charging Sessions from 12 Residential Locations in Norway". In: *Data in Brief* 57 (Dec. 1, 2024), p. 110883. ISSN: 2352-3409. DOI: 10.1016/j.dib.2024.110883.

[50] Statbel. *Census 2021 - Nombre de Logements Classiques Occupés Selon Le Type de Bâtiment*. Jan. 1, 2021. URL: https://statbel.fgov.be/fr/themes/census/logement/type-de-logement#panel-12 (visited on 02/28/2025).

[51] Statbel. *Census 2021 - Nombre de Logements Selon Le Nombre de Pièces*. Jan. 1, 2021. URL: https://statbel.fgov.be/sites/default/files/files/documents/Census2021/T04_NOR_BE_FR.XLSX (visited on 02/28/2025).

[52] Statbel. *Census 2021 - Nombre de Logements Selon Le Type de Chauffage*. Jan. 1, 2021. URL: https://statbel.fgov.be/sites/default/files/files/documents/Census2021/T04_TOH_BE_FR.XLSX (visited on 02/28/2025).

[53] V., Sébastien. *Les machines à laver consomment-elles plus l'hiver que l'été ?* ENGIE. Feb. 11, 2020. URL: https://www.engie.be/fr/blog/conseils-energie/impact-temperature-eau-sur-consommation-machine-laver/ (visited on 10/21/2024).

[54] S. I. Vagropoulos and A. G. Bakirtzis. "Optimal Bidding Strategy for Electric Vehicle Aggregators in Electricity Markets". In: *IEEE Transactions on Power Systems* 28.4 (Nov. 2013), pp. 4031–4041. ISSN: 1558-0679. DOI: 10.1109/TPWRS.2013.2274673.

[55] D. Vanneste, I. Thomas, and L. Goossens. *Le Logement En Belgique*. Bruxelles: SPF Economie, 2007.

[56] Vlaams Energie- en Klimaatagentschap. *Zonnepanelen in Vlaanderen*. 2024. URL: https://apps.energiesparen.be/energiekaart/vlaanderen/zonnepanelen (visited on 02/28/2025).

[57] VREG. *Energieverbruik*. 2024. URL: https://www.vlaamsenutsregulator.be/nl/energieverbruik (visited on 09/24/2025).