

Distill2Explain: Differentiable decision trees for explainable reinforcement learning in energy application controllers

Gargya Gokhale

gargya.gokhale@ugent.be
IDLab, Ghent University – imec
Ghent, Belgium

Bert Claessens

IDLab, Ghent University – imec
Beebop.ai
Antwerp, Belgium

Seyed Soroush Karimi Madahi

IDLab, Ghent University – imec
Ghent, Belgium

Chris Develder

IDLab, Ghent University – imec
Ghent, Belgium

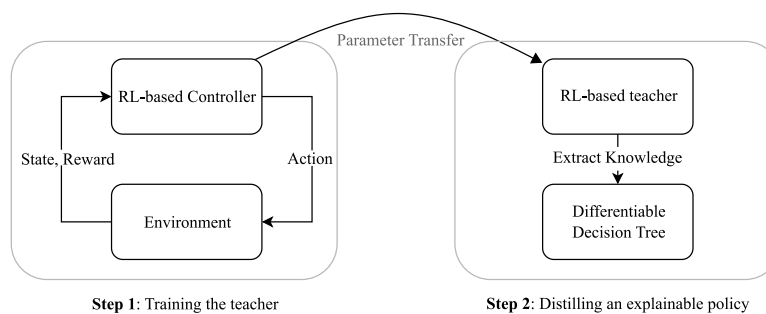


Figure 1: Proposed policy distillation approach using differentiable decision trees to obtain explainable reinforcement learning controllers.

ABSTRACT

Demand-side flexibility is gaining importance as a crucial element in the energy transition process. Accounting for about 25% of final energy consumption globally, the residential sector is an important (potential) source of energy flexibility. However, unlocking this flexibility requires developing a control framework that (1) easily scales across different houses, (2) is easy to maintain, and (3) is simple to understand for end-users. A potential control framework for such a task is data-driven control, specifically model-free reinforcement learning (RL). Such RL-based controllers learn a good control policy by interacting with their environment, learning purely based on data and with minimal human intervention. Yet, they lack explainability, which hampers user acceptance. Moreover, limited hardware capabilities of residential assets forms a hurdle (e.g., using deep neural networks). To overcome both those challenges, we propose a novel method to obtain explainable RL policies by using differentiable decision trees. Using a policy distillation approach,

we train these differentiable decision trees to mimic standard RL-based controllers, leading to a decision tree-based control policy that is data-driven and easy to explain. As a proof-of-concept, we examine the performance and explainability of our proposed approach in a battery-based home energy management system to reduce energy costs. For this use case, we show that our proposed approach can outperform baseline rule-based policies by about 20-25%, while providing simple, explainable control policies. We further compare these explainable policies with standard RL policies and examine the performance trade-offs associated with this increased explainability.

CCS CONCEPTS

• **Theory of computation** → **Reinforcement learning**; • **Hardware** → **Smart grid**; • **Computing methodologies** → *Rule learning*; *Apprenticeship learning*.

KEYWORDS

Reinforcement Learning, Explainable AI, Home Energy Management, Differentiable decision trees, Policy Distillation

ACM Reference Format:

Gargya Gokhale, Seyed Soroush Karimi Madahi, Bert Claessens, and Chris Develder. 2024. Distill2Explain: Differentiable decision trees for explainable reinforcement learning in energy application controllers. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24)*, June 04–07, 2024, Singapore, Singapore. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3632775.3661937>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

E-Energy '24, June 04–07, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0480-2/24/06

<https://doi.org/10.1145/3632775.3661937>

1 INTRODUCTION

The ongoing shift towards sustainable energy is leading to a significant restructuring of the energy sector: large-scale integration of distributed renewable energy sources, increased electrification, phasing out of fossil fuel-based generation, etc. [18]. As a result of these changes, there is a growing need for grid balancing services and demand-side flexibility to ensure reliable and secure functioning of the grid. Conventionally, large industries and big consumers were the primary source of such demand-side flexibility. However, another important and as-of-yet untapped source of flexibility is the residential sector [17].

Households account for about 25% of the final energy consumption and with growing adoption of rooftop solar PVs, home batteries, heat pumps, etc., represent an appealing source of flexibility [9]. Usually, exploiting this flexibility entails optimizing the use of a battery or other flexible assets to shift the real-time consumption of households while ensuring user comfort [20]. In most cases, the primary objective is to minimize the energy bill of the household, however, prior research has also investigated other objectives such as maximizing self-consumption or participation in other explicit demand response services [22, 28].

An important component for extracting this household flexibility is a home energy management system (HEMS), responsible for solving the underlying, non-linear sequential decision-making problem and calculating the necessary control actions to be taken in real-time. Developing HEMS has been a major research area, with works such as [15, 42] providing an overview of techniques used in literature.

A prominent research direction in this context is the use of model-predictive control (MPC) algorithms. MPC forms an advanced control framework that relies on a model of the system to predict the system's behavior and uses the model to analytically obtain optimal actions [8]. Works such as [10, 14, 31] have investigated the application of MPC in both simulation and real-world scenarios, showing significant performance improvements in such systems. However, as highlighted in [2, 41], accurate models—which an MPC requires—of the system are often difficult to obtain in the residential sector, significantly limiting widespread adoption of MPC-based solutions in this sector.

The residential sector necessitates control frameworks that can easily scale to many, potentially diverse households. This has led to an increased interest in data-driven control frameworks, especially based on reinforcement learning (RL). RL-based controllers work by continuously interacting with the environment (i.e., the household), collecting experience (data) from these interactions, and using them to learn a control policy that maximizes a predefined reward [36]. Thus, with little human intervention and relying completely on data, such RL-based controllers can learn good control policies. Previous works on RL-based HEMS controllers such as [1, 7, 25] have shown significant improvements over baseline scenarios.

However, most RL-based research is limited to simulation environments or specialized buildings. As discussed in [32], this is due to two main factors: (i) the data inefficiency of RL training, and (ii) the opaque nature of obtained control policies. To address (i), i.e., the high amount of data required for training RL-based controllers, previous works such as [3, 39, 44] propose different solutions. However, (ii) raises another important concern related

to RL, i.e., the non-interpretable/non-explainable nature of their policies, especially when based on deep neural networks. With limited prior works in this area, we identify this as a significant gap in existing literature and thus introduce our innovative approach to specifically address the (lack of) explainability of RL-based HEMS.

More specifically, we propose a policy distillation framework using differentiable decision trees [11, 19]. The key idea is to distill information from pre-trained RL-based controllers into an explainable decision tree, leading to control policies that are explainable and perform nearly as good as the original RL-based policies. To the best of our knowledge, this is one of the first works in the energy field to adopt policy distillation using differentiable decision trees for explainable RL. Our main contributions can be summarized as:

- (1) We propose a novel framework for explainable RL that uses differentiable decision trees and policy distillation for converting black-box RL policies into explainable decision trees (§4).
- (2) Using different case studies, we detail the explainability of our proposed method, contrasting it with conventional RL-based policies (§6.2).
- (3) We compare the performance of our method with conventional RL-based policies and other baselines to show the performance trade-off that results from the increased explainability (§6.1).

The primary emphasis of this paper is to introduce a novel method for obtaining explainable RL-based control policies. As a proof-of-concept, we validate our proposed approach on a battery-based home energy management scenario using real-world data and present our preliminary findings. Section §7.1 outlines the future work in terms of additional investigation of this method and its application to other, more complex scenarios.

2 RELATED WORK

Designing control algorithms for unlocking flexibility in households has been a major field of research, with works such as [15, 42] providing an exhaustive review of prior works including heuristics-based controllers, MPCs, and data-driven algorithms. As discussed in §1, our work focuses on improving the explainability of reinforcement learning-based controllers and hence this section focuses on developments in the fields of reinforcement learning-based control, policy distillation, and explainable AI. We refer interested readers to [15, 28] for more comprehensive reviews of other relevant methods in the context of HEMS and demand-side flexibility.

2.1 Data-driven Home Energy Management Systems

A recent research direction in HEMS has been the use of data-driven and mainly reinforcement learning-based controllers [5]. RL-based controllers rely primarily on past data and have minimal modeling requirements as compared to prominent control techniques such as MPCs. For example, works such as [1, 7, 25], demonstrate the applications of RL-based controllers in the context of HEMS. In most of these cases, the RL-based controllers rely on state-of-the-art RL algorithms such as deep Q -networks (DQN) [30], deep deterministic policy gradient (DDPG) [23] and use control policies based

on deep neural networks to achieve significant performance improvements (~5-16% as reported in these works). While these deep neural networks are beneficial for achieving good performance, a common drawback associated with their use is their opaque (or black-box) control policy [38]. We aim to address this challenge associated with the explainability of RL-based controllers, providing a framework for distilling a standard RL control policy into an explainable policy.

2.2 Explainable AI

Providing explainability for AI-based technology is an important and necessary issue to address for large-scale deployment of machine learning-based solutions, especially in the context of the energy sector. We refer to more exhaustive reviews [24, 29] of available techniques, metrics, and methodologies across different fields such as image recognition, natural language processing, etc. However, as discussed in [27], in the context of energy, research on explainable AI has been largely restricted to applications such as forecasting, modeling, or fault diagnosis. While few works such as [21, 40, 43], present explainable RL-based controllers, they largely rely on decomposition methods or utilize post-hoc explanation frameworks such as SHAPley values, feature importances, LIME, etc. Although useful, such post-hoc explanations are typically designed for experts and are not easily accessible to the average end-user, such as a homeowner. Our proposed method differs from such approaches by distilling the deep RL-based control policy into an explainable architecture in the form of differentiable decision trees. Thus, the resulting control policies are structurally explainable, i.e., in the form of rather simple *if-then-else* rules, that can be easily (a) explained to non-expert end users, and (b) deployed on simple hardware.

2.3 Policy Distillation and Differentiable Decision Trees

As discussed in §1, our approach employs policy distillation to trained RL-based controllers and distills their knowledge into a differentiable decision tree structure. This closely follows prior works that have used knowledge distillation strategies to (1) compress large neural networks, or (2) combine knowledge from model ensembles into a single model [4, 16]. Differing from these, works such as [6, 12, 34] adopt knowledge distillation in RL to transform the architecture of the final policy, e.g., into a fuzzy inference system. We follow a similar approach and distill an RL-based control policy into a differentiable decision tree. This enables us to extract knowledge from standard RL-based controllers into simple decision trees which are structurally easy-to-explain and simple to understand. This choice is closely related to the objective of obtaining control policies that are easy to explain (to both energy experts and end-users).

Differentiable decision trees (DDTs) or soft decision trees are variants of binary decision trees that can be trained using gradient descent [11, 19]. Prior works such as [16, 26] have applied DDTs to computer vision and regression tasks. For our energy use case, we follow the approach of [6], using DDTs to distill RL-based control policies. However, our proposed approach differs from [6] in two ways: (i) we learn deterministic decision trees instead of the soft decision trees, and (ii) we learn using observed, explainable features

(as opposed to rather indirect pixel-based learning in [6]). This enables us to learn DDTs using gradient descent and then convert them into simple decision decision trees for inference (as detailed further in §4).

3 PRELIMINARIES

The proposed differentiable decision tree-based policy distillation framework is examined in the context of a home energy management scenario. This section describes the problem formulation for this proof-of-concept and introduces basic concepts related to reinforcement learning (RL).

3.1 Problem Formulation

In the context of home energy management, we consider an average Belgian household with a rooftop solar PV installation (with generated power P_t^{PV}), non-flexible electrical load (P_t^{con}), and a home battery. We assume that this household is exposed to varying BELPEX day-ahead prices (λ_t^{con}) and a capacity tariff based on peak power. This leads to a joint optimization problem, where the HEMS must minimize the daily cost of both the energy consumption (c_t^{eng}) and the peak power (c_t^p). This optimization problem is modeled as:

$$\min_{u_1, \dots, u_T} \sum_{t=1}^T c_t^{\text{eng}} + c_t^p \quad (1a)$$

$$\text{s.t.}: c_t^{\text{eng}} = \begin{cases} \lambda_t^{\text{con}} P_t^{\text{agg}} \Delta t & : P_t^{\text{agg}} \geq 0 \\ \lambda_t^{\text{inj}} P_t^{\text{agg}} \Delta t & : P_t^{\text{agg}} < 0 \end{cases} \quad \forall t \quad (1b)$$

$$c_t^p = \lambda^{\text{cap}} \max(P_t^{\text{agg}}, P_{\min}^{\text{agg}}) \quad (1c)$$

$$P_t^{\text{agg}} = P_t^{\text{con}} + P_t^{PV} + u_t \quad \forall t \quad (1d)$$

$$E_{t+1} = \begin{cases} E_t + \eta u_t \Delta t & : u_t \geq 0 \\ E_t + \frac{1}{\eta} u_t \Delta t & : u_t < 0 \end{cases} \quad \forall t \quad (1e)$$

$$0 \leq E_t \leq E^{\text{max}}; u^{\text{min}} \leq u_t \leq u^{\text{max}} \quad \forall t. \quad (1f)$$

The battery is modeled using a linear model (Eq. (1e) with charging/discharging actions u_t and current energy level (E_t)). The cost of energy consumed (c_t^{eng}) depends on the power consumed (P_t^{agg}) and the current injection and consumption prices (λ_t^{inj} and λ_t^{con} respectively). Similarly, the capacity cost (c_t^p) depends on the actual power consumed and the minimum power capacity contracted [37]. Furthermore, we assume $T = 24$ hours and a time resolution $\Delta t = 1$ hour.

The above-mentioned problem illustrates a real-world scenario that is pertinent in the present day where a household's HEMS needs to efficiently leverage the home battery to reduce the energy bill, taking charging/discharging actions dependent on the real-time price, solar PV production, and daily load consumption patterns. Accordingly, we further assume that the HEMS can only take discrete actions (a total of 5 related to 2 charging modes, 2 discharging modes, and 1 'do nothing' mode). Nonetheless, our method can be easily extended to other action spaces as well.

3.2 Markov Decision Process

We model the sequential decision-making problem presented in §3.1 as a Markov Decision Process (MDP) [36]. The states ($\mathbf{x}_t \in \mathbf{X}$) consist of the current price (λ_t^{con}), battery state-of-charge, non-flexible demand (P_t^{con}), and solar PV generation (P_t^{PV}). The actions ($u_t \in \mathbf{U}$) are the charging/discharging signals given to the battery. As stated above, we assume a discrete action space of 5 elements (i.e., $\mathbf{U} = \{-1, -0.5, 0, 0.5, 1\}$), with the possibility of extending it reserved for future work. The reward function ($\rho : \mathbf{X} \times \mathbf{U} \rightarrow \mathbb{R}$) is defined as the cost incurred for each time step t and is modeled based on Eq. (1b), Eq. (1c). The transition function (f) models the dynamics of the household, taking into account the (controllable) behavior of the battery along with (uncontrollable) real-time solar PV generation, and non-flexible power consumption.

3.3 Reinforcement Learning

In RL, the goal of an agent is to find a policy $\pi : \mathbf{X} \rightarrow \mathbf{U}$ that minimizes the expected T -step cost (J^π) starting from an initial state $\mathbf{x}_0 \in \mathbf{X}$ (Eq. (2)).

$$J^\pi = \sum_{t=0}^T \gamma^t \rho(\mathbf{x}_t, \pi(\mathbf{x}_t), \omega) \quad (2)$$

This expected cost J^π can be expressed as a recursive function using a state-action value function, called Q -function:

$$Q^\pi(\mathbf{x}_t, \mathbf{u}_t) = \mathbb{E}[\rho(\mathbf{x}_t, \mathbf{u}_t, \omega) + \gamma Q^\pi(\mathbf{x}_{t+1}, \pi(\mathbf{x}_{t+1}))]. \quad (3)$$

Here, ω represents the stochasticity in the transition function (f) and can be attributed to exogenous factors. The discount factor is represented as γ .

For our work, we focus on the deep- Q network (DQN) algorithm [30], where the Q -function is iteratively estimated using a deep neural network as a function approximator. The neural network-based Q -function (parameterized as \hat{Q}_θ) is trained on a batch of data (\mathcal{F}) with the following loss term:

$$\mathcal{L} = \mathbb{E} \left[\left(\hat{Q}_\theta(\mathbf{x}_t, \mathbf{u}_t) - \left(c_t + \min_{\mathbf{u} \in \mathbf{U}} \hat{Q}_\theta(\mathbf{x}_{t+1}, \mathbf{u}) \right) \right)^2 \right], \quad (4)$$

where, $c_t = \rho(\mathbf{x}_t, \mathbf{u}_t, \omega)$ is the observed cost value during the state transition from \mathbf{x}_t to \mathbf{x}_{t+1} and the expectation is over all elements of \mathcal{F} . For more details related to the DQN algorithm, we refer the readers to [30]. Note that our proposed method is agnostic to the choice of the RL algorithm and can be easily extended to other RL algorithms as well.

4 METHODOLOGY

This section details our proposed approach. We first mathematically formulate the differentiable decision tree architecture, followed by the policy distillation process.

4.1 Differentiable Decision Trees (DDTs)

Differentiable decision trees or soft decision trees are a variant of ordinary decision trees, introduced in prior works such as [11, 19]. We follow the work presented in [35], where a DDT is formulated as a directed, acyclic graph consisting of nodes and edges. There are two types of nodes in a DDT: (1) decision nodes, characterized by a feature selection weights (β) and a threshold (ϕ); and (2) leaf nodes

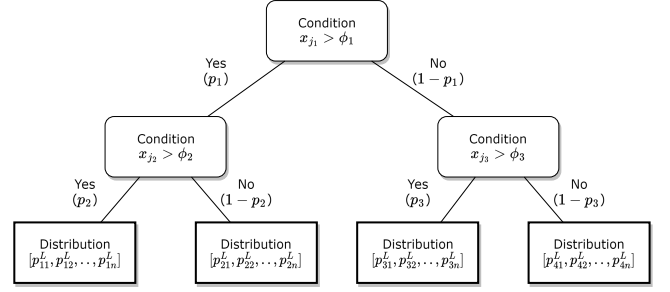


Figure 2: Illustration of a DDT of depth 2. The rounded boxes depict the decision nodes and the rectangles depict leaf nodes. All p_i represent the path probabilities and p_{jk}^L denotes the leaf probability distributions (with $n = |\mathbf{U}|$ i.e., the total number of actions).

containing a weight vector (\mathbf{w}^L) to express the probability distribution. While ordinary decision trees have decision nodes represented using a boolean function, DDTs implement a ‘soft’ decision using the *sigmoid* function (represented as σ). Consequently, each path (or edge) going out of a decision node carries a probability value that is based on the condition evaluated at that decision.

4.1.1 Decision Node. A decision node (represented as rounded boxes in Fig. 2) is modeled as:

$$p = \sigma(\beta \mathbf{x} - \phi) \quad (5a)$$

$$p^{\text{left}} = p \quad (5b)$$

$$p^{\text{right}} = 1 - p \quad (5c)$$

Here, β and ϕ are trainable parameters representing the feature selection weight and the cut thresholds respectively. Each decision node evaluates a condition based on the selected feature and cut threshold and gives path probabilities for going left (the condition is likely to be True) and going right (the condition is likely to be False).

4.1.2 Leaf Nodes. A leaf node l contains a weight vector (\mathbf{w}_l^L) that leads to an output probability distribution modeled using a *SoftMax* function. In our case, each leaf output is the probability distribution over all actions in the action space (\mathbf{U}), however, this can be extended to estimate exact values (for continuous actions) as well. For this leaf node, the probability for each action $u_m \in \mathbf{U}$ is calculated using Eq. (6)

$$p_{lm}^L = \frac{e^{-w_m}}{\sum_{\kappa=1}^{|\mathbf{U}|} e^{-w_\kappa}} \quad \forall m \in \{1, 2, \dots, |\mathbf{U}|\} \quad (6)$$

4.1.3 Creating a DDT. Eq. (5) and Eq. (6) are combined to implement a DDT of required depth. To illustrate this, we now present the formulation of a DDT of depth 2 (as shown in Fig. 2). Such a DDT contains 3 decision nodes and 4 leaf nodes. For each decision node, we have feature selection vectors ($\beta_1, \beta_2, \beta_3$) and cut-thresholds (ϕ_1, ϕ_2, ϕ_3); each leaf node contains weight vectors (\mathbf{w}_k^L). The tree is built based on algorithm 1. This formulation is used to perform a forward pass of the DDT and train the parameters using gradient

Algorithm 1 Depth 2 DDT Formulation

```

1: Initialize:  $\beta_i, \phi, \mathbf{w}_k^L$ , where  $i = \{1, 2, 3\}$  (decision nodes) and
    $k = \{1, 2, 3, 4\}$  (leaf nodes)
2: Input: State  $\mathbf{x}$ 
3: for all  $i$  do
4:   Feature Selection:  $x_j = \beta_i \mathbf{x}$ 
5:   Evaluate Condition:  $p_i = \sigma(x_j - \phi_i)$ 
6: end for
7: Calculating Path Probabilities:  $\mathbf{p} = \begin{bmatrix} p_1 & 0 \\ 0 & 1-p_1 \end{bmatrix} \cdot \begin{bmatrix} p_2 & 1-p_2 \\ p_3 & 1-p_3 \end{bmatrix}$ 
8: for all  $k$  do
9:   Calculate Leaf Probabilities:  $\mathbf{p}_k^L = \{p_{k1}^L, p_{k2}^L, \dots, p_{kn}^L\}$  based
     on Eq. (6)
10: end for
11: Output:  $o = \mathbf{p}[1, 1]\mathbf{p}_1^L + \mathbf{p}[1, 2]\mathbf{p}_2^L + \mathbf{p}[2, 1]\mathbf{p}_3^L + \mathbf{p}[2, 2]\mathbf{p}_4^L$ 

```

descent. At inference, each node is converted from the ‘soft’ version into a crisp node, resembling an ordinary decision tree. This includes reducing all feature selection parameters (β) into one-hot representations (using *argmax*) and converting all probabilities into ‘crisp’, boolean values. Note that this method of creating a DDT decomposes all computations into differentiable computations and allows to parallelize them. Additionally, while a DDT of any depth can be implemented based on Eq. (5), Eq. (6), for this work we restrict the scope to trees of depth 2 and 3. This choice is primarily driven by the ease of explainability for such (shallow) trees.

4.2 Policy Distillation

Distillation is a method for transferring knowledge from a teacher model T to a student model S [34]. In the context of reinforcement learning, this refers to transferring knowledge related to a control policy from a trained teacher agent (π^T) to a student agent (π^S). Typically, this leads to a classification problem where targets are obtained using the outputs of the trained agent.

We follow the approach presented in [34], where a DQN-based teacher agent is trained first and then using a batch of observations (\mathcal{F}), a student policy is distilled based on the teacher agent. First, the trained teacher agent is used to create a new batch of training data of the form $\mathcal{D} = \{\mathbf{x}_i, \mathbf{q}_i\}_{i=1}^{|\mathcal{F}|}$. Here, \mathbf{q}_i is the vector corresponding to Q -values for all actions for a state $\mathbf{x}_i \in \mathcal{F}$, obtained using the teacher agent (i.e., $\mathbf{q}_i = \{Q^T(\mathbf{x}_i, u_i) \mid \forall u_i \in \mathbf{U}\}$). Following this, the student agent is trained to mimic this distribution using Kullback-Leibler (KL) divergence with temperature (τ) as presented in Eq. (7).

$$\mathcal{L}_{\theta_s} = \text{softmax}\left(\frac{\mathbf{q}_i}{\tau}\right) \cdot \ln\left(\frac{\text{softmax}\left(\frac{\mathbf{q}_i}{\tau}\right)}{\text{softmax}\left(\frac{\mathbf{q}_i^S}{\tau}\right)}\right) \quad (7)$$

Note that \mathbf{q}_i^S is the output of the student model parameterized by θ_s . The temperature τ is used to adjust the ‘smoothness’ of the Q -function distribution.

Algorithm 2 Training algorithm for our proposed method

```

1: Initialize: Teacher agent  $T$ , DDT student  $S$ , buffer  $\mathcal{F}$ .
2: Train Teacher:
   Use  $\mathcal{F}$  and Eq. (4) to train teacher i.e., obtain  $\pi^T$  and  $Q^T$ 
3: Create Distillation Batch:
   Distillation batch  $\mathcal{D} = \{\mathbf{x}_i, \mathbf{q}_i\}_{i=1}^{|\mathcal{F}|}$ 
   where  $\mathbf{q}_i = \{Q^T(\mathbf{x}_i, u_i) \mid \forall u_i \in \mathbf{U}\}$ 
4: Train Student DDT:
   Use  $\mathcal{D}$  and Eq. (7) to train the student ( $\pi^S$ ) using gradient
   descent.

```

4.3 Our Approach

For our work, we assume a teacher agent (policy π^T and Q -function Q^T) as a standard DQN agent, and the student agent (π^S) consists of the DDT architecture. First, the teacher agent is trained independently using DQN, to obtain a control policy. Following this, the trained teacher is used to create target distributions using data collected from previous interactions with the environment. This data is then used to train the student DDT-based agent. Algorithm 2 outlines the training procedure for our proposed approach.

5 EXPERIMENT SETUP

We validate our proposed approach on a home energy management scenario using a battery as the source of flexibility. This section presents the simulation environment and details the training and experimental scenarios used.

5.1 Simulator Setup

We use a Python-based simulation environment to validate and compare our proposed approach with standard RL-based controllers. This simulator is derived from a real-world Belgian household with rooftop solar PV and is modeled based on Eq. (1). Real demand and solar PV profiles are used along with the battery model presented in Eq. (1e). Additionally, we use hourly, real-world BELPEX prices as consumption prices (λ_t^{con}) and a capacity tariff structure based on [37]. The battery parameters are detailed in Appendix A.1. Further, we assume the injection price (λ_t^{inj}) is 25% of the consumption price i.e., ($\lambda_t^{\text{inj}} = 0.25 \lambda_t^{\text{con}}$).

5.2 Training Setup

The training is divided into two parts: (i) training the teacher agent; and (ii) policy distillation to train student agent. For the teacher agent, we follow the standard DQN implementation and use an ϵ -greedy training strategy to train the DQN-based teacher agent. Following this, the buffer generated by the DQN-based teacher is used to create the distillation dataset (\mathcal{D}). The student agent is then trained using this dataset. To improve the stability of the training process, we set the temperature (τ) from Eq. (7) to 0.03 to obtain a sharp Q -function distribution. We list all the hyperparameters used in Appendix A.2. For each agent (teacher and student) we perform 5 seeded runs and compare the mean values over the 5 runs.

5.3 Experiment Scenarios

The primary goal of this work is to present a novel approach for obtaining explainable, RL-based policies. We investigate the HEMS scenario described in §3.1, where one intentionally simplified scenario is used to effectively assess the explainability of our proposed approach. We specifically investigate two key scenarios:

5.3.1 Scenario 1: Performance Comparison. In this scenario, we investigate the performance of our proposed approach and evaluate whether our method can achieve satisfactory performance compared to standard DQN agents and baseline rule-based controller. For this, we consider a realistic HEMS scenario and use real-world data for load profiles, solar PV, and prices as described in §5.1. The performance is quantified as the total cost for a day, comprising both energy and capacity costs. As baselines, we use the teacher agents as the upper bound of performance and a rule-based control (RBC) policy as the lower bound. The RBC policy is designed similar to the typical built-in control policy of home batteries and aims to maximize self-consumption. We consider two key variants of price profiles: (i) an artificial, square wave price profile (resembling day-night tariff); and (ii) an actual real-world day ahead price profile. The artificial price profile is a simplified scenario with clear peaks and valleys in the price to provide unambiguous opportunities for energy arbitrage.

5.3.2 Scenario 2: Explainability Assessment. To further assess the explainability of our method, we consider a simplified scenario where we exclude solar PV from the system and reduce the state features to 3 components i.e., battery state-of-charge, price, and demand. This simplification enables clear visualization of the learned policies, contrasting them with standard DQN policies to qualitatively investigate the explainability of our proposed method.¹

6 RESULTS

This section presents the results obtained for the different scenarios discussed in §5.3.

6.1 Performance Evaluation

The performance of our proposed approach using DDTs of depth 2 and 3 is presented in Fig. 3. We note two key observations: (i) both DDT agents clearly outperform the baseline RBC controller; (ii) while the DQN-based teacher performs better than the DDTs, the performance difference (mean) is quite small ($\sim 5\%$). This indicates that our proposed approach can learn satisfactory control policies that outperform an RBC included with standard batteries. Additionally, the DDTs can mimic the teacher agents well and sustain minimal deterioration in performance.

While the overall performance is satisfactory, Fig. 3 indicates some (training) stability issues with the DDT-based controllers. This is particularly apparent in Fig. 3a, where both DDTs demonstrate a strong performance for 3 of the runs, while the other two instances do not fare as well. This problem can be attributed to the training process, where changes in ‘upstream’ or hierarchically higher features could have a disproportionate impact on the output

¹Quantitatively assessing the explainability of AI methods remains an open question with most prior works relying on either qualitative methods or user studies for assessment [33].

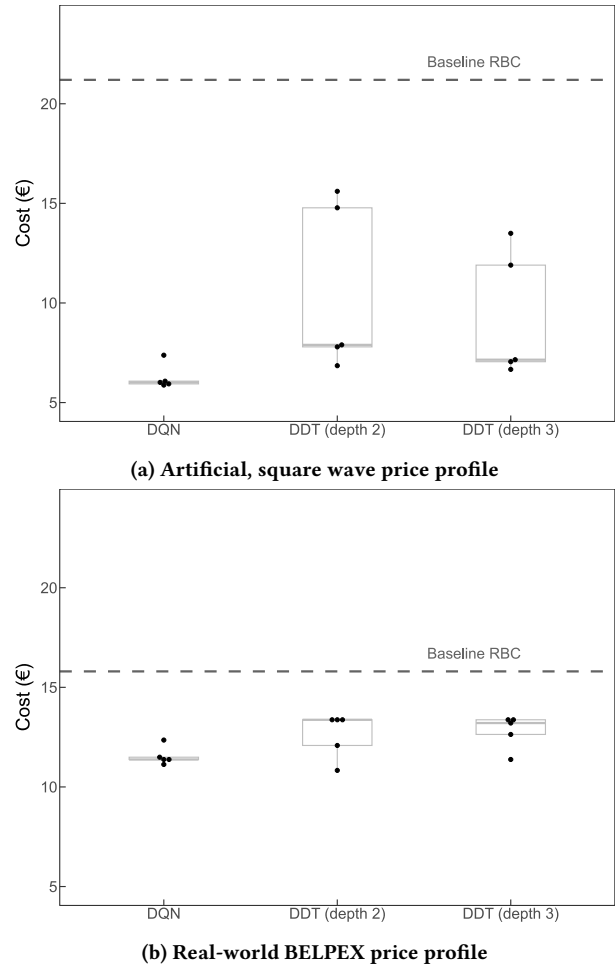


Figure 3: Performance of DDT-based students as a HEMS on different price scenarios. The dots represent the actual performance of individual models and the box plots show the aggregate performance. The student agents are benchmarked using teacher agent ‘DQN’ and a RBC.

distributions. This needs to be investigated further and will be part of future work as discussed in §7.1.

Furthermore, examples of learned DDTs of depth 2 are presented in Fig. 4. Note that these DDTs are randomly initialized and over the course of training learn the feature selection (e.g., choosing ‘demand’ or ‘solar PV’ as the feature for the first decision node) and the respective cut thresholds via gradient descent. We observed that both DDTs are straightforward to understand, easily ‘explaining’ how the controller takes an action. Additionally, the actions taken are intuitive and follow human intuition – e.g., in Fig. 4b, the controller decides to take a charging action only when solar PV generation is high (greater than 0.47) while demand is low (less than 0.37). Likewise, in Fig. 4a, the controller discharges with maximum power when both price and demand are high while only discharging by half the power when price is high but demand is

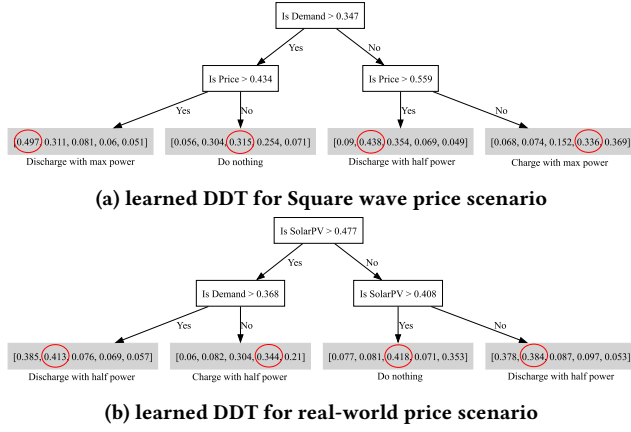


Figure 4: Visual representation of learned decision trees of depth 2 for both price scenarios. The decision nodes are depicted with unshaded boxes and contain the learned features and the threshold values. The leaf nodes are depicted by grey boxes and contain the learned distribution. The annotations highlight the actions related to each leaf node.

low, showing that the learned policy adjusts its decision based on the current as well as expected future demand.

We conclude that the results presented in Fig. 3 and Fig. 4 validate the performance and explainability of our proposed DDT-based approach and show that the DDTs learn a simple, easy-to-explain policy and achieve satisfactory control performance.

6.2 Explainability Comparison

While §6.1 investigated the control performance of our method, we now examine the explainability of the obtained policies. As described in §5.3, we consider a reduced problem where a house without solar PV is exposed to an artificial square wave price profile. Despite being hypothetical, this scenario reduces the dimensionality of the state space (now reduced to battery state-of-charge, price and non-flexible demand) and allows us to examine and compare the explainability of the learned DDT policy with that of the DQN policy. While the visual representation of the policy as shown in Fig. 4 is useful, we cannot directly compare it with the teacher policy (which is a neural network). Consequently, we make use of policy heatmaps to visualize different control policies [33]. Figure 5 illustrates such heatmaps comparing the teacher (DQN) policy with depth 2 and depth 3 DDT policies. These heatmaps are generated by evaluating the controller’s policy on all possible states (in a fixed subset of the state space) and provide an overview of how an agent would react for different states.

Based on Fig. 5, we observe that the DDT heatmaps (top rows of the figure) are consistent, straightforward, and can be easily decomposed into a few rules based on demand, price or state-of-charge. Contrary to this, the DQN-based policy is complex and often non-intuitive in terms of actions taken in specific regions. E.g., the DQN policy in the low-demand region prefers to discharge the battery even in the regions where the price is quite low (e.g., regions where price is less than 0.25 and the state of charge is greater than 0.75). Such behavior is counter-intuitive and difficult to understand

Table 1: Comparison of DQN and DDTs based on computational metrics

Algorithm	Number of Parameters	Storage Size
DQN (teacher agent)	4.8k	22KB
DDT – depth 2	10 (38)	4KB
DDT – depth 3	22 (82)	7KB

even for experts, not to mention everyday homeowners (who will actually use such a system). This further highlights the increased explainability achieved using our proposed approach.

6.3 Compute performance

Besides explainability, the proposed DDT-based method is computationally light and easy to deploy on any edge device given that it reduces the control policy into a limited set of if-then-else rules. As a comparison, Table 1 lists the number of parameters used and the storage footprint of the teacher agents and the distilled DDTs used in §6.1. For DDTs, the number of parameters used during training are represented inside parenthesis along with the parameters used during inference. Unlike DQN which uses the same set of parameters during training and inference, DDTs require fewer parameters for inference – e.g., at any decision node, the feature selection parameters can be reduced to a single parameter representing the selected feature. From this table, it can be clearly observed that the proposed DDTs have a significantly smaller compute footprint due to the reduced number of parameters, leading to trained models which are about 200 times smaller than the teacher DQN agents. To conclude, the comparison in Table 1 further underscores the potential for deploying such controllers in real-world scenarios.

7 CONCLUSION

Through this work, we introduced a novel method for obtaining explainable RL-based control policies using differentiable decision trees and policy distillation. The key idea of our work is to distill knowledge from a standard RL-based controller into a simple, easy-to-explain decision tree architecture by purely relying on data. For this, we use differentiable decision trees in a policy distillation setup, training the decision trees using a standard (pre-trained) RL-based controller and gradient descent. The policy distillation step allows extracting knowledge from an RL-based controller, while the differentiable decision tree architecture constrains the policy to be simple and explainable at all times.

We validated our method on a battery-based home energy management problem and investigated the control performance and explainability of our proposed approach. As presented in §6, our proposed approach learns a control policy that performs comparable to the teacher DQN agent, while being simple (i.e., ~ 200 times reduction in number of parameters) and easy-to-explain. Furthermore, the performance of our DDT-based controllers surpasses the performance of commonly observed RBC, performing ~ 20 – 25% better than the RBC.

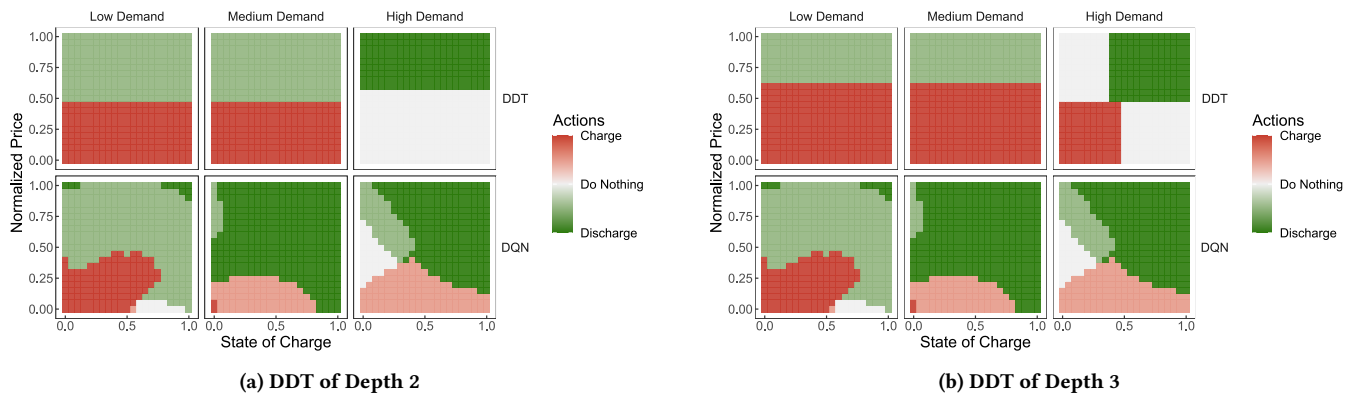


Figure 5: Visualizing the trained policy of DDT and DQN-based agent on a simplified HEMS scenario. The heatmaps show the actions chosen by the agents for different values of state-of-charge and price across different demand regions. The bottom row depicts the DQN policy and the top rows show the policy of our proposed DDT-based controllers

7.1 Limitations and Future Work

As discussed in §1, the goal of this work was to introduce this novel method and highlight its potential for future applications in the energy domain. In support of this objective, we identify some limitations within the current work and outline areas for future investigations.

The initial consideration pertains to the problem formulation discussed in §3.1, which we will further expand to include thermal models and joint optimization with comfort constraints. While the current problem mimics a real-world house in the present times, future application scenarios will require more elaborate HEMS that can optimize cost by leveraging flexibility from different sources including building thermal mass, batteries, and electric vehicles. To efficiently deal with such complex scenarios, our future work will explore two main aspects: (i) extending the policy distillation set-up to multi-agent RL settings, where simple, shallow DDTs can be trained per flexibility asset; and (ii) domain knowledge induced feature engineering (using previous works such as [13]) to compress information and allow the use of shallow DDTs. While large DDTs can be trained for such complex scenarios, we intend to focus on “shallow” DDTs that are intuitively easier to explain (or more explainable) as compared to “deep” trees.

Besides this, another limitation of our current approach is the occasional instability in the training process related to the DDTs. As noted in §6, this training instability could be attributed to the tree structure of the DDT with features hierarchically higher up in the tree significantly affecting the outputs. This needs to be investigated further to identify possible solutions to stabilize the learning process. This includes effective regularization strategies, warm starting, or constraining the decisions being learned. The latter seems particularly useful for DDTs of higher depth, where some decisions are conflicting or redundant (as shown in Appendix B).

The third area that needs to be addressed further is the deployment of such an algorithm in real-world scenarios and performing a user trial to further validate the explainability of our method. While non-trivial, such a pilot study is needed to investigate the acceptance of such a HEMS as well as the challenges associated with maintaining such a system. This will further allow us to investigate

more advanced approaches such as human-in-loop training and intervention strategies to maximize the decision tree architecture and develop a robust, data-driven controller that can be widely deployed across houses.

ACKNOWLEDGMENTS

This research has received funding from the Horizon 2020 Project RENERgetic (grant no. 957845) and Energy Transition Fund’s FlexMyHeat project.

REFERENCES

- [1] Aya A Amer, Khaled Shaban, and Ahmed M Massoud. 2022. DRL-HEMS: Deep reinforcement learning agent for demand response in home energy management systems considering customers and operators perspectives. *IEEE Transactions on Smart Grid* 14, 1 (2022), 239–250.
- [2] Mohak Bhardwaj, Sanjiban Choudhury, and Byron Boots. 2020. Blending mpc & value function approximation for efficient reinforcement learning. *arXiv preprint arXiv:2012.05909* (2020).
- [3] Wenqi Cai, Shambhuraj Sawant, Dirk Reinhardt, Soroush Rastegarpour, and Sebastien Gros. 2023. A Learning-Based Model Predictive Control Strategy for Home Energy Management Systems. *IEEE Access* (2023).
- [4] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. 2021. Distilling knowledge via knowledge review. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5008–5017.
- [5] Xin Chen, Guannan Qu, Yujie Tang, Steven Low, and Na Li. 2022. Reinforcement learning for selective key applications in power systems: Recent advances and future challenges. *IEEE Transactions on Smart Grid* 13, 4 (2022), 2935–2958.
- [6] Youri Coppens, Kyriakos Efthymiadis, Tom Lenaerts, Ann Nowé, Tim Miller, Rosina Weber, and Daniele Magazzeni. 2019. Distilling deep reinforcement learning policies in soft decision trees. In *Proceedings of the IJCAI 2019 workshop on explainable artificial intelligence*. 1–6.
- [7] Hongyuan Ding, Yan Xu, Benjamin Chew Si Hao, Qiaoqiao Li, and Antonis Lentzakis. 2022. A safe reinforcement learning approach for multi-energy management of smart home. *Electric Power Systems Research* 210 (2022), 108120.
- [8] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L Vrabie, et al. 2020. All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50 (2020), 190–232.
- [9] eurostat. 2023. *Energy consumption in households*. Retrieved January 28, 2024 from https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_consumption_in_households
- [10] Laura Fiorini and Marco Aiello. 2020. Predictive multi-objective scheduling with dynamic prices and marginal co2-emission intensities. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems*. 196–207.
- [11] Nicholas Frosst and Geoffrey Hinton. 2017. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784* (2017).
- [12] Arne Gevaert, Jonathan Peck, and Yvan Saeys. 2022. Distilling Deep RL Models Into Interpretable Neuro-Fuzzy Systems. In *2022 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*. IEEE, 1–8.

- [13] Gargya Gokhale, Bert Claessens, and Chris Develder. 2022. Physics informed neural networks for control oriented thermal modeling of buildings. *Applied Energy* 314 (2022), 118852. <https://doi.org/10.1016/j.apenergy.2022.118852>
- [14] ILR Gomes, MG Ruano, and AE Ruano. 2023. MILP-based model predictive control for home energy management systems: A real case study in Algarve, Portugal. *Energy and Buildings* 281 (2023), 112774.
- [15] Binghui Han, Younes Zahraoui, Marizan Mubin, Saad Mekhilef, Mehdi Seyedmohmoudian, and Alex Stojcevski. 2023. Home Energy Management Systems: A Review of the Concept, Architecture, and Scheduling Strategies. *IEEE Access* (2023).
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [17] IEA. 2023. *Energy Technology Perspectives 2023, IEA, Paris*. Retrieved January 28, 2024 from <https://www.iea.org/reports/energy-technology-perspectives-2023>
- [18] IRENA. 2023. *World Energy Transitions Outlook 2023: 1.5°C Pathway, Volume 2, International Renewable Energy Agency, Abu Dhabi*. Retrieved January 28, 2024 from <http://crma.stanford.edu/~jos/bayes/bayes.html>
- [19] Ozan Irsoy, Olcay Taner Yildiz, and Ethem Alpaydm. 2012. Soft decision trees. In *Proceedings of the 21st international conference on pattern recognition (ICPR2012)*. IEEE, 1819–1822.
- [20] Rune Grønborg Junker, Armin Ghasem Azar, Rui Amaral Lopes, Karen Byskov Lindberg, Glenn Reynders, Rishi Relan, and Henrik Madsen. 2018. Characterizing the energy flexibility of buildings and districts. *Applied energy* 225 (2018), 175–182.
- [21] Hareesh Kumar, Priyanka Mary Mammen, and Krithi Ramamritham. 2019. Explainable AI: Deep reinforcement learning agents for residential demand side cost savings in smart grids. *arXiv preprint arXiv:1910.08719* (2019).
- [22] Joaquim Leitao, Paulo Gil, Bernardete Ribeiro, and Alberto Cardoso. 2020. A survey on home energy management. *IEEE Access* 8 (2020), 5699–5722.
- [23] Timothy P Lillcrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* (2015).
- [24] Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis. 2020. Explainable ai: A review of machine learning interpretability methods. *Entropy* 23, 1 (2020), 18.
- [25] Paulo Lissa, Conor Deane, Michael Schukat, Federico Seri, Marcus Keane, and Enda Barrett. 2021. Deep reinforcement learning for home energy management system control. *Energy and AI* 3 (2021), 100043.
- [26] Haoran Luo, Fan Cheng, Heng Yu, and Yuqi Yi. 2021. SDTR: Soft decision tree regressor for tabular data. *IEEE Access* 9 (2021), 55999–56011.
- [27] R Machlev, L Heistrene, M Perl, KY Levy, J Belikov, S Mannor, and Y Levron. 2022. Explainable Artificial Intelligence (XAI) techniques for energy and power systems: Review, challenges and opportunities. *Energy and AI* 9 (2022), 100169.
- [28] Bandana Mahapatra and Anand Nayyar. 2022. Home energy management system (HEMS): Concept, architecture, infrastructure, challenges and energy management schemes. *Energy Systems* 13, 3 (2022), 643–669.
- [29] Dang Minh, H Xiang Wang, Y Fen Li, and Tan N Nguyen. 2022. Explainable artificial intelligence: a comprehensive review. *Artificial Intelligence Review* (2022), 1–66.
- [30] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [31] Prateek Munankarmi, Hongyu Wu, Annabelle Pratt, Monte Lunacek, Sivasathya Pradha Balamurugan, and Paul Spitsen. 2022. Home energy management system for price-responsive operation of consumer technologies under an export rate. *IEEE Access* 10 (2022), 50087–50099.
- [32] Zoltan Nagy, Gregor Henze, Sourav Dey, Javier Arroyo, Lieve Helsen, Xiangyu Zhang, Bingqing Chen, Kadir Amasyali, Kuldeep Kurte, Ahmed Zamzam, et al. 2023. Ten questions concerning reinforcement learning for building energy management. *Building and Environment* (2023), 110435.
- [33] Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt, Jörg Schlöterer, Maurice van Keulen, and Christin Seifert. 2023. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable ai. *Comput. Surveys* 55, 13s (2023), 1–42.
- [34] Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2015. Policy distillation. *arXiv preprint arXiv:1511.06295* (2015).
- [35] Andrew Silva, Matthew Gombolay, Taylor Killian, Ivan Jimenez, and Sung-Hyun Son. 2020. Optimization methods for interpretable differentiable decision trees applied to reinforcement learning. In *International conference on artificial intelligence and statistics*. PMLR, 1855–1865.
- [36] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [37] VREG. 2023. *What is the capacity rate and how is it calculated? (2023) VREG. Available at: (Accessed: 29 January 2024)*. Retrieved January 29, 2024 from <https://www.vreg.be/nl/wat-zijn-de-nieuwe-nettarieven-en-hoe-woorden-ze-bereken>
- [38] Zhe Wang and Tianzhen Hong. 2020. Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy* 269 (2020), 115036.
- [39] Luolin Xiong, Yang Tang, Chensheng Liu, Shuai Mao, Ke Meng, Zhaoyang Dong, and Feng Qian. 2023. Meta-Reinforcement Learning-Based Transferable Scheduling Strategy for Energy Management. *IEEE Transactions on Circuits and Systems I: Regular Papers* 70, 4 (2023), 1685–1695.
- [40] Lingxiang Yun, Di Wang, and Lin Li. 2023. Explainable multi-agent deep reinforcement learning for real-time demand response towards sustainable manufacturing. *Applied Energy* 347 (2023), 121324.
- [41] Eva Záčková, Zdeněk Váňa, and Jiří Cigler. 2014. Towards the real-life implementation of MPC for an office building: Identification issues. *Applied Energy* 135 (2014), 53–62.
- [42] Usman Zafar, Sertac Bayhan, and Antonio Sanfilippo. 2020. Home energy management system concepts, configurations, and technologies for the smart grid. *IEEE access* 8 (2020), 119271–119286.
- [43] Ke Zhang, Jun Zhang, Pei-Dong Xu, Tianlu Gao, and David Wenzhong Gao. 2021. Explainable AI in deep reinforcement learning models for power system emergency control. *IEEE Transactions on Computational Social Systems* 9, 2 (2021), 419–427.
- [44] Xiangyu Zhang, Xin Jin, Charles Tripp, David J Biagioni, Peter Graf, and Huaiguang Jiang. 2020. Transferable reinforcement learning for smart homes. In *Proceedings of the 1st international workshop on reinforcement learning for energy management in buildings & cities*. 43–47.

A HYPERPARAMETERS

A.1 Home Energy Management Environment

For the HEMS simulator described in §5.1, we consider a nominal power rating (P_t^{agg}) of 4kW, prompting the agents to perform peak shaving operations to avoid exceeding this value. The agents can do this by controlling the battery (modeled based on Eq. (1e)), the details of which are tabulated in Table 2. The action space represents the charging or discharging signal that is given to the battery (and not the actual power).

Table 2: Parameters related to the Battery model used in the Home Energy Management Simulator

Parameter	Value
Max Capacity	10 kWh
Max Power	4 kW
Efficiency (round trip)	0.9
Action Space	{-1, -0.5, 0, 0.5, 1}

A.2 DQN-based Teacher Agents

The hyperparameters used for the DQN-based teacher agents are listed in Table 3. Additionally, during the distillation process, the temperature τ is set to 0.03.

A.3 Baseline Rule-based Controller (RBC)

We compare the performance of our proposed DDTs with the corresponding teacher agents and a baseline RBC. This baseline is modeled based on built-in controllers that are available with commercially available batteries and are generally designed to maximize the self-consumption of solar PV. Such an RBC is modeled as:

$$u_i = \begin{cases} -1 & : P_t^{\text{agg}} \leq -P^{\text{max}} \\ 1 & : P_t^{\text{agg}} \geq P^{\text{max}} \\ \frac{P_t}{P^{\text{max}}} & : \text{otherwise} \end{cases}, \quad (8)$$

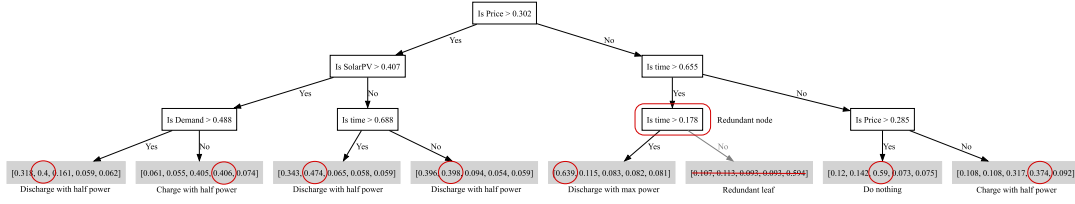


Figure 6: Example of a learned DDT for depth 3 for the real-world BELPEX price scenario.

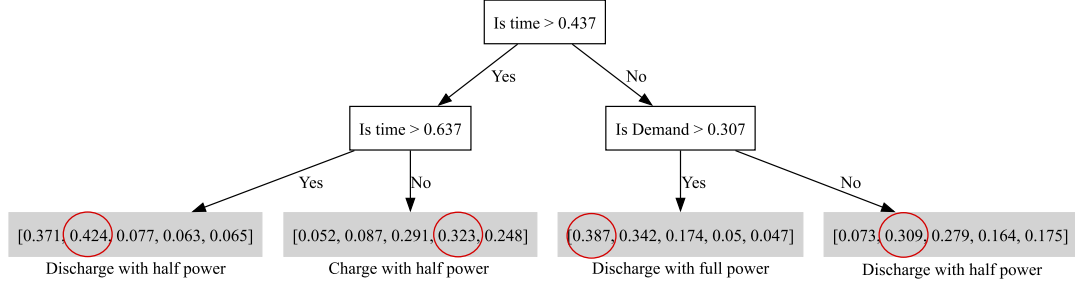


Figure 7: Example of a learned DDT for depth 2 for the real-world BELPEX price scenario. This particular policy learns an unusual set of rules, yet performs well.

Table 3: Hyperparameters for Q-network (DQN)

Parameter	Value
Optimizer	Adam
Learning Rate	0.001
Activation Function	ReLU
Mini Batch Size	1000
Output Size	5 (= U)
Hidden Layers	[64, 64]
Target Update (τ)	0.1
Buffer Size	5000

exploiting the underlying correlation between time, price, and solar PV generation. Hence, while the rules learned are unusual, this policy still shows good performance. Such unusual rules are related to high correlation between various features and to improve the rules (making them more intuitive) we need to carefully select the relevant features from the perspective of information content and explainability.

where $P_t = P_t^{\text{con}} + P_t^{\text{pv}}$ represents the power required or left over after self-consumption.

B ADDITIONAL EXAMPLE OF LEARNED DIFFERENTIABLE DECISION TREES

This section provides some more examples related to the learned DDT policies. Figure 6 depicts a depth 3 DDT student. The shown policy (decision tree) takes intuitive actions, e.g., charging the battery when the price is low or when solar PV is high. However, as compared to the depth 2 DDT (shown in Fig. 4b), this policy is a bit more difficult to understand. Further, as annotated in the figure, some decision nodes learn redundant rules, which may lead to some training instability. Besides this, for some instances, the DDTs may learn rules that are not obvious or human intuitive, an example of such a scenario is presented in Fig. 7. As shown in the figure, the policy bases its decisions on time as a feature, possibly