

# Control Policy Correction Framework for Reinforcement Learning-based Energy Arbitrage Strategies

Seyed Soroush Karimi Madahi  
seyedsoroush.karimimadahi@ugent.be  
IDLab, Ghent University – imec  
Ghent, Belgium

Gargya Gokhale  
IDLab, Ghent University – imec  
Ghent, Belgium

Marie-Sophie Verwee  
IDLab, Ghent University – imec  
Ghent, Belgium

Bert Claessens  
Beebop  
Belgium

Chris Develder  
IDLab, Ghent University – imec  
Ghent, Belgium

## ABSTRACT

A continuous rise in the penetration of renewable energy sources, along with the use of the single imbalance pricing, provides a new opportunity for balance responsible parties to reduce their cost through energy arbitrage in the imbalance settlement mechanism. Model-free reinforcement learning (RL) methods are an appropriate choice for solving the energy arbitrage problem due to their outstanding performance in solving complex stochastic sequential problems. However, RL is rarely deployed in real-world applications since its learned policy does not necessarily guarantee safety during the execution phase. In this paper, we propose a new RL-based control framework for batteries to obtain a safe energy arbitrage strategy in the imbalance settlement mechanism. In our proposed control framework, the agent initially aims to optimize the arbitrage revenue. Subsequently, in the post-processing step, we correct (constrain) the learned policy following a knowledge distillation process based on properties that follow human intuition. Our post-processing step is a generic method and is not restricted to the energy arbitrage domain. We use the Belgian imbalance price of 2023 to evaluate the performance of our proposed framework. Furthermore, we deploy our proposed control framework on a real battery to show its capability in the real world.

## CCS CONCEPTS

• **Theory of computation** → **Reinforcement learning**; *Apprenticeship learning*; • **Hardware** → **Smart grid**; **Batteries**.

## KEYWORDS

Battery energy storage systems, distributional reinforcement learning, energy arbitrage, interpretable reinforcement learning, knowledge distillation, safe reinforcement learning

### ACM Reference Format:

Seyed Soroush Karimi Madahi, Gargya Gokhale, Marie-Sophie Verwee, Bert Claessens, and Chris Develder. 2024. Control Policy Correction Framework

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](mailto:permissions@acm.org).

*E-Energy '24, June 4–7, 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/10.1145/3632775.3661948>

for Reinforcement Learning-based Energy Arbitrage Strategies. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24)*, June 4–7, 2024, Singapore, Singapore. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3632775.3661948>

## 1 INTRODUCTION

Countries are making progress in transitioning toward a decarbonized electricity grid by adopting a larger amount of renewable energy sources (RES). However, the rise in shares of RES leads to an increasing mismatch between generation and consumption, given the dependence of RES generation on weather conditions. This mismatch poses challenges to transmission system operators (TSOs) in maintaining the balance of the grid. TSOs rely on balance responsible parties (BRPs) to assist in keeping supply and demand balance, by penalizing unbalanced BRPs in a periodic (typically 15 min based) imbalance settlement scheme [5]. Europe's electricity balancing guideline (EBGL) stipulates calculating the imbalance cost based on a single imbalance price, implying that both negative and positive imbalances are penalized equally [13]. In this pricing method, BRPs can reduce their cost while assisting TSOs in maintaining grid balance by deviating from their day-ahead nomination. The given imbalance pricing scheme, and the increased need for balancing because of higher RES penetration presents a new opportunity for BRPs to reduce their cost through energy arbitrage in imbalance settlement.

The energy arbitrage problem is a sequential complex one, given the highly uncertain imbalance prices and the nearly real-time decision-making that is required. Most previous research is based on model-based optimization methods to obtain energy arbitrage strategies [7, 20, 32]. These methods formulate the energy arbitrage problem as a nonlinear programming problem which is typically non-convex, meaning that its optimal solution cannot be directly found. For this reason, linearization techniques (such as piecewise linear approximation) are used to transform the nonlinear problem into a linear or mixed-integer convex problem. However, using these linearization techniques might lead to an imprecise approximation or an intractable optimization problem. Stochastic optimization and robust optimization are the most popular model-based optimization methods. However, stochastic optimization has a high computational burden due to numerous scenarios, while a robust optimization solution tends to be extremely cautious [43]. Reinforcement learning (RL) methods can deal with such model-based related problems: model-free RL methods do not require prior knowledge

or a detailed model of the environment. The agent interacts with the environment to capture stochasticity in the environment and learn a (near-)optimal strategy. Also, RL does not have any specific hypothesis concerning the reward function adapting to any non-linear rewards. Furthermore, RL agents directly learn a control policy, without the need for repeatedly solving an optimization (as seen in model-based methods), making them suitable for real-time control. RL methods have achieved state-of-the-art performance on many energy applications such as control of building systems [42], electricity market modeling [17], and voltage control [14].

In spite of the impressive performance of RL in simulations, it is challenging to deploy RL in real-world applications. Indeed, vanilla RL methods cannot guarantee the safety and correctness of the learned policy for unseen states. Safe RL aims to address this by maximizing the cumulative reward while satisfying safety constraints [1]. Most safe RL methods try to constrain the policy during the learning process [3, 10, 40]. While effective, a major drawback of such methods is that, since the agent is trained to satisfy some predefined constraints, the learned constrained policy cannot be scaled to other similar settings: for redeploying the learned policy in other similar settings, the agent needs to be retrained (even if there is a slight change in the constraints). However, retraining RL agents is usually time-consuming and computationally expensive.

To reduce computation at inference time, knowledge distillation in neural networks was first introduced in [18]. The key idea of knowledge distillation is that a larger teacher model transfers its knowledge to a smaller student model to achieve competitive performance and faster inference. Knowledge distillation is commonly used to address the huge computational burden and memory requirement of large models. It is shown that knowledge distillation improves model generalization as the student model is trained using soft targets instead of hard targets [35]. The teacher model tries during its training phase to assign the highest probability to the correct class (or best action) and small probabilities to incorrect classes (or other actions). These relative probabilities of incorrect classes provide more information about how the teacher model tends to generalize. Therefore, training the student model with soft targets transfers this generalization ability of the teacher model efficiently to the student model.

In this paper, we introduce the extension of the standard knowledge distillation process by adding an optimization layer to the student model to correct and constrain the RL policy. Building upon this idea, we propose a new RL-based control framework for batteries to obtain a safe energy arbitrage strategy in the imbalance settlement mechanism. In the proposed control framework, the agent is initially trained to maximize the arbitrage profit (Section 3). Afterwards, in the post-processing step, the aim is to transfer knowledge from the trained agent to the student agent while correcting the policy of the student agent (Section 4). The policy correction of the student agent is based on human intuitive constraints, to make the final policy rational from a human perspective. The main advantage of the proposed framework over previous online safe RL methods lies in its greater flexibility for reusing pretrained RL agents. By applying the post-processing step, BRPs can avoid retraining from scratch and effortlessly reuse pretrained RL agents in accordance with their own defined constraints and preferences.

We employ distributional deep Q learning (DDQN), a state-of-the-art RL method. The main advantage of distributional RL over standard RL is its ability to estimate the complete probability distribution of returns instead of relying on a single value expectation, resulting in superior performance. However, the proposed policy correction process and control framework can be applied to all RL methods. The performance of the proposed control framework is validated using the Belgian imbalance price of 2023. Moreover, a real-time experimental study using a real battery is conducted to better demonstrate the capability of the proposed control framework in the real world (Sections 5 and 6). Our main contributions in this paper are:

- We introduce a new policy correction step that can be applied to any RL method to ensure the correctness and safety of the final policy;
- We propose a distributional RL-based control framework for a battery to obtain a safe energy arbitrage strategy in the imbalance settlement mechanism based on properties that adhere to human intuition;
- We implement the proposed control framework on a real battery to evaluate its performance in the real world.

## 2 BACKGROUND AND RELATED WORK

**Imbalance Settlement.** BRPs are responsible for consistently balancing their individual consumption and generation. However, a deviation from their day-ahead nomination is inevitable because of uncertainties in RES. To correct the system imbalance, a TSO activates reserve capacities offered in the balancing market, and BRPs incur charges from the TSO for their imbalance at the end of the imbalance settlement period (15 mins in most European countries) [19]. This mechanism is called imbalance settlement. The imbalance price is determined based on the reserve volume and direction activated by the TSO [38]. Two main imbalance pricing methodologies are used in different countries: (1) *dual pricing*, where the price is different for positive and negative imbalances, and (2) *single pricing*, where the price is the same for both imbalance directions and is determined by the total system imbalance. As mentioned earlier, the objective of ENTSO-E is to standardize the imbalance settlement mechanism in Europe by adopting the single pricing method for calculating the imbalance price for each 15-minute imbalance settlement period. Hence, we focus on the single pricing methodology. The Belgian imbalance settlement mechanism was chosen as a case study for our research [5].

Energy arbitrage refers to a technique to gain financial profits by buying energy at lower prices and selling it at higher prices. Due to the high volatility in imbalance prices and the need for near real-time sequential decision-making, energy arbitrage is challenging. For this reason, together with the recent change in the imbalance pricing methodology, only few research studies have been conducted on the arbitrage in the imbalance settlement mechanism [7, 20, 23, 32]. A new tailored encoder-decoder architecture was implemented in [7] to generate improved probabilistic predictions of the future system imbalance. Afterward, a bi-level robust optimization problem was solved to optimize the profit of a BESS in the imbalance settlement. A novel risk-aware stochastic model

predictive control (MPC) approach was introduced in [32] to maximize the revenue of BESS in the imbalance settlement mechanism while considering battery degradation costs.

Most of the mentioned studies have solved the arbitrage problem by applying model-based optimization methods. However, the main drawback of these methods is that they need linearization techniques to estimate the nonlinear problem as a linear (or mixed-integer) convex problem that can lead to an inaccurate approximation. To address such problems of model-based optimization methods, we deploy model-free RL methods. Several other research works already studied energy arbitrage using model-free RL. For example, [23] proposes a battery control framework based on distributional RL for a risk-sensitive energy arbitrage in the imbalance settlement mechanism, taking into account a cycle constraint. In [9], an RL-based method was proposed to optimize battery energy arbitrage in the day-ahead market, taking into account an accurate battery degradation model. The authors in [22] present an RL-based battery bidding strategy in the real-time and frequency control ancillary services markets, using a transformer-based temporal feature extractor. In [27], learning-based control algorithms were proposed to obtain an optimal policy for home batteries. Further, [39] used a deep-RL approach to solve the electricity arbitrage problem in the day-ahead market.

Our current work is complementary to all of the above works on RL for energy arbitrage, and particularly to our earlier research in distributional RL [23]. The latter's main objective was to establish which distributional RL algorithm (distributional SAC vs. DQN) performed best, for a battery controller with cycle constraints in imbalance settlement. Yet, our current work specifically looks at the problem that learned RL policies may exhibit erratic behavior, in that they take surprising (if not non-optimal) actions in certain regions of the system state space, especially those that have been only infrequently visited during training. We specifically propose a post-processing step to correct those policies, adopting a distillation setup.

**Safe RL.** In RL literature, the concept of safety is used in opposition to risk, and it is not solely confined to physical damage [15]. In stochastic environments, the learned optimal policy may result in poor performance as the learned policy is not necessarily robust against the rare occurrence of large negative returns. The risk stems from uncertainties in the environment. Some safe RL works focused on *domain knowledge* of the problem to guarantee safety, such as safe exploration [11], designing a safety shield [3], and human interventions [30]. In these works, the safety model is assumed to be known priori. On the other hand, there is another class of safe RL work that focus on *constrained* optimization. In these studies, the safe RL problems are commonly formulated as a constrained Markov decision process (CMDP) (e.g., [1]). Usually, a Lagrange-multiplier method is applied to transform the constrained problem into a non-constrained one (e.g., [40]). In [10], convex constraints are enforced in learned policies by incorporating a differentiable projection layer within a neural network-based policy. The authors in [29] repair the RL policy using expert knowledge by solving a convex optimization problem. They define the optimization problem using a fuzzy model with triangular membership functions to approximate the policy. In [28], an implicit optimization layer was

employed for projecting the actions taken by an RL-based controller to ensure satisfaction of electric vehicle constraints.

The main disadvantage of the previous safe RL works is that they generally constrain the policy during the *training phase*. Thus, it is not straightforward to (re)deploy the learned constrained policy in other similar settings with slight changes in constraints. A new agent needs to be trained for each new setting, even if changes in constraints are minor (e.g., boundary values). In some domains, it is easy to distinguish sub-optimal actions from catastrophic actions – for instance, an autonomous vehicle driving too slowly is sub-optimal, while the car driving into a group of pedestrians is clearly catastrophic. However, in the energy arbitrage domain, the distinction between sub-optimal and catastrophic actions is challenging, since safety and risk preferences may significantly differ among BRPs. This problem becomes more pronounced when a safe agent is trained for energy arbitrage in a highly volatile market, such as the imbalance settlement mechanism. We therefore propose a *new post-processing* step to constrain (correct) a pretrained unconstrained policy. The main advantage of our framework is that it avoids retraining the agent from scratch, a process that is typically time-consuming and computationally expensive. Instead, our proposed framework constrains pretrained agents according to the risk preference of BRPs in a post-processing step. This provides an opportunity to reuse a pretrained agent in similar settings with different constraints. It is worth noting that our proposed post-processing step can be applied to any RL method.

**Knowledge distillation.** It is usually preferred to train cumbersome deep neural networks with a strong regularizer such as dropout because overparameterization increases the generalization performance [16]. However, their computational complexity and slow inference limit their usage in many applications. Knowledge distillation is a technique used to efficiently compress the capabilities of a larger teacher model into a smaller student model. Different forms of knowledge can be transferred from the teacher model to the student model, e.g., logits [18], feature maps [26], and relations between pairs of feature maps [41]. Apart from compaction, knowledge distillation is also used to improve generalization [35], reproducibility [4], data augmentation [21], and defend from adversarial attacks [25]. The teacher model based on its confidence in the ground-truth class rescales gradients of the student model. Also, the teacher model's probability mass on incorrect classes reflects class relationships, offering more guidance to the student model [35]. These effects contribute to improving the generalization of the student model. In this paper, we correct (constrain) the policy when the logits knowledge is transferred from the pretrained unconstrained neural network-based policy to the student model.

### 3 PROBLEM FORMULATION

In this section, the energy arbitrage problem in the imbalance settlement mechanism is formulated as a Markov decision process (MDP) (Section 3.1) and the RL method used for solving the problem is explained in detail (Section 3.2).

#### 3.1 MDP Formulation

An MDP presents stochastic sequential decision-making problems as a mathematical framework. The MDP problem is defined by a

tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma)$ , where  $\mathcal{S}$  represents the state space,  $\mathcal{A}$  denotes the (discrete) action space,  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is the instantaneous reward function,  $\mathcal{P} : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  represents the unknown state transition probability distribution, and a discount factor  $\gamma \in (0, 1]$  [34]. At each time step  $t$ , the environment state  $s_t \in \mathcal{S}$  is the observation of the agent. After taking action  $a_t \in \mathcal{A}$ , the environment provides a reward value  $\mathcal{R}(s_t, a_t)$  for the agent. The state transition probability distribution  $\mathcal{P}(s_{t+1}|s_t, a_t)$  determines the probability of moving to a new state  $s_{t+1} \in \mathcal{S}$ . In our energy arbitrage problem, an instantaneous action taken by the agent affects the overall expected profit of the agent. If the agent prioritizes immediate rewards and discharges the battery completely, it can miss high price periods and make less overall profit. The agent in the energy arbitrage problem makes a decision (i.e., action) at each time step regarding the charging/ discharging of battery. Electricity markets and the grid are considered the environment the agent interacts with. We define the MDP formulation for the energy arbitrage problem in the imbalance settlement mechanism as:

- (i) *State*: The state at each time step (which is considered to be 2 minutes) is given by

$$s_t = (T_{\text{qh}}, qh, mo, SOC_t, \hat{\pi}_t^{\text{imb}}) \quad (1)$$

where  $T_{\text{qh}} \in [0, 14]$  is the minute of the quarter hour,  $qh \in [0, 95]$  denotes the quarter hour of the day,  $mo$  represents the month of the year,  $SOC_t$  is the state of charge (SoC) of battery at time  $t$ . Finally,  $\hat{\pi}_t^{\text{imb}}$  is the *indicative* imbalance price of the current quarter hour  $qh$ . Indeed, the *actual* imbalance price of each quarter hour is only known at the end of the quarter hour. For this reason, our agent can only observe an indicative imbalance price for the current quarter hour. Due to uncertainty in the imbalance price of the current quarter hour, our defined arbitrage problem is stochastic.

- (ii) *Action*: Our action space is discrete, consisting of three possible actions, defined as follows:

$$a_t \in \mathcal{A}, \quad \mathcal{A} = \{-P_{\text{max}}, 0, P_{\text{max}}\} \quad (2)$$

where  $P_{\text{max}}$  is the maximum (dis-)charging power of the battery. A positive action  $a_t$  corresponds to charging the battery, while a negative action means discharging the battery at time  $t$ . The emergence of bang-bang behavior was investigated in continuous control RL by [31]. They showed that RL methods with discrete action space can achieve competitive performance on standard continuous control benchmarks. For this reason, we use a discrete action space in this paper.

- (iii) *Reward*: The agent aims to maximize the profit by purchasing energy at cheap imbalance prices and selling it at expensive imbalance prices. Therefore, the reward function is defined as the negative of the energy cost:

$$r_t = -a_t \pi_{\text{qh}}^{\text{imb}}, \quad (3)$$

where  $\pi_{\text{qh}}^{\text{imb}}$  represents the real imbalance price of the quarter hour in which  $t$  lies.

- (iv) *State transition function*: A state transition probability function  $\mathcal{P}$  describes system dynamics. This probability function is unknown for the agent in our problem due to uncertainties in the imbalance price. More specifically, the probability

distribution of  $\hat{\pi}_{t+1}^{\text{imb}}$  given  $\hat{\pi}_t^{\text{imb}}$  ( $\mathcal{P}(\hat{\pi}_{t+1}^{\text{imb}}|\hat{\pi}_t^{\text{imb}})$ ) is the only source of stochasticity in the agent state and it is independent from the taken action. Nevertheless, the state transition for  $SOC_t$  is explicitly calculated as shown below, since it is influenced by  $a_t$ .

$$SOC_{t+1} = \begin{cases} SOC_{t+1}^{\text{temp}} & : 0 < SOC_{t+1}^{\text{temp}} < 1 \\ 0 & : SOC_{t+1}^{\text{temp}} < 0 \\ 1 & : SOC_{t+1}^{\text{temp}} > 1 \end{cases} \quad (4)$$

$$SOC_{t+1}^{\text{temp}} = SOC_t + (\max(a_t, 0)\eta_{\text{cha}} + \frac{\min(a_t, 0)}{\eta_{\text{dis}}}) \frac{\Delta t}{C_{\text{BESS}}} \quad (5)$$

where  $C_{\text{BESS}}$  is the maximum capacity of the battery, and  $\eta_{\text{cha}}$  and  $\eta_{\text{dis}}$  are the charging and discharging efficiency ( $\in [0, 1]$ ) of the battery, respectively. The interaction between the agent and the environment helps the agent to estimate the transition probability distribution. Although the part of the state transition related to the SoC is deterministic, the agent is unaware of it and needs to learn it.

### 3.2 Distributional Deep Q Learning

We will solve the arbitrage problem, formulated as the MDP above, using a RL method. RL agents learn a policy to maximize the expected long-term reward. Classical tabular RL methods, e.g., Q-learning, cannot be applied to problems with high-dimensional or continuous state space due to the curse of dimensionality. Moreover, for replacing the tabular Q values with a function approximator, features need to be manually extracted [8]. To deal with these limitations, the deep Q learning (DQN) method suggests an idea of using a deep neural network as a function approximator to estimate the Q function. To learn the Q function  $Q_\theta(s_t, a_t)$ , the following loss function is minimized:

$$L_Q(\theta) = \mathbb{E}_{(s_t, a_t, r_t, s_{t+1}) \sim \mathcal{D}} \left[ (r_t + \gamma \max_a Q_{\theta'}(s_{t+1}, a) - Q_\theta(s_t, a_t))^2 \right]. \quad (6)$$

$$\theta' = \tau\theta + (1 - \tau)\theta' \quad (7)$$

The DQN method is stable in learning because of using the target Q function  $Q_{\theta'}$  for calculating next state-action values in Eq. (6), and training by a mini-batch sampled from an experience replay buffer  $\mathcal{D}$  [24]. Furthermore, this method can avoid overfitting the policy because it is an off-policy method that can learn from historical data, not just current experiences [33].

The authors in [6] introduced for the first time a distributional perspective on RL. These methods learn the probability distribution over returns instead of a single-value return. Distributional RL methods have various benefits, including the mitigation of Q-value overestimation [12], facilitating learning risk-sensitive policies [36], and improving training stability [6].

Extending beyond the basic DQN method, in DDQN, the probability distribution of returns ( $\mathcal{Z}_\theta$ ) is learned using the distributional Bellman equation below [6]:

$$L_{\mathcal{Z}}(\theta) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} [D_{\text{KL}}(\mathcal{T}\mathcal{Z}_{\theta'}(s_t, a_t) || \mathcal{Z}_\theta(s_t, a_t))] \quad (8)$$

$$\mathcal{T}\mathcal{Z}(s_t, a_t) \stackrel{D}{=} r_t + \gamma \max_a \mathbb{E}_{\mathcal{Z} \sim \mathcal{Z}_{\theta'}} [Z(s_{t+1}, a)] \quad (9)$$

In Eqs. (8) and (9)  $\mathcal{T}\mathcal{Z}_\theta$  is the probability distribution of  $\mathcal{T}\mathcal{Z}$ , and  $A \stackrel{D}{=} B$  indicates the equality of probability distributions for two

random variables  $A$  and  $B$ .  $D_{\text{KL}}(\cdot)$  denotes Kullback-Leibler (KL) divergence loss. we formulate the distribution over returns  $Z$  as a categorical distribution:

$$Z(s_t, a_t) = \left\{ z_i \mid z_i = V_{\min} + \frac{V_{\max} - V_{\min}}{N - 1} i, 0 \leq i < N \right\} \quad (10)$$

where  $V_{\min}$  and  $V_{\max}$  are the maximum and minimum values of random returns, respectively, and  $N$  is the number of bins.

#### 4 POLICY CORRECTION

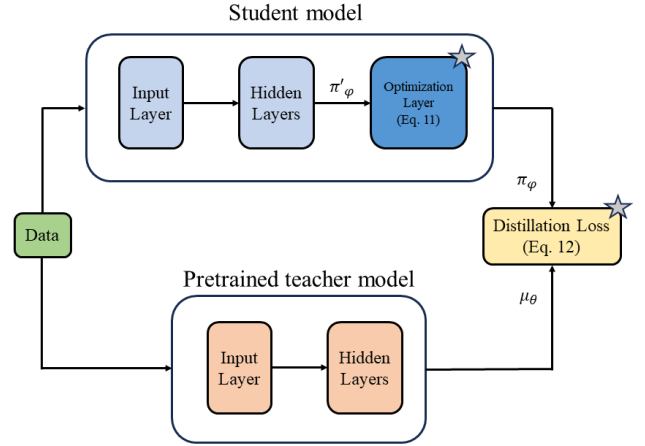
As mentioned earlier, vanilla RL methods cannot guarantee good performance in improbable states due to several reasons. First, as these states occur rarely (or sometimes they do not exist in the training set), the agent cannot learn a good action for them and tries to take action based on the agent's generalization ability. Second, the agent takes an action to maximize the *expected* return, not based on the *worst-case scenario* return, leading to taking catastrophic actions. In such a case, the agent might observe a few times that taking the specific action can cause a large negative return. However, since most of the time taking this action results in a high positive return, the agent still prefers this action over other possible actions. In this section, we propose a post-processing step to correct these catastrophic actions in our energy arbitrage problem. The purpose is to post-process the learned RL policy in a way that yields a correct and interpretable policy from a human perspective.

A human-intuitive energy arbitrage policy for our problem needs to possess three key properties: (1) charge the battery at very low prices; (2) discharge the battery at very high prices; (3) be monotonic with respect to price and SoC. The first two properties refer to the fact that the agent always must react properly to extremely rare prices regardless of the SoC level. The third property is the most crucial feature that makes the policy human-intuitive. For instance, we expect that when the agent decides to charge the battery at a specific price, it should also charge the battery at prices lower than this specific price, assuming other elements in the state remain constant. The following optimization problem ensures that the resulting neural network-based policy  $\pi_\phi$  has all the above-mentioned properties.

$$\begin{aligned} & \min_{\phi} \mathcal{L}_d[\mu_\theta(s_t), \pi_\phi(s_t)], s_t \in \mathcal{S} \\ \text{s.t. Property 1: } & \operatorname{argmax} \pi_\phi(s_t) = P_{\max}(\text{charging}), \\ & \text{if } \hat{\pi}_t^{\text{imb}} \leq \underline{\pi}^{\text{imb}} \\ \text{Property 2: } & \operatorname{argmax} \pi_\phi(s_t) = -P_{\max}(\text{discharging}), \quad (11) \\ & \text{if } \hat{\pi}_t^{\text{imb}} \geq \bar{\pi}^{\text{imb}} \\ \text{Property 3: } & \operatorname{argmax} \pi_\phi(s_t) \geq \operatorname{argmax} \pi_\phi(s'_t), \\ & \text{if } \text{SOC}_t \leq \text{SOC}'_t \text{ and } \hat{\pi}_t^{\text{imb}} \leq \hat{\pi}'_t^{\text{imb}} \end{aligned}$$

In Eq. (11),  $\underline{\pi}^{\text{imb}}$  and  $\bar{\pi}^{\text{imb}}$  are lower and upper bound thresholds for the imbalance price, respectively,  $\mathcal{L}_d(\cdot)$  is a loss function defined in Eq. (12), and  $\mu_\theta(\cdot)$  is the unconstrained neural network-based (pretrained) policy. Note that for the last property, other elements in  $s_t$  and  $s'_t$  are the same.

To solve the optimization problem in Eq. (11), we formulate it as a knowledge distillation process. Figure 1 shows the proposed framework for knowledge distillation. During knowledge distillation, the



**Figure 1: The proposed extension of knowledge distillation for policy correction. The starred blocks indicate our new addition to the standard distillation process.**

teacher model  $\mu_\theta$  is fixed and only the student model  $\pi_\phi$  is trained. The last layer of the student model is a differentiable optimization layer. The optimization layer enforces the three human-intuitive properties into the final policy. As these three properties are formulated as a least squares problem with linear constraints in this paper, the resulting optimization problem is convex and can be solved using the optimization layer. We use the *cvxpylayers* package to implement this optimization layer [2]. In the forward pass, the optimization layer solves the convex optimization problem, while the weights of the student model are trained end-to-end based on the gradient of the optimization solution during backpropagation.

$$\mathcal{L}_d[\mu_\theta(s_t), \pi_\phi(s_t)] = \omega D_{\text{KL}}(\pi_\phi \parallel \mu_\theta) + D_{\text{KL}}(\pi'_\phi \parallel \pi_\phi) \quad (12)$$

Equation (12) formulates the loss function in Eq. (11) where  $\pi'_\phi$  is the output of the student model before the optimization layer and  $\omega$  is a parameter to balance the two KL divergence losses. The first term in the loss function ensures that the student model closely mimics the behavior of the teacher model. On the other hand, the second term tries to keep the output of the student model before and after the optimization layer as close as possible. Although the optimization layer provides considerable flexibility in imposing convex constraints on the final policy, its computational burden to the student model makes the model inefficient, especially for real-world applications. To address this problem, the second term is added to the loss function. In that way, the student model is forced to train its weights to meet constraints in the absence of the optimization layer making the policy  $\pi'_\phi$  as close as possible to the constrained one  $\pi_\phi$ . Consequently, during inference, the optimization layer can be ignored, and the output of the student model before optimization layer ( $\pi'_\phi$ ) is used for decision-making.

Our intuition for using knowledge distillation for the policy correction has two parts. First, knowledge distillation provides a framework for reusing pretrained RL agents. In this way, BRPs can avoid retraining from scratch and effortlessly reuse pretrained

RL agents based on their own defined constraints and preferences. Second, the optimization layer is computationally expensive. In this arbitrage problem, the agent requires 50 000 episodes to converge. Considering a minibatch size of 16 384, training the teacher model with the optimization layer for 50 000 episodes would drastically increase the runtime to a point where the training loop may become intractable. On the other hand, in the knowledge distillation process, the student model only needs 600 epochs to be trained. Therefore, knowledge distillation can significantly reduce the computational time in our arbitrage problem.

In the end, it is noteworthy to highlight that our proposed post-processing step is a generic framework and it is not limited to energy domain applications. Using our proposed framework, the policy of pretrained RL agents can be corrected by a set of defined convex constraints. Moreover, the optimization problem defined in Eq. (11) is an example of constraints (rules) that can make the final energy arbitrage policy interpretable to humans. BRPs can have their individual preferences and define their own set of constraints. Also, these constraints depend on the inputs of the controller, and therefore can be adapted to those inputs.

## 5 RESULTS

In this section, we evaluate the performance of our proposed control framework, as explained in Sections 3 and 4, through simulation and experimental results.

### 5.1 Simulation Setup

We use the Belgian imbalance price of 2023 for evaluating our proposed framework.<sup>1</sup> BRPs imbalances are settled at 15-minute-based prices, which are calculated at the end of each quarter-hour period. To provide more information to BRPs, the TSO (Elia) also publishes 1-minute-based indicative prices in real-time which are calculated based on the instantaneous system imbalance and prices of cumulative activated regulation volumes on a minute basis.<sup>2</sup> We use these non-validated prices as a indicator of the real imbalance price for the related quarter hour period. By getting closer to the end of the quarter hour, the indicative price becomes closer to the real imbalance price. Thus, the time resolution of decision-making needs to be sufficiently short to benefit from the most recent situation of the grid using the indicative price and to take action accordingly. The resolution for the RL agent is hence set to 2 min. The characteristics of the simulated battery are 4MW/8MWh with 90% round-trip efficiency. To increase the lifetime of the battery, a minimum SoC is set at 10%. For the sake of proof-of-concept study, we assume that the battery does not participate in the day-ahead or other markets.

To train the initial unconstrained RL agent, the price dataset is split as follows: the first 20 days of each month as a training set, the 21st to the 25th of each month as a validation set, and the remaining days as a test set. The RL methods are trained with 50 000 episodes, where each episode constitutes a single day. The discount factor  $\gamma$ , the soft update factor  $\tau$ , the experience replay buffer size, and the mini-batch size are set to 0.999, 0.1,  $1 \times 10^6$ , and 16 384, respectively.

<sup>1</sup><https://opendata.elia.be/pages/home/>

<sup>2</sup>[https://www.elia.be/-/media/project/elia/elia-site/grid-data/balancing/20190827\\_end-user-documentation-elia1-minute-publications.pdf](https://www.elia.be/-/media/project/elia/elia-site/grid-data/balancing/20190827_end-user-documentation-elia1-minute-publications.pdf)

In both vanilla DQN and DDQN methods, the Q-value function and target Q function are modeled by a fully connected neural network that has two hidden layers with 256 and 128 neurons, respectively. The learning rate of the networks is  $5 \times 10^{-4}$ . In the DDQN method,  $V_{\max} = -V_{\min} = 1 \times 10^5$  and  $N = 51$ .

For the knowledge distillation process, a 2-layer fully connected neural network with hidden layer dimensions of 64 and 32 is used as the student model  $\pi_\phi$ . The teacher model architecture is the same as the model trained in the previous step. The student model is trained for 600 epochs with the learning rate of  $1 \times 10^{-3}$  using the Adam optimizer. Tuning  $\omega$  plays an important role in the training of the student model. Choosing large  $\omega$  results in a student model that more frequently violates constraints. Conversely, a small  $\omega$  leads to a student model that is less similar to the teacher model. To keep this balance, we used  $\omega = 1 \times 10^{-4}$ . The values of  $\underline{\pi}^{\text{imb}}$  and  $\bar{\pi}^{\text{imb}}$  in Eq. (11), are set to  $-500$  and  $1500$ , respectively. The main decision boundaries for the decision-making are learned by the RL agent:  $\underline{\pi}^{\text{imb}}$  and  $\bar{\pi}^{\text{imb}}$  just define safety thresholds for the agent. For this reason, these safety thresholds must correspond to extremely rare prices to avoid impacting the main policy and primary decision boundaries. However, the definition of these safety thresholds can vary among BRPs due to their different risk preferences. The PyTorch package in Python is used to implement our proposed control framework.

To benchmark the proposed control framework, a rule-based controller (RBC) is introduced as a baseline method. The RBC in this paper is a rudimentary threshold-based controller with two cutoff points based on the statistical analysis of imbalance prices. Although the RBC does not guarantee the optimal operation of the battery, it is widely used in the real-world due to its simplicity and ease of implementation [27]. These thresholds classify imbalance prices into three categories: cheap, normal, and expensive. In the RBC, the battery is (dis)charged with the maximum power when the price is (above) below a (upper) lower bound. The RBC action is formulated as follows:

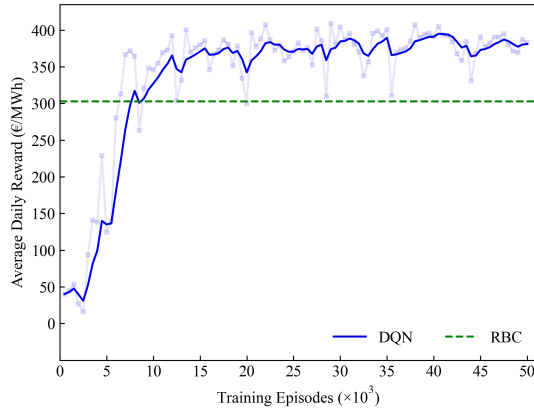
$$a_t = \begin{cases} P_{\max} & : \hat{\pi}_t^{\text{imb}} < \underline{\lambda} \\ 0 & : \underline{\lambda} \leq \hat{\pi}_t^{\text{imb}} \leq \bar{\lambda} \\ -P_{\max} & : \hat{\pi}_t^{\text{imb}} > \bar{\lambda} \end{cases} \quad (13)$$

where  $\underline{\lambda}$  and  $\bar{\lambda}$  represent lower and upper bounds, respectively. These thresholds are determined according to the distribution of the Belgian imbalance price in 2023: we set the upper and lower bounds to the first and third quartiles of the 2023 prices, which are equal to 11 €/MWh and 179 €/MWh, respectively.

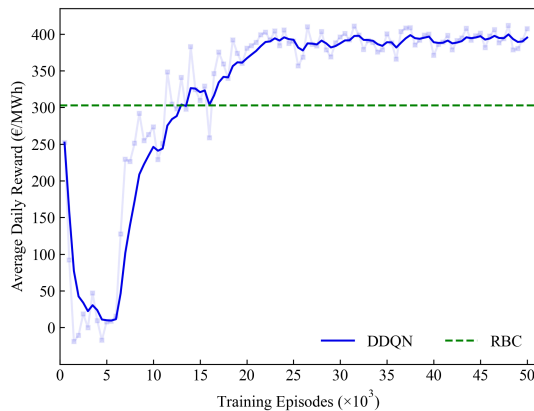
### 5.2 Simulation Results

The learning process of the RL methods is illustrated in Fig. 2, which shows the performance of the RBC and two RL methods on the validation set during the training. Table 1 shows the performance of the RBC and trained RL methods on the test set. The DDQN method increases the average daily profit by 32.2% and 9.2% compared to the RBC and DQN methods, respectively. The reason behind this is that the DDQN method estimates the probability distribution of returns instead of the expectation of returns. In this way, distributional RL methods mitigate instability in the Bellman optimality operator.

The proposed post-processing step improves the performance by 3.2% compared to that of the original DDQN model.



(a) DQN Agent



(b) DDQN Agent

**Figure 2: The learning process of (a) DQN and (b) DDQN. The solid and faded lines show smoothed and actual learning curves, respectively.**

**Table 1: Evaluation of RBC and RL methods on the test set.**

Method	RBC	DQN	DDQN	DDQN after policy correction
Profit (€/ day/ MWh)	341.1	413.1	450.9	<b>465.2</b>

Figure 3 illustrates how the policy correction step improves the performance by showing the DDQN policy heatmaps before and after applying the policy correction step. As the most determinative features for the DDQN agent are indicative imbalance prices and SoC, we plot the learned policy with respect to these two input

features for various times and months. It can be observed that the learned policy in some areas does not align with human intuition. For instance, there is a significant idle/ charge area in the policy for the 11th month at 10:00 when the price exceeds 1000 €/MWh, or there is a discharge region in the policy for the 1st month at 20:00 when the price is lower than  $-400$  €/MWh. The reason for these imperfection areas can be found in the probability distribution of the price, shown in Fig. 4. Prices below  $-200$  €/MWh rarely occur (with a probability of 3%), while prices above 1000 €/MWh are even rarer (with a probability of 0.04%). Thus, the Q function  $\mu_\theta$  overestimates Q values for out-of-distribution (OOD) actions in these rarely seen (or even unseen) states. As a result of this overestimation due to the max operator in the Bellman equation (Eq. (9)), the agent takes unexpected actions in these rarely seen states. As Fig. 3b shows, after our post-processing step, all imperfections in the policy are removed. The distilled (corrected) policy successfully mimics the main decision thresholds of the teacher model, while replacing all OOD actions with correct actions. As a result of replacing these OOD actions with correct actions, the student agent achieves higher performance than the original teacher model, as indicated in Table 1. Furthermore, applying the policy correction step makes the final policy interpretable with clear decision boundaries.

To compare the performance of the agent before and after the post-processing step, the student and teacher models are tested using data from September 29, 2023. As Fig. 5 shows, both models can effectively react to nearly all fluctuations in the price: they appropriately respond to two major peaks from 6:30 to 8:00 and from 10:30 to 12:00 by discharging the battery, and react to one major valley from 2:00 to 4:00 by charging the battery. However, the teacher model takes wrong actions between 8:30 and 10:00: the teacher agent decides to charge the battery between 8:15 and 8:30 when the price is approximately 130 €/MWh, while it does nothing from 8:30 to 10:00 when the price mostly hovers around 70 €/MWh and sometimes even dips to  $-20$  €/MWh. On the other hand, the corrected model continuously charges the battery from 8:00 to 10:00, which ultimately results in a higher profit since the battery has more energy to discharge in the following hours where prices are higher (from 10:00 to 12:00).

### 5.3 Experimental Results

We implement our proposed control framework on a real battery to better demonstrate its capabilities in a real-world setting. We use a residential AlphaESS battery with 4kW/ 8kWh, as shown in Fig. 6, which is installed in the imec/ Ghent University HomeLab. HomeLab is a real house offering a unique residential test environment for IoT services and smart home services.<sup>3</sup> For implementation, an API is available for the HomeLab which contains functionalities for the battery and other HomeLabGym devices [37]. This HomeLab API will forward the command to a battery-specific API. This codebase is running on a Raspberry Pi inside the HomeLab. The battery-specific codebase includes the translation needed from Python to the Modbus protocol in order to protocol to read and write appropriate registers on the battery energy management system (e.g., battery power) to the corresponding registers on the battery. The

<sup>3</sup><https://homelab.ilabt.imec.be/>

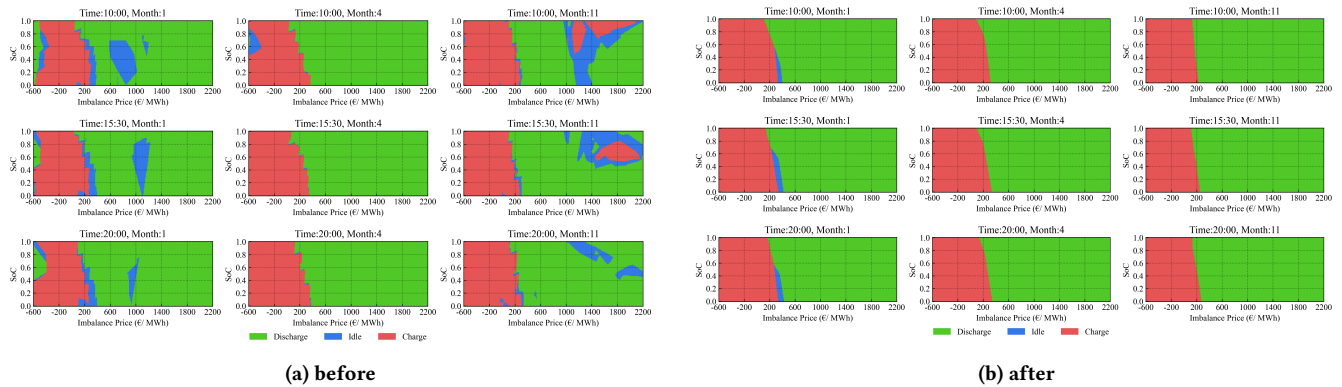


Figure 3: The learned policy (a) before and (b) after applying the policy correction step.




Figure 4: The probability distribution of the Belgian imbalance price in 2023.

DDQN model, trained on the imbalance prices of 2023, after applying the post-processing step is used to control the HomeLab battery with a granularity of 2 minutes. We control the battery from January 21, 2024 to January 28, 2024 with Elia’s real-time imbalance prices.

Table 2 lists the average daily profit during the trial period. The revenue for the real-world implementation is 6.7% lower than that for the simulation implementation (in the simulation implementation, we use the same price data as in the real-world implementation). There are three main reasons for this drop in the revenue: first, it takes 3s to take an action (includes doing calculations and fetch price data from Elia), 2s to send the taken action to the battery (communication delay), and finally, the battery needs an average of 5s to change it. Hence, it takes an average of 10s to select and execute the appropriate action, which accounts for nearly 8% of the 2-minute time step. Second, from January 22 at 22:00 to January 23 at 7:00, Elia faced technical issues that caused a delay in publishing imbalance price data. This led to our controller not receiving the necessary values to take appropriate actions. Third, sometimes




Figure 5: The performance of the proposed controller on September 29, 2023.

excessive frequent changes between charging and discharging can cause the temperature of the battery to rise, resulting in decreased battery efficiency. Figure 7 demonstrates a snapshot of the experimental results on January 28 between 10:00 and 14:00. The bottom row figure displays the actual power consumed by the HomeLab battery, with colors indicating the actions sent to the battery. This experimental study successfully demonstrated the deployability of our framework.

## 6 CONCLUSION

In this paper, we proposed a new RL-based control framework for batteries to perform energy arbitrage in the imbalance settlement mechanism. In our proposed control framework, we first train the

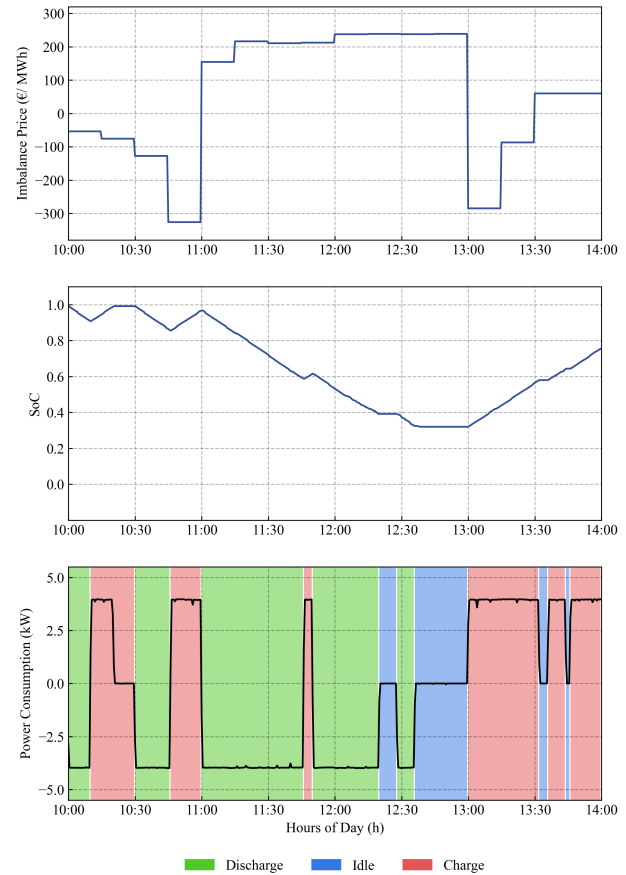


**Table 2: Evaluation of the proposed control framework on the real battery.**

Method	DDQN after policy correction (simulation)	DDQN after policy correction (real world)
Profit (€/ day)	6	5.6

**Figure 6: HomeLab battery used for testing the proposed control framework in real world**

agent to maximize the arbitrage revenue. Then, the post-processing step corrects (constrains) the learned policy during a knowledge distillation process based on defined human-intuitive properties. We extended the standard knowledge distillation process by incorporating an optimization layer into the student model and modifying the standard distillation loss to remove the (slow) differentiable layer at inference. The performance of the proposed control framework was evaluated through both simulation and experimental results using the Belgian imbalance price of 2023. The results, using DDQN as the RL algorithm, demonstrated that adding the post-processing step outperforms all of the RBC, vanilla DQN, and vanilla DDQN methods. The DDQN method with the policy correction step could improve the average daily profit by 36.4% and 12.6% compared to the RBC and DQN methods, respectively. This improvement stems from two factors: (i) the *distributional* perspective diminishes instability in the Bellman optimality operator by learning the full probability

**Figure 7: The experimental results on January 28, 2024 from 10:00 to 14:00.**

distribution of returns rather than a single value expectation of returns; (ii) the *post-processing* step deals with Q-value overestimation by replacing OOD actions with correct (human-intuitive) actions. During the post-processing step, the student model effectively distilled the knowledge from the pretrained teacher model regarding the main decision boundaries, while correcting the policy based on the defined human-intuitive properties. We deploy our proposed control framework in a real-world experimental setup, i.e., on the HomeLab battery to investigate its performance in a real-world environment. The experimental performance is 6.7% lower than the simulated environment, due to delays in action calculation and execution, missing data in real-time imbalance prices, and a decrease in battery efficiency as a result of temperature rise. Note that although we used DDQN in this paper, the proposed control framework is applicable to all RL methods.

In future research, we will focus on developing an online knowledge distillation process wherein both teacher and student models are trained end-to-end during the RL training loop. Another direction for future work is to consider adding more constraints to the proposed control policy, such as a (daily) cycle constraint for the battery.

## ACKNOWLEDGMENTS

We would like to thank AlphaESS and iinno-benelux companies for providing the battery and their technical support. This research has received funding from the Energy Transition Fund (FOD Economy) via the FlexMyHeat project. Part of the research leading to these results has received funding from Agentschap Innoveren & Ondernemen (VLAIO) as part of the Strategic Basic Research (SBO) program under the InduFlexControl-2 project.

## REFERENCES

- [1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained policy optimization. In *International conference on machine learning*. PMLR, 22–31.
- [2] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J Zico Kolter. 2019. Differentiable convex optimization layers. *Advances in neural information processing systems* 32 (2019).
- [3] Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. 2018. Safe reinforcement learning via shielding. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [4] Rohan Anil, Gabriel Pereyra, Alexandre Passos, Robert Ormandi, George E Dahl, and Geoffrey E Hinton. 2018. Large scale distributed neural network training through online distillation. *arXiv preprint arXiv:1804.03235* (2018).
- [5] Jens Baetens, Joannes Laveyne, Greet Van Eetvelde, and Lieven Vandevelde. 2020. Imbalance pricing methodology in Belgium: Implications for industrial consumers. In *2020 17th International Conference on the European Energy Market (EEM)*. IEEE, 1–6.
- [6] Marc G Bellemare, Will Dabney, and Rémi Munos. 2017. A distributional perspective on reinforcement learning. In *International conference on machine learning*. PMLR, 449–458.
- [7] Jérémie Bottieau, Louis Hubert, Zacharie De Grève, François Vallée, and Jean François Toubeau. 2020. Very-Short-Term Probabilistic Forecasting for a Risk-Aware Participation in the Single Price Imbalance Settlement. *IEEE Transactions on Power Systems* 35 (3 2020), 1218–1230. Issue 2.
- [8] Di Cao, Weihao Hu, Junbo Zhao, Guozhou Zhang, Bin Zhang, Zhou Liu, Zhe Chen, and Frede Blaabjerg. 2020. Reinforcement learning and its applications in modern power and energy systems: A review. *Journal of modern power systems and clean energy* 8, 6 (2020), 1029–1042.
- [9] Jun Cao, Dan Harrold, Zhong Fan, Thomas Morstyn, David Healey, and Kang Li. 2020. Deep reinforcement learning-based energy storage arbitrage with accurate lithium-ion battery degradation model. *IEEE Transactions on Smart Grid* 11, 5 (2020), 4513–4521.
- [10] Bingqing Chen, Priya L Donti, Kyri Baker, J Zico Kolter, and Mario Bergés. 2021. Enforcing policy feasibility constraints through differentiable projection for energy optimization. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*. 199–210.
- [11] Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. 2018. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757* (2018).
- [12] Jingliang Duan, Yang Guan, Shengbo Eben Li, Yangang Ren, Qi Sun, and Bo Cheng. 2021. Distributional soft actor-critic: Off-policy reinforcement learning for addressing value estimation errors. *IEEE transactions on neural networks and learning systems* 33, 11 (2021), 6584–6598.
- [13] European Network of Transmission System Operators for Electricity. 2018. Explanatory document to all TSOs' proposal to further specify and harmonise imbalance settlement in accordance with Article 52(2) of Commission Regulation (EU) 2017/2195 of 23 November 2017, establishing a guideline on electricity balancing. (2018).
- [14] Jie Feng, Yuanyuan Shi, Guannan Qu, Steven H Low, Anima Anandkumar, and Adam Wierman. 2023. Stability constrained reinforcement learning for decentralized real-time voltage control. *IEEE Transactions on Control of Network Systems* (2023).
- [15] Javier Garcia and Fernando Fernández. 2015. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* 16, 1 (2015), 1437–1480.
- [16] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. 2021. Knowledge distillation: A survey. *International Journal of Computer Vision* 129 (2021), 1789–1819.
- [17] Nick Harder, Anke Weidlich, and Philipp Staudt. 2023. Modeling Participation of Storage Units in Electricity Markets using Multi-Agent Deep Reinforcement Learning. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 439–445.
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [19] Jesus Lago, Ksenia Poplavskaya, Gowri Suryanarayana, and Bart De Schutter. 2021. A market framework for grid balancing support through imbalances trading. *Renewable and Sustainable Energy Reviews* 137 (2021), 110467.
- [20] Jesus Lago, Gowri Suryanarayana, Ecem Sogancioglu, and Bart De Schutter. 2021. Optimal Control Strategies for Seasonal Thermal Energy Storage Systems with Market Interaction. *IEEE Transactions on Control Systems Technology* 29 (9 2021), 1891–1906.
- [21] Hankook Lee, Sung Ju Hwang, and Jinwoo Shin. 2020. Self-supervised label augmentation via input transformations. In *International Conference on Machine Learning*. PMLR, 5714–5724.
- [22] Jinhao Li, Changlong Wang, Yanru Zhang, and Hao Wang. 2022. Learn to bid: Deep reinforcement learning with transformer for energy storage bidding in energy and contingency reserve markets. In *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning*.
- [23] Seyed Soroush Karimi Madahi, Bert Claessens, and Chris Develder. 2023. Distributional Reinforcement Learning-based Energy Arbitrage Strategies in Imbalance Settlement Mechanism. *arXiv preprint arXiv:2401.00015* (2023).
- [24] 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.
- [25] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE symposium on security and privacy (SP)*. IEEE, 582–597.
- [26] Peyman Passban, Yimeng Wu, Mehdi Rezagholizadeh, and Qun Liu. 2021. Alp-kd: Attention-based layer projection for knowledge distillation. In *Proceedings of the AAAI Conference on artificial intelligence*, Vol. 35. 13657–13665.
- [27] Baihong Qi, Mohammad Rashedi, and Omid Ardakanian. 2019. Energyboost: Learning-based control of home batteries. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*. 239–250.
- [28] Saidur Rahman, Linda Punt, Omid Ardakanian, Yashar Ghiassi, and Xiaoqi Tan. 2022. On efficient operation of a V2G-enabled virtual power plant: when solar power meets bidirectional electric vehicle charging. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 119–128.
- [29] Frederik Ruelens, Bert J Claessens, Stijn Vandael, Bart De Schutter, Robert Babuška, and Ronnie Belmans. 2016. Residential demand response of thermostatically controlled loads using batch reinforcement learning. *IEEE Transactions on Smart Grid* 8, 5 (2016), 2149–2159.
- [30] William Saunders, Girish Sastry, Andreas Stuhlmüller, and Owain Evans. 2017. Trial without error: Towards safe reinforcement learning via human intervention. *arXiv preprint arXiv:1707.05173* (2017).
- [31] Tim Seyde, Igor Gilitschenski, Wilko Schwarting, Bartolomeo Stellato, Martin Riedmiller, Markus Wulfmeier, and Daniela Rus. 2021. Is bang-bang control all you need? solving continuous control with bernoulli policies. *Advances in Neural Information Processing Systems* 34 (2021), 27209–27221.
- [32] Ruben Smets, Kenneth Bruninx, Jérémie Bottieau, Jean-François Toubeau, and Erik Delarue. 2023. Strategic Implicit Balancing with Energy Storage Systems via Stochastic Model Predictive Control. *IEEE Transactions on Energy Markets, Policy and Regulation* (4 2023), 1–14.
- [33] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [34] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT Press.
- [35] Jiayi Tang, Rakesh Shivanna, Zhe Zhao, Dong Lin, Anima Singh, Ed H Chi, and Sagar Jain. 2020. Understanding and improving knowledge distillation. *arXiv preprint arXiv:2002.03532* (2020).
- [36] Thibaut Théate and Damien Ernst. 2023. Risk-sensitive policy with distributional reinforcement learning. *Algorithms* 16, 7 (2023), 325.
- [37] Toon Van Puyvelde, Marie-Sophie Verwee, Gargya Gokhale, Mehran Zareh Esghdoust, and Chris Develder. 2024. HomeLabGym: A real-world testbed for home energy management systems. *arXiv preprint arXiv:2404.14110* (2024).
- [38] Behzad Vatandoust, Bashir Bakhshideh Zad, François Vallée, Jean-François Toubeau, and Kenneth Bruninx. 2023. Integrated Forecasting and Scheduling of Implicit Demand Response in Balancing Markets Using Inverse Optimization. In *2023 19th International Conference on the European Energy Market (EEM)*. IEEE, 1–6.
- [39] Gaoyuan Xu, Jian Shi, Jiaman Wu, Chenbei Lu, Chenye Wu, Dan Wang, and Zhu Han. 2024. An optimal solutions-guided deep reinforcement learning approach for online energy storage control. *Applied Energy* 361 (2024), 122915.
- [40] Qisong Yang, Thiago D Simão, Simon H Tindemans, and Matthijs TJ Spaan. 2021. WCSAC: Worst-case soft actor critic for safety-constrained reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 10639–10646.
- [41] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. 2017. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.

- 4133–4141.
- [42] Tianyu Zhang, Gaby Baasch, Omid Ardakanian, and Ralph Evins. 2021. On the joint control of multiple building systems with reinforcement learning. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*. 60–72.
- [43] Xianqiu Zhao, Yongbiao Yang, Minglei Qin, and Qingshan Xu. 2023. Day-ahead dispatch of novel battery charging and swapping station based on distributionally robust optimization. *Journal of Energy Storage* 63 (2023), 107080.