Reinforcement Learning Based Optimal Battery Control Under Cycle-based Degradation Cost

Kyung-bin Kwon, *Student Member, IEEE*, and Hao Zhu, *Senior Member, IEEE*

arXiv:2108.02374v3 [math.OC] 2 Jun 2022

Abstract—Battery energy storage systems are providing increasing level of benefits to power grid operations by decreasing the resource uncertainty and supporting frequency regulation. Thus, it is crucial to obtain the optimal policy for utility-level battery to efficiently provide these grid-services while accounting for its degradation cost. To solve the optimal battery control (OBC) problem using the powerful reinforcement learning (RL) algorithms, this paper aims to develop a new representation of the cycle-based battery degradation model according to the rainflow algorithm. As the latter depends on the full trajectory, existing work has to rely on linearized approximation for converting it into instantaneous terms for the Markov Decision Process (MDP) based formulation. We propose a new MDP form by introducing additional state variables to keep track of past switching points for determining the cycle depth. The proposed degradation model allows to adopt the powerful deep Q-Network (DQN) based RL algorithm to efficiently search for the OBC policy. Numerical tests using real market data have demonstrated the performance improvements of the proposed cycle-based degradation model in enhancing the battery operations while mitigating its degradation, as compared to earlier work using the linearized approximation.

Index Terms—Energy storage, reinforcement learning, battery degradation, rainflow algorithm, deep Q-networks (DQN)

I. INTRODUCTION

Battery energy storage systems as flexible resources are a key technology to enable the decarbonization of electricity infrastructure in future [1], [2]. Particularly, utility-level battery systems can be used to increase the payoff from electricity market via energy arbitrage [3], while contributing to the grid's power balance through participating in ancillary services [4]. It is crucial to develop effective strategies for real-time battery operations in order to utilize its flexibility potentials to mitigate the increasing uncertainty introduced by renewable or non-controllable loads.

The optimal battery control (OBC) problem for determining the (dis)charging policies has been popularly considered to reduce a combination of battery operational costs. It aims to reduce the net cost for electricity usage and frequency regulation (FR) penalty, as well as possible violations of network constraints; see e.g., [5]–[8]. In addition, battery's cycle life as characterized by the degradation cost is especially needed when participating in FR or other fast services [5], [7]. Unlike other costs that mostly depend on the instantaneous battery

status, the modeling of battery degradation is *cycle-based* according to the full trajectory of battery's state of charge (SoC). It requires the identification of all charging/discharging cycles using the so-termed *rainflow* algorithm [9]. Thus, degradation-aware OBC problem results in increased complexity as shown by [5], [7].

Due to the fast dynamics in prices or load demands, the OBC solution can be greatly affected by the uncertainty of future information. To address this issue, a model-predictive control (MPC) framework has been widely used by optimizing the current action according to the predicted input values for a fixed time window; see e.g., [10]–[12]. Nonetheless, the FR signal exhibits very minimal temporal correlation [13], leading to significant difficulty in predicting it and thus applying MPC for reducing FR penalty. Furthermore, even though battery health has been considered in MPC-based OBC work [14], [15], the cycle-based degradation model is largely missing.

Recent advances on reinforcement learning (RL) [16] enable the effective search of optimal control policies directly using real data samples to address the uncertainty issue in dynamical systems. This data-driven framework helps to bypass the hurdles in formulating the complex models of system dynamics or estimating the statistical information on the uncertainty. Specifically for the OBC problem, it allows to flexibly incorporate a variety of operational objectives, and several RL techniques have been widely used, such as the Deep Q-Network (DQN) [17], SARSA [18] and $TD(\lambda)$-learning [19]. However, a majority of these techniques have not considered the battery degradation cost, due to the difficulty of representing cycle-based model in the RL formulation. Very recently, [20] has developed a linearized approximation for the cycle-based degradation cost, which converts it to an instantaneous degradation coefficient that can be used by the RL algorithms. Nevertheless, the accuracy of this approximation method depends on a given sample trajectory based on which the linearization is performed. The resultant modeling mismatch can limit the RL iterations from finding the best policy within the full search space. Thus, it is still an open problem of effectively incorporating the accurate battery degradation cost into the search of OBC policy.

The goal of our work is to develop a modeling approach to precisely represent the battery degradation cost and use it for the design of RL-based OBC algorithm. The overall objective includes the net electricity cost, FR penalty, and cycle-based degradation cost. The main modeling challenge lies in the latter as it is determined by the battery's full SoC trajectory. Based on the rainflow algorithm, the complex process of material fatigue is associated with the stress level

of each individual charging or discharging cycle [21]. Thus, the degradation cost is an exponentially increasing function of cycle depth [9], and the latter strongly depends on the past trajectory of battery status. This leads to a pronounced mismatch with the Markov Decision Process (MDP) form used by RL algorithms, as the latter would represent the problem objective as functions of instantaneous states and actions only. The aforementioned approach of linearizing the degradation cost as in [20], [22] fails to recognize this exponential relation with the cycle depth, and unfortunately can lead to deep (dis)charging cycles that may not be overall profitable.

To this end, we have analytically shown that it is possible to keep track of the battery cycles by augmenting the state with the more recent switching points (SPs) along the SoC trajectory. These critical transition points between charging and discharging sessions are extremely useful for identifying the correct cycle depth according to the rainflow condition. In addition, they allow for decomposing the degradation cost of a full (dis)charging cycle into incremental differences between consecutive time instances in the form of instantaneous cost. This proposed representation of battery degradation cost helps to deploy state-of-the-art RL algorithms to learn the OBC policy. We have used the DQN technique to search for the parameters of the action-value function, or Q-function, associated with the resultant MDP form.

To sum up, the main contribution of the present paper is two-fold. First, we have developed an instantaneous cost modeling of the battery degradation with *guaranteed equivalence* to the original cycle-based representation based on rainflow algorithm. Second, the proposed degradation cost is successfully applied to form an MDP, allowing to develop efficient RL algorithms for the OBC problem. Our numerical tests using real data of electricity prices and FR signals have validated the performance improvement of the proposed cycle-based cost model in accurately representing battery degradation and effectively generating profit-seeking OBC policies.

The rest of the paper is organized as follows. Section II introduces the key variables for modeling the battery control problem into the MDP form. In Section III, we model the cycle-based degradation cost using the rainflow algorithm, and develop a new approach to represent it as instantaneous cost through state augmentation. Section IV formalizes the OBC problem and presents the DQN method as the RL solution technique. Numerical results using real-world data are presented in Section V to validate the performance improvement of the proposed degradation model, as compared to earlier approach using linearized approximation. Finally, the paper is wrapped up in Section VI.

II. SYSTEM MODELING

This paper considers the optimal battery control (OBC) for maximizing the economic pay-off while accounting for the battery degradation. The pay-off is from energy market participation and also the provision of FR service, as discussed later. One notable feature of the present work is the consideration of battery degradation cost, which can greatly increase the life-cycle under any general pay-off model [5]. A list of symbol notation and description is tabulated in Table I.

TABLE I: List of symbols

Notation	Description
c_t	state of charge (SoC) of the battery
$\bar{c}, \underline{c}$	maximum/minimum capacity of the battery
p_t	electricity market price
f_t	frequency regulation (FR) signal
s_t	battery full state
a_t	battery charging/discharging action
b_t	battery charging/discharging power
γ	discount factor
$\mathcal{T}$	the exploration time-horizon
h_t^e	net cost for electricity usage
h_t^f	frequency regulation penalty
h_t^d	battery degradation cost
δ	frequency regulation penalty coefficient
$\Phi(d)$	degradation cost for a cycle of depth d
α_d, β	degradation coefficient based on battery types
$c_t^{(0)}, c_t^{(1)}, c_t^{(2)}$	SoC level of switching points (SPs)

To determine the battery's effective (dis)charging power $b_t \in [\underline{b}, \ \bar{b}]$ at each discrete-time instance $t = 0, 1, \ldots$, we introduce a list of state variables based on battery status or external inputs.

- $c_t \in [\underline{c}, \ \bar{c}]$: normalized state of charge (SoC) of the battery;
- p_t: electricity market price;
- f_t: frequency regulation (FR) signal.

Note that the SoC is normalized by the maximum capacity; i.e., $c_t \in [0, 1]$. It is also an internal battery state affected by the past actions $\{b_\tau\}$, whereas the other states are received from grid operators and thus are not directly action-dependent.

To leverage reinforcement learning algorithms for this problem, we consider it as a Markov Decision Process (MDP) [23, Ch. 3] denoted by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, as detailed here.

State space $\mathcal{S}$ contains the set of feasible values for the system state s_t, including both the SoC c_t, and the other inputs p_t and f_t which affect the economic benefits. Additional state variables will be specified in Section IV for representing cycle-based degradation cost. State dynamics need to follow the Markov property as discussed soon.

Action space $\mathcal{A}$ includes the set of decisions that battery can take. We consider a discrete multi-level set with a total of $|\mathcal{A}|$ actions, as

$$a_t \in \mathcal{A} = \{a^{(1)}, a^{(2)}, \cdots, a^{(|\mathcal{A}|)}\} \qquad (1)$$

with normalized actions $a^{(n)} \in [-1, \ 1]$. Accordingly, the normalized (dis)charging power $b_t \in [\underline{b}, \ \bar{b}]$ is set to be

$$b_t = \begin{cases} \min\{\bar{c} - c_t, \bar{b}a_t\} & \text{if } a_t \geq 0, \\ \max\{\underline{c} - c_t, \underline{b}a_t\} & \text{if } a_t < 0. \end{cases} \qquad (2)$$

Continuous action space that directly determines $b_t = a_t$ is also possible. While this paper focuses on a discrete $\mathcal{A}$, the RL algorithm can be generalized to continuous a_t as well.

Transition kernel $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ captures the system dynamics under the Markov property [23, Ch. 3]. For the input states such as price p_t, we assume $\Pr(p_{t+1}|\{p_\tau\}_{\tau=1}^t) = \Pr(p_{t+1}|p_t)$; and similarly for f_t. This is reasonable as the market price has very short-term memory [24], while FR signal

f_t can be modeled as a white noise sequence of no memory [25]. A longer memory is possible too; such as the prices that follow $\Pr(p_{t+1}|\{p_\tau\}_{\tau=1}^t) = \Pr(p_{t+1}|p_t,\ p_{t-1})$. In this case, both p_t and p_{t-1} are included as the part of the state per time t to satisfy the Markov transition property.

Using Eq. (2), the SoC state c_t transitions as

$$c_{t+1} = c_t + b_t, \text{ with } b_t \text{ given in (2).} \qquad (3)$$

For general action space with any $b_t \in [\underline{b}, \bar{b}]$, c_t is updated by

$$c_{t+1} = \begin{cases} \bar{c} & \text{if } \bar{c} - c_t \le b_t, \\ \underline{c} & \text{if } \underline{c} - c_t \ge b_t, \\ c_t + b_t & \text{otherwise.} \end{cases} \qquad (4)$$

Reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ captures the learning objective. Notably, it is always the accumulated reward consisting of *instantaneous* terms, where per time t the latter only depends on the current state and action as

$$r_t = r_t(s_t, a_t) \qquad (5)$$

In the following we will minimize the objective cost function h_t as negative reward, where its instantaneous property will be ensured after introducing additional state variables as detailed in Section IV.

Discount factor $\gamma \in (0, 1]$ is a constant to accumulate the total reward along the time horizon. Smaller γ values imply that future rewards are less important than current ones at a discounted rate [23, Ch. 3]. As we adopt a finite exploration time-horizon $\mathcal{T} = [1, \ldots, T]$ for the OBC problem, for simplicity $\gamma = 1$ will be used.

III. MODELING OF BATTERY DEGRADATION COST

We consider three types of operational cost related to battery management. The energy cost relates to the electricity price according to (dis)charging, while the FR cost is based on its fast-varying flexibility. Under a contract of providing FR service, the battery would follow the f_t signal sent by the market operator as much as possible [5]. These two costs can be simply obtained by the state variables discussed so far. First, the net cost for electricity usage under (dis)charging power b_t can be represented as

$$h_t^e(p_t, b_t) = p_t b_t, \quad \forall\, t \in \mathcal{T}. \qquad (6)$$

Second, using a penalty coefficient δ for deviation from FR signal f_t, one can form

$$h_t^f(f_t, b_t) = \delta|f_t - b_t|, \quad \forall t \in \mathcal{T}. \qquad (7)$$

The energy cost in (6) is typically negative due to the energy arbitrage capability, while the FR penalty in (7) is always positive. This is because the additional economic benefit by participating in the FR contract is not included here. Overall, a battery should receive positive pay-off from these two tasks.

Remark 1. (Frequency regulation signal) *In practice, the FR signal is much faster than other system dynamics. For example, the real-time price is typically updated every 5 minutes, while the FR signal may be at 2-second rate [26]. To reduce the*
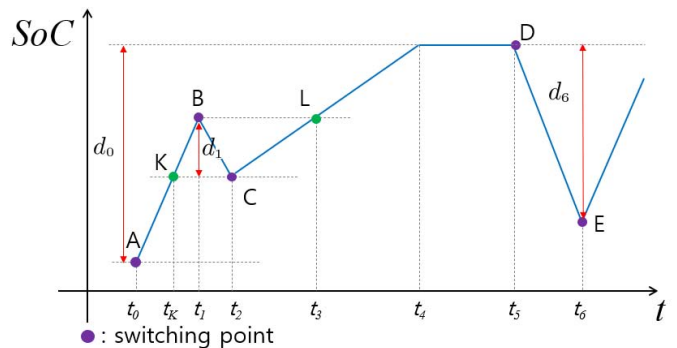

Fig. 1: An example of battery SoC trajectory used for modeling the battery degradation cost based on the rainflow algorithm.

complexity of the training computation later on, we will down-sample the FR signal to attain $\{f_t\}$ at a slower rate for searching the policy. In testing and implementing the resultant policy, the original fast FR signal will be instead used to realistically evaluate the performance of the RL approach.

As for the battery degradation cost, there are several stress factors affecting the battery lifetime such as temperature, high C-rates, average SoC, and Depth of Discharge (DoD) [9], [27]. During daily battery operations, the DoD stress model is considered the most relevant while other factors may be minimally affected. This will be shown numerically in Section V. According to the DoD stress model, the aging of battery cells mainly depends on material fatigue as a result of (dis)charging cycles of the SoC trajectory, especially due to following the FR signal [13]. Since this cycle-based degradation constitutes as a key battery lifetime consideration [28], the proposed OBC formulation to reduce it can greatly increase the battery's lifetime revenue.

The rainflow algorithm [9] is widely used for computing the cycle-based degradation cost. Fig. 1 illustrates an example of battery SoC trajectory which consists of several charging and discharging cycles. The switching points (SPs), labeled by $A - E$, correspond to the transitions between charging and discharging and will be used for identifying the cycles by rainflow algorithm. For example, the trajectory $A - B - C - D$ consists of a long charging cycle with a small discharging part from $B - C$. The respective depths of these two cycles, defined as the absolute SoC differences between the start and end SPs, are d_0 and d_1. As d_1 is smaller than the difference between $A - B$ and that between $C - D$, this trajectory is thus divided into the *full cycle* from $K - B - C$ of depth d_1 and the other *half cycle* from $A - K(C) - D$ of depth d_0. This is the so-called rainflow condition as stated in Lemma 1; see e.g., [5].

Lemma 1. *The SoC values of the last three SPs by time t are sufficient for evaluating the rainflow condition and determining the depth of (dis)charging cycles.*

Based on the cycle depth $d > 0$, the associated degradation cost is given by

$$\Phi(d) = \alpha_d e^{\beta d} \qquad (8)$$

with positive constant coefficients α_d and β based on battery

types [9], [22]. Recalling the normalized SoC $c_t \in [0,1]$, we have the cycle depth $d \in [0,1]$ as well. Note that for any pair in $\mathcal{D} := \{(d_1, d_2) : d_1, d_2 \geq 0, d_1 + d_2 \leq 1\}$, we can show that $e^{(d_1+d_2)} \leq e^{d_1} + e^{d_2}$. This is because the maximum value of the function $g(d_1, d_2) := e^{(d_1+d_2)} - e^{d_1} - e^{d_2}$ for the simplex $\mathcal{D}$ equals to $(e - 2e^{0.5}) < 0$, which is attained at $(d_1, d_2) = (0.5, 0.5)$. Thus, to reduce the degradation cost a single (dis)charging cycle that is longer and deeper is typically preferred, as opposed to the combination of multiple shorter cycles. This intuitive rule for cycle-based degradation model will be demonstrated later on in numerical tests. Unfortunately, this cycle-based degradation cost depends on the past SoC trajectory, and unfortunately, it does not follow the accumulated form of instantaneous terms as in Eq. (5).

Linearized degradation model has been developed in [20] to compute the averaged degradation coefficient from past SoC trajectory. Specifically, a degradation coefficient α_d is first determined using a given SoC trajectory over $\mathcal{T}$ as

$$a_d = \frac{\sum_{i=0}^{\bar{N}} \Phi(\bar{d}_i)}{\sum_{t=0}^{T} |\bar{b}_t|} \tag{9}$$

by averaging the total degradation costs of the $(\bar{N}+1)$ cycles over the accumulative absolute charging power throughout the sample trajectory. This way, the instantaneous degradation cost for any new SoC trajectory is approximated by

$$h_t^d(b_t) \cong -a_d|b_t|. \tag{10}$$

This linearized degradation cost model can be easily computed once a_d is known. However, this approximation inexplicitly assumes that the new trajectory should be very similar to the given sample trajectory for computing a_d. To implement the RL algorithm later on, the coefficient a_d will be updated using the most recent trajectory during the sampling process. Nonetheless, as an approximation it does not represent the actual cycle-based degradation cost and thus limits the RL algorithm's search for the best SoC trajectory.

Cycle-based degradation model will be pursued instead to address the approximation issue by augmenting the state s_t with the last three SPs before time t. As stated in Lemma 1, they are sufficient information for checking the rainflow condition. The state s_t now includes three additional variables, $c_t^{(0)}$, $c_t^{(1)}$, and $c_t^{(2)}$, as the SoC from the oldest SP to the latest one. Note that they may overlap if there are less than three SPs before time t. For example, at point K in Fig. 1, these three SP states all equal to the SoC of point A; and similarly for point C, we have $c_t^{(1)} = c_t^{(2)}$ equal to the SoC of B. The latest SP's SoC $c_t^{(2)}$ can be used to identify if the current instance t is a new SP, using the rule

$$b_t(c_t - c_t^{(2)}) < 0. \tag{11}$$

If Eq. (11) holds, we have a new SP and will update $\{c_t^{(i)}\}$ based on whether the rainflow condition is satisfied.

To update $\{c_t^{(i)}\}$, Fig. 2 illustrates two cases where the rainflow condition is not satisfied. Fig. 2(a) shows the SoC of the point C_1 is not within the range between A and B, while Fig. 2(b) indicates the SoC of the new SP D is within the range between C_2 and B. These cases are denoted by


Fig. 2: Two cases of rainflow condition not satisfied: (a) case NR_a and (b) case NR_b.

cases NR_a and NR_b, respectively. In either case, the oldest SP A will be removed while the remaining SPs will be used to update $\{c_{t+1}^{(i)}\}$, as listed in Table II. In addition, case RA denotes the scenario of rainflow condition being satisfied, such as the point D in Fig. 1. This is because at SP D: i) the third SP C lies between the first two SPs A and B; and ii) the current SoC at SP D exceeds the range between the latest two SPs B and C. Hence, the trajectory $A - B - C - D$ is divided in to the long half-cycle $A - K - C - L$, and another full-cycle $K - B - C$ of depth d_1. After the RA case is satisfied, the two SPs B and C will be removed from the record. The SoC state updates for all three cases are summarized in Table II.

Interestingly, the state transitions in Table II also allow for decomposing the cycle-based degradation cost into instantaneous difference term for each instance t. As the cycle depth changes according to b_t only, the degradation cost in Eq. (8) can be modeled by accumulating the following incremental term per time t:

$$h_t^d(b_t, c_t, c_t^{(2)}) = \alpha_d e^{\beta|c_t + b_t - c_t^{(2)}|} - \alpha_d e^{\beta|c_t - c_t^{(2)}|}. \tag{12}$$

Basically, the instantaneous degradation model in Eq. (12) accounts for difference of $\Phi(\cdot)$ due to the change of cycle-depth, which can be computed based on the latest SP state $c_t^{(2)}$. Therefore, summing up all instantaneous terms in Eq. (12) yields the total degradation cost, as formally stated in **Proposition 1** with the proof provided in the Appendix.

Proposition 1. *Under Lemma 1, the summation of the instantaneous terms in Eq.* (12) *throughout the time-horizon* $\mathcal{T} = [1, \cdots, T]$ *is exactly equivalent to the total cycle-based degradation cost along* $\mathcal{T}$.

TABLE II: State transitions at a new SP identified at time t

Next state	NR_a	NR_b	RA
$c_{t+1}^{(0)}$	$c_t^{(1)}$	$c_t^{(1)}$	$c_t^{(0)}$
$c_{t+1}^{(1)}$	$c_t^{(1)}$	$c_t^{(2)}$	$c_t^{(0)}$
$c_{t+1}^{(2)}$	$c_t^{(1)}$	c_t	$c_t^{(0)}$

IV. OPTIMAL BATTERY CONTROL ALGORITHM

Thanks to our proposed model of instantaneous degradation cost, we can define the MDP form for the OBC problem. First, each state is given by

$$s_t = [p_t, f_t, c_t, c_t^{(0)}, c_t^{(1)}, c_t^{(2)}], \quad \forall t \in \mathcal{T} \quad (13)$$

which is used to determine the action a_t based on the policy of interest. The transition kernel $\mathcal{P}$ now includes the updates in Table II, while the instantaneous reward is given by

$$r_t(s_t, a_t) = -h_t^e - h_t^f - h_t^d, \quad \forall t \in \mathcal{T}. \quad (14)$$

The battery control problem now becomes to determine the best policy π for forming the action as $a_t \sim \pi(s_t)$ with s_t given in Eq. (13). To simplify the policy search, we are particularly interested in the set of parameterized policies given by $\pi_\theta(\cdot) = \pi(\cdot; \theta)$, with parameter θ optimized through

$$\max_\theta \ \mathbb{E}_{\pi_\theta} \left[\sum_{t=1}^T \gamma^t r_t(s_t, \ a_t) \right]. \quad (15)$$

To solve Eq. (15), we can adopt certain RL algorithms to search for the optimal parameter θ; see e.g., [16]. We use the deep Q-networks (DQNs) [23, Ch. 9] here as a popular RL approach based on nonlinear neural network modeling. Accordingly, the parameter θ represents the DQN weights to be learned, and the DQN is used to obtain the so-termed Q-network that models the MDP's *action-value, or Q-function*, namely the expected total future reward under a given pair of state and action:

$$Q(s_t, a_t) := \mathbb{E}_{\pi_\theta} \left[\sum_{\tau=t}^T \gamma^{(\tau-t)} r_\tau(s_\tau, a_\tau) \Big| s_t, a_t \right]. \quad (16)$$

For the optimal Q-function, the Bellman optimality condition [16] states that:

$$Q^*(s_t, a_t) = r_t(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1}} \left[\max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1}) \Big| s_t, a_t \right]. \quad (17)$$

To find the optimal Q^*, we parameterize the action-value Q-function using θ as the NN weights, as denoted by $Q(s_t, a_t; \theta)$. The Bellman optimality in Eq. (17) can be used to develop iterative *gradient descent* updates to obtain the best θ. At each update, the Q-network on the right-hand side of Eq. (17) is kept constant as the *target network*, whereas the other one is varied to minimize the difference between both sides. Letting θ' denote the latest NN weights, we design the loss function for DQN training as the expected squared difference:

$$\mathcal{L}(\theta) = \mathbb{E}_{\{s_t, a_t, s_{t+1}\}} \left[\left(r_t + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}; \theta') - Q(s_t, a_t; \theta) \right)^2 \right]. \quad (18)$$

To minimize $\mathcal{L}(\theta)$, one can need to compute its gradient over the parameter θ given by

$$\nabla_\theta \mathcal{L}(\theta) = \mathbb{E}_{\{s_t, a_t, s_{t+1}\}} \Big[-2 \Big(r_t + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}; \theta') - Q(s_t, a_t; \theta) \Big) \nabla_\theta Q(s_t, a_t; \theta) \Big]. \quad (19)$$

Each gradient-based update relies on the estimate from sampling the trajectory such that the expectation in Eq. (19) is replaced by the sample average. To this end, the action a_t is sampled for given state s_t based on θ' as $a_t^* = \arg\max_{a_t} Q(s_t, a_t; \theta'), \forall t$. To ensure adequate exploration of the state space, the ϵ-greedy method [23, Ch. 2] can be used to randomize the action by selecting a_t^* with probability $(1 - \epsilon)$ at every time. The value of ϵ would decrease as the DQN updates continue, typically at an exponential decreasing rate $\kappa \in (0, 1)$. This method can improve the exploration process at the beginning phase while eventually picking the optimal actions to attain convergence.

To improve the efficiency and stability of DQN implementation, we introduce two additional techniques. First, we implement the *experience replay* method [29] to efficiently use the past samples by storing all the past samples in the memory $\mathfrak{D} := \{(s_t, a_t, s_{t+1}, r_t)\}$ along the trajectory. When computing the loss function Eq. (18), a subset of samples denoted by mini-batch $\mathfrak{J}$ is randomly picked from $\mathfrak{D}$ and used as the samples for gradient estimation. This method can improve the training efficiency by selectively reusing past samples. In addition, we advocate the *fixed target network* approach [30] by keeping the target network parameter fixed for several updates. To this end, let θ^- denote the target network parameter, which is only updated once every N_o iterations. This technique could mitigate any potential instability issue by changing the DQN target weights less frequently. By adopting *experience replay* method and *fixed target network* approaches, we obtain the estimates of loss function and its gradient as

$$\nabla_\theta \hat{\mathcal{L}}(\theta) = (1/|\mathfrak{J}|) \sum_{t \in \mathfrak{J}} \Big[-2 \Big(r_t + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}; \theta^-) - Q(s_t, a_t; \theta) \Big) \nabla_\theta Q(s_t, a_t; \theta) \Big]. \quad (20)$$

The detailed algorithmic steps for DQN-based OBC algorithm are tabulated in Algorithm 1. As mentioned earlier, the state variables p_t and f_t are not action dependent. Thus, their transitions are obtained from the profiles given as the algorithm input, such as real data provided by the market operators. For the convergence of DQN algorithms, the total number of episodes N is typically chosen to be large enough in practice. For each episode n, there are total T samples from $t = 1$ to $t = T$. Note that Algorithm 1 can be used to search for the best policy under the linearized degradation cost as well, by using this simpler degradation cost model in Eq. (10). The ensuing section will compare these two degradation models numerically.

V. NUMERICAL TESTS

We have compared the proposed RL-based battery control algorithm under cycle-based degradation cost with the

---

Algorithm 1: DQN-based Optimal Battery Control

---

1 **Hyperparameters:** discount factor $\gamma = 1$, learning rate $\eta > 0$, ϵ-greedy coefficient $\kappa \in (0,1)$, mini-batch size $|\mathfrak{J}|$, target network update interval N_o, and maximum number of episodes N.

2 **Input:** training profiles of prices and FR signals with the exploration time horizon T.

3 **Initialize:** the ϵ-greedy probability $\epsilon \in (0,1)$, replay memory $\mathfrak{D} = \emptyset$, initial action-value function $Q(s, a; \theta')$ with a random θ' and the target network parameter $\theta^- = \theta'$ at episode $n = 0$.

4 **while** $n \leq N$ **do**

5   **for** $t=1, \cdots, T$ **do**

6   Select a random action a_t with probability ϵ; otherwise, use the action $a_t^* = \arg\max_{a_t} Q(s_t, a_t; \theta')$.

7   Implement the action a_t to obtain the ensuing state s_{t+1} based on the transitions of both Eq. (4) and Table II, and by using the input profiles of p_{t+} and f_{t+1}.

8   Compute the instantaneous reward r_t in Eq. (14).

9   Store the tuple (s_t, a_t, s_{t+1}, r_t) in $\mathfrak{D}$.

10   Select a random mini-batch $\mathfrak{J}$ with size $|\mathfrak{J}|$ from $\mathfrak{D}$.

11   Compute the gradient estimate using Eq. (20).

12   Update the parameter $\theta' \leftarrow \theta' - \eta \nabla \hat{\mathcal{L}}(\theta')$.

13   **if** t/N_o *is an integer* **then**

14   Update the target network parameter $\theta^- \leftarrow \theta'$.

15   **end**

16   Update $\epsilon = \kappa\epsilon$.

17   **end**

18   Update the episode number $n \leftarrow n + 1$.

19 **end**

TABLE III: Parameter settings for DQN training

Parameter	Value
Number of hidden layers	2
Number of nodes	[128, 32]
Activation function	ReLU
Learning rate	0.001
Optimizer	Adam
Epsilon (ϵ)	0.001
Batch size (J)	256
Maximum number of episodes (N)	2000
Number of daily profiles	7

(a)

(b)

Fig. 3: Comparisons of the total reward trajectory between (a) cases LD1 and CD1 (α_d) and (b) case LD2 and CD2 ($2\alpha_d$)

linearized approximation one [20]. Actual data of electricity market prices and FR signals have been used, respectively from the ERCOT's market data depository [31] and PJM's ancillary service datasets [32]. Each time instance corresponds to a 5-minute interval. The FR signal is normalized to indicate either maximum charging or discharging for the battery. As mentioned in Remark 1, the fast FR signal at 2-second rates is averaged over a 10-second interval for the training phase, while the original data rate is maintained for the testing phase. We have used a 200kWh-capacity battery with (dis)charging rate of 120kW and minimum SoC of 20kWh, which takes 90 minutes to fully (dis)charge. The multiple discrete action space is adopted with overall 11 actions, as $\mathcal{A} = \{-1, -0.8, \cdots, 0.8, 1\}$. The parameters associated with battery degradation are set to $\alpha_d = 4.5 \times 10^{-3}$ and $\beta = 1.3$, as used in [22].

The DQN **Algorithm 1** has been implemented in Python with the popular NN toolboxes Tensorflow and Keras [33]. Table III lists the parameter settings for the DQN training, which uses 7 daily profiles of $\{p_t, f_t\}$. There are a total of $T = 8,640$ time instances for each exploration episode. Upon the convergence of Q-network, it is used for determining the optimal (dis)charging actions for each 2-second interval of 60 days of testing data, while each testing trajectory having 43,200 time instances.

Training Comparisons. To compare the proposed cycle-based degradation cost with the linear one (denoted by CD and LD, respectively) in terms of battery control performance, We have considered two levels of degradation coefficients, at α_d and $2\alpha_d$, respectively. Fig. 3 illustrates the training comparisons of the actual episode rewards (all based on cycle-based degradation) for all the three cases. Clearly, all the DQN iterations are convergent as the total reward trajectories tend to be non-decreasing till reaching the highest values. Moreover, while the CD cases using our proposed *instantaneous cycle-based degradation* modeling outperform the LD counterparts, especially at larger degradation cost. This comparison validates the advantages of the proposed degradation model in terms of accurately representing the battery cost and thus leading to

(a)

(b)

Fig. 4: The total reward differences (positive difference indicating higher reward for CD) between (a) cases LD1 and CD1 (α_d) as well as (b) case LD2 and CD2 ($2\alpha_d$)

TABLE IV: Reward differences between CD and LD

| Cases (A-B)	Mean	Max	Min	Mean (A > B)	Mean (A < B)
CD1-LD1	84.45	133.09	-97.04	111.38	-72.12
CD2-LD2	176.51	315.15	-198.85	172.79	-134.19

(a)

(b)

Fig. 5: The battery degradation cost differences (positive differences indicating lower cost for CD) between (a) cases LD1 and CD1 (α_d) as well as (b) case LD2 and CD2 ($2\alpha_d$)

effective control policies.

Testing Comparisons. We have further compared the testing performance of the learned Q-networks using both CD and LD based models. Fig. 4 plots the total reward differences between the CD and LD solutions (positive differences indicating higher reward for CD) for each test trajectory under the two levels of degradation coefficients. The proposed CD based control leads to higher total reward for at least 73.33% or 81.67% of test scenarios, respectively for the two α_d levels. This result confirms the earlier observations in training phase that proposed solution is more attractive for larger degradation cost. Table IV indicates the total mean, maximum, minimum values and average values of the cases when CD has better reward than LD, and vice versa. As shown in the table, the overall mean value increases as the degradation coefficient increases and the battery degradation cost affects more in total cost accordingly. In addition, the maximum, minimum and mean values show that even though there are some cases that LD show better performance than CD, the number of these cases is very small. Similar comparison is also observed in differences of battery degradation performance only (again, positive differences indicating lower degradation cost for CD), as illustrated by Fig. 5. Clearly, the proposed CD solutions overwhelmingly improve the battery degradation performance and accordingly the total reward, as compared to the existing LD-based approximation. In addition, Fig. 4 and Fig. 5 share very similar pattern, which implies that the increase in total reward is mostly caused by the decrease in the battery degradation cost and has least impacts on the decrease in the rewards regarding the net cost for electricity usage or frequency regulation penalty.

Fig. 6 plots the selected testing SoC trajectory along with the electricity price to better illustrate the improvement of the

proposed CD-based policy, corresponding to the two choices of degradation parameter (α_d and $2\alpha_d$). Clearly, both trajectories show that the CD-based policy leads to less number of cycles with long depth as compared to the LD one, which reduces the overall degradation cost especially for hours between [3, 13]. In addition, during high-price hours in [12, 15], the CD trajectory has one smooth and long discharging cycle and this pattern is amendable to mitigating battery degradation. In contrast, the LD one has frequent, noticeable fluctuations during this period. Because of the linearized approximation, the LD-based policy leads to eight more noticeable (dis)charging cycles of considerate depth than the CD one. This speaks for the capability of the proposed CD model in effectively removing some unnecessary cycles of moderate depth, thanks to the accurate representation of rainflow-based degradation. In addition, in the post-peak hours [15, 20], the LD based policy produces a couple of cycles of moderate depth which are not very profitable. The proposed CD based policy is able to successfully remove these nonprofitable cycles and does not lead to any considerate cycles.

Interestingly, by mitigating cycle-based degradation the proposed CD approach can potentially contribute to the improvement of other degradation factors too. Table V compares the proposed CD with the LD approach on the degradation related to high C-rates [27] and SoC stress [9], both of which have been numerically improved by the CD-based policies. The high C-rate based degradation depends on the
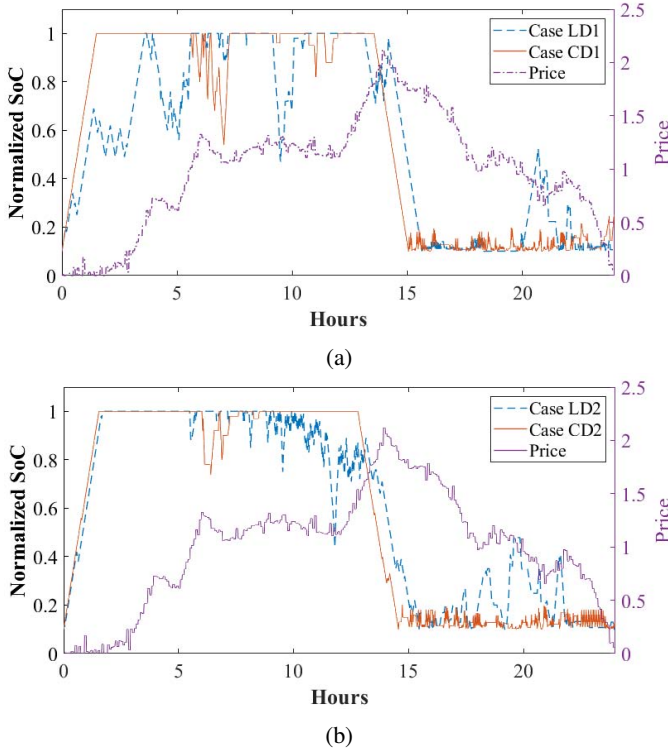
(a)

(b)

Fig. 6: Comparison of selected SoC trajectories in testing between (a) cases LD1 and CD1 (α_d) and (b) cases LD2 and CD2 ($2\alpha_d$)

TABLE V: Degradation comparisons between CD and LD

Degradation factor	LD1	CD1	LD2	CD2
High C-rates	0.0563	0.0342	0.0463	0.0369
SoC stress	0.0141	0.0108	0.0182	0.0107

total DoD summed over all cycles of the trajectory. Intuitively, a concise list of smooth and long (dis)charging cycles attained by CD-based policy can reduce both the number of cycles and their DoD, thus beneficial for the high C-rate metric. Similarly, as CD-based policy has also been observed to remove unnecessary cycles in the post-peak hours, the average SoC level decreases which relieves the SoC stress. These intuitions corroborate the claim in Section III that the cycle-based DoD stress model is most relevant for the fast battery control problem.

To sum up, the numerical results have validated the performance improvement attained by the proposed *instantaneous* cycle-based degradation model, by exactly representing the rainflow conditions. The proposed approach effectively leads to battery control trajectories that reduce unnecessary fluctuations or improve the overall economical profits.

VI. CONCLUSION AND FUTURE WORK

The paper proposes an accurate model of cycle-based degradation cost in order to allow for efficient battery control designs using reinforcement learning (RL). In order to model the degradation which depends on the full cycle, we introduce additional state variables to judiciously keep track of important switching points of SoC trajectory for effectively identifying

(dis)charging cycles. This way, the actual degradation cost is separated into instantaneous terms along with other operation costs such as the net cost for electricity usage and FR penalty, such that powerful DQN based RL algorithms are readily applicable. Numerical tests confirm the effectiveness of proposed cycle-based degradation model and demonstrate the performance improvements in effectively mitigating battery degradation over existing linearized approximation approach.

Exciting future research directions open up on expanding the battery operations to support other ancillary services such as peak-shaving. The key question will be how to model the peak-shaving cost as instantaneous reward. In addition, the power network constraints such as voltage limit are of high interest in practice and will be considered as well.

APPENDIX

The goal is to show that Eq. (12) represents the exact incremental degradation cost from time t to $(t+1)$. With the current SoC denoted as c_t at time t, the current cycle depth equals to $d_t = c_t - c_t^{(2)}$ starting from the SP $c_t^{(2)}$ in time t_0. Thus, the resultant degradation cost is given by

$$\Phi(d_t) = \alpha_d e^{\beta d_t} = \alpha_d e^{\beta|c_t - c_t^{(2)}|}$$

with the absolute difference capturing cycle depth. When transitioning to time $(t+1)$, the incremental difference in degradation cost due to action b_t becomes

$$\Delta_t = \Phi(d_{t+1}) - \Phi(d_t) = \Phi(d_t + b_t) - \Phi(d_t)$$
$$= \alpha_d e^{\beta|c_t + b_t - c_t^{(2)}|} - \alpha_d e^{\beta|c_t - c_t^{(2)}|}.$$

Therefore, summing up these differences leads to

$$\Phi(d_{t+1}) = \Phi(d_t) + \Delta_t = \sum_{\tau=t_0+1}^{t+1} \Delta_\tau \qquad (21)$$

with the degradation cost at time t_0 initialized by zero. Notably, there are two types of scenarios when considering this summation at time $(t+1)$, as detailed here.

(1) Case NR_a and case NR_b: In either case, rainflow condition is not satisfied as depicted in Fig. 2, and thus Eq. (21) accumulates the total degradation thus far for this cycle.

(2) Case RA: Without loss of generality, consider the RA cycle $A-K-B-C-L-D$ as shown in Fig. 1. As the rainflow condition is satisfied at point L, the total degradation cost of this cycle equals to $\bar{\Phi} = \Phi(d_0) + 2\Phi(d_1)$. Following from Eq. (21), the degradation cost from point A to L is obtained by

$$\Phi^{A \to L} = \Phi(d_K + d_1) + 2\Phi(d_1).$$

Similarly, from point L to D, the cycle continues on as $A - B - L - D$ due to the satisfaction of rainflow condition, and Eq. (21) leads to an additional degradation cost as

$$\Phi^{L \to D} = \Phi(d_0) - \Phi(d_K + d_1).$$

Together, the total degradation cost for the cycle $A - K - B - C - L - D$ equals to $\bar{\Phi}$, which completes the proof for Proposition 1.   □

References

[1] M. Arbzabzadeh, R. Sioshansi, J. X. Johnson, and G. A. Keoleian, "The role of energy storage in deep decarbonization of electricity production," *Nature Communications*, vol. 10, no. 1, p. 3413, Jul 2019. [Online]. Available: https://doi.org/10.1038/s41467-019-11161-5

[2] P. Denholm, E. Ela, B. Kirby, and M. Milligan, "Role of energy storage with renewable electricity generation," National Renewable Energy Lab (NREL), Tech. Rep., 2010.

[3] D. Krishnamurthy, C. Uckun, Z. Zhou, P. R. Thimmapuram, and A. Botterud, "Energy storage arbitrage under day-ahead and real-time price uncertainty," *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 84–93, 2018.

[4] N. Padmanabhan, M. Ahmed, and K. Bhattacharya, "Battery energy storage systems in energy and reserve markets," *IEEE Transactions on Power Systems*, vol. 35, no. 1, pp. 215–226, 2020.

[5] Y. Shi, B. Xu, Y. Tan, D. Kirschen, and B. Zhang, "Optimal battery control under cycle aging mechanisms in pay for performance settings," *IEEE Transactions on Automatic Control*, vol. 64, no. 6, pp. 2324–2339, 2019.

[6] B. Xu, Y. Shi, D. S. Kirschen, and B. Zhang, "Optimal battery participation in frequency regulation markets," *IEEE Transactions on Power Systems*, vol. 33, no. 6, pp. 6715–6725, 2018.

[7] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 2882–2894, 2018.

[8] S. Gupta, V. Kekatos, and W. Saad, "Optimal real-time coordination of energy storage units as a voltage-constrained game," *IEEE Transactions on Smart Grid*, vol. 10, no. 4, pp. 3883–3894, 2019.

[9] B. Xu, A. Oudalov, A. Ulbig, G. Andersson, and D. S. Kirschen, "Modeling of lithium-ion battery degradation for cell life assessment," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 1131–1140, 2018.

[10] T. Morstyn, B. Hredzak, R. P. Aguilera, and V. G. Agelidis, "Model predictive control for distributed microgrid battery energy storage systems," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 3, pp. 1107–1114, 2017.

[11] A. Oshnoei, M. Kheradmandi, and S. M. Muyeen, "Robust control scheme for distributed battery energy storage systems in load frequency control," *IEEE Transactions on Power Systems*, vol. 35, no. 6, pp. 4781–4791, 2020.

[12] K. Meng, Z. Y. Dong, Z. Xu, and S. R. Weller, "Cooperation-driven distributed model predictive control for energy storage systems," *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2583–2585, 2015.

[13] B. J. Kirby, "Frequency regulation basics and trends," 5 2005. [Online]. Available: https://www.osti.gov/biblio/885974

[14] C. Zou, C. Manzie, and D. Nešić, "Model predictive control for lithium-ion battery optimal charging," *IEEE/ASME Transactions on Mechatronics*, vol. 23, no. 2, pp. 947–957, 2018.

[15] D. M. Rosewater, D. A. Copp, T. A. Nguyen, R. H. Byrne, and S. Santoso, "Battery energy storage models for optimal control," *IEEE Access*, vol. 7, pp. 178 357–178 391, 2019.

[16] B. Recht, "A tour of reinforcement learning: The view from continuous control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, no. 1, pp. 253–279, 2019. [Online]. Available: https://doi.org/10.1146/annurev-control-053018-023825

[17] A. S. Zamzam, B. Yang, and N. D. Sidiropoulos, "Energy storage management via deep q-networks," in *2019 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2019, pp. 1–5.

[18] J. Liu, H. Tang, M. Matsui, M. Takanokura, L. Zhou, and X. Gao, "Optimal management of energy storage system based on reinforcement learning," in *Proceedings of the 33rd Chinese Control Conference*, 2014, pp. 8216–8221.

[19] B. V. Mbuwir, F. Ruelens, F. Spiessens, and G. Deconinck, "Battery energy management in a microgrid using batch reinforcement learning," *Energies*, vol. 10, no. 11, 2017. [Online]. Available: https://www.mdpi.com/1996-1073/10/11/1846

[20] J. Cao, D. Harrold, Z. Fan, T. Morstyn, D. Healey, and K. Li, "Deep reinforcement learning-based energy storage arbitrage with accurate lithium-ion battery degradation model," *IEEE Transactions on Smart Grid*, vol. 11, no. 5, pp. 4513–4521, 2020.

[21] R. Spotnitz, "Simulation of capacity fade in lithium-ion batteries," *Journal of Power Sources*, vol. 113, no. 1, pp. 72–80, 2003. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775302004901

[22] Y. Shi, B. Xu, Y. Tan, and B. Zhang, "A convex cycle-based degradation model for battery energy storage planning and operation," in *2018 Annual American Control Conference (ACC)*, 2018, pp. 4590–4596.

[23] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. The MIT Press, 2018. [Online]. Available: http://incompleteideas.net/book/the-book-2nd.html

[24] T. Jónsson, "Forecasting and decision-making in electricity markets with focus on wind energy," Ph.D. dissertation, 2012.

[25] D. Zhu and Y.-J. A. Zhang, "Optimal coordinated control of multiple battery energy storage systems for primary frequency regulation," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 555–565, 2019.

[26] PJM Manual 12: Balancing Operations. https://pjm.com/~/media/documents/manuals/m12-redline.ashx. Accessed: 2022-03-01.

[27] J. Wang, P. Liu, J. Hicks-Garner, E. Sherman, S. Soukiazian, M. Verbrugge, H. Tataria, J. Musser, and P. Finamore, "Cycle-life model for graphite-LiFePO4 cells," *Journal of power sources*, vol. 196, no. 8, pp. 3942–3948, 2011.

[28] B. Foggo and N. Yu, "Improved battery storage valuation through degradation reduction," *IEEE Transactions on Smart Grid*, vol. 9, no. 6, pp. 5721–5732, 2017.

[29] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Machine Learning*, vol. 8, no. 3, pp. 293–321, May 1992. [Online]. Available: https://doi.org/10.1007/BF00992699

[30] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb 2015. [Online]. Available: https://doi.org/10.1038/nature14236

[31] *ERCOT Market Prices*. [Online]. Available: http://www.ercot.com/mktinfo/prices.

[32] *PPJM Ancillary Services*. [Online]. Available: https://www.pjm.com/markets-and-operations/ancillary-services.aspx.

[33] *Keras: The Python Deep Learning Library*. [Online]. Available: https://keras.io.

Biographies

Kyung-bin Kwon (S'21) received the B.S. and M.S. degrees in electrical and computer engineering from Seoul National University, Seoul, South Korea, in 2012 and 2014, respectively. He is currently working towards the Ph.D. degree at The University of Texas at Austin, Austin, USA. His current research focuses on reinforcement learning applications in power system and data-driven control of networked system.

Hao Zhu (M'12–SM'19) is an Associate Professor of Electrical and Computer Engineering (ECE) at The University of Texas at Austin. She received the B.S. degree from Tsinghua University in 2006, and the M.Sc. and Ph.D. degrees from the University of Minnesota in 2009 and 2012. From 2012 to 2017, she was a Postdoctoral Research Associate and then an Assistant Professor of ECE at the University of Illinois at Urbana-Champaign. Her research focus is on developing innovative algorithmic solutions for problems related to learning and optimization for future energy systems. Her current interest includes physics-aware and risk-aware machine learning for power system operations, and energy management system design under the cyber-physical coupling. She is a recipient of the NSF CAREER Award and an invited attendee to the US NAE Frontier of Engr. (USFOE) Symposium, and also the faculty advisor for three Best Student Papers awarded at the North American Power Symposium. She is currently an Editor of *IEEE Trans. on Smart Grid* and *IEEE Trans. on Signal Processing*.