Deep Q-Learning for Aggregator Price Design

Aisling Pigott, Kyri Baker, and Cory Mosiman

Department of Civil, Environmental, and Architectural Engineering
University of Colorado Boulder, Boulder, CO, USA
{aisling.pigott, kyri.baker, cory.mosiman}@colorado.edu

***Abstract*—This paper develops a reinforcement learning (RL) aggregator agent that intelligently determines real-time electricity prices to achieve grid-level goals. A diverse community of home energy management (HEMS) systems, equipped with model predictive control (MPC) capabilities, respond to the prices without revealing their state, resources, constraints, or objectives to the aggregator. Only the aggregate response of the community is visible to the RL agent, preserving the privacy of each homeowner. The results indicate that the RL agent is successful in improving the flattening the aggregate community demand while adhering to individual user constraints without requiring two-way negotiations, transactive frameworks, or any knowledge about individual homes.**

I. Introduction

Currently options for dynamic residential electricity pricing are generally limited to predetermined time-of-use pricing (TOU) structures or in very limited cases, real-time pricing. However, TOU pricing can have natural limitations and adverse effects as the number of participating consumers increases [1], [2], and existing real-time pricing frameworks do not adapt to consumers' willingness to pay; rather, they are often based on locational marginal prices (LMPs) [3]. Dynamic pricing which reflects current objectives of the system operator while adapting to the balance of supply and demand is still an open research question, one which grows more relevant as we approach highly variable, zero-marginal cost grid operation [4]. In particular, CAISO recognizes short and steep ramping periods as one of three major challenges for communities with high penetration renewable resources [5]. Recent work suggests the use of batteries to replace gas-fired generators that are typically used to meet peak demand, but batteries are also limited in their effective ramping rates and are resource intensive [6].

In this paper, we develop a reinforcement learning (RL) framework which generates a real-time electricity price signal sent by an aggregator in attempts to elicit a desired aggregate load profile through purely thermal energy storage. Each home, equipped with a model predictive control (MPC)-based home energy management system (HEMS), attempts to optimize its own cost function while adhering to individual user constraints. The proposed model-free framework of RL does not require knowledge of the objectives or constraints of the homes, does not need off-line training data, and is scalable enough to calculate real-time prices.

Prior work has been performed exploring the use of RL in designing real time prices for residential consumers. The objective for learning electricity prices is generally given as a function of the cost incurred from the wholesale LMP and revenue from the consumer [7]–[9]. Much of the work that uses RL at the aggregator level also use RL at the HEMS level. These HEMS assume that community members are willing to cooperate by sharing information about their desired demand, to achieve a flattened aggregate demand with no external motivating factors. Some of these HEMS also neglect realistic demand response models. These simplified techniques include randomly sampled instantaneous demand [9] or using a set percent of instantaneous demand as a load shift resource with a discounted consumption rate [8]. However, more realistic modeling of the consumer-side response is important to understanding the benefits and limitations of RL for price determination.

As in [1], we also acknowledge that the efficacy of pricing signals is limited by the response of houses in the network. Therefore we aim to validate the use of RL generated signals at the aggregator level when the community responds with home energy management systems that reflect typical system dynamics and obey occupant comfort constraints. We look to the linear reduced order models (ROMs) of HVAC systems, domestic water heaters, and photovoltaic (PV) systems [10], [11] for use in an MPC based HEMS for each consumer.

The proposed approach offers the following benefit over the aforementioned works that optimize price signals:

1) In reality, building models are not readily available to an aggregator. RL is a model-free way of optimizing price signals; the dynamics of the model can be arbitrary (non-convex, probabilistic, no closed form, etc).
2) RL naturally adapts to changes in model parameters; i.e. as equipment ages, or additional assets are added.
3) Unlike many other frameworks which require consumers to submit bid curves or cost functions, the proposed RL framework does not require any two-way interaction with consumers or any information to be exchanged between consumer and utility other than power consumption, which can be read from smart electric meters.
4) The aggregator price signal is uniform for all consumers, in pursuit of fairness.

The remainder of the paper is formulated as follows. In Section II, the reinforcement learning agent is formulated and the action space is defined. In Section III, a generic home energy management system is developed for use in the analysis of community response. Finally, in Section IV we present our results, and Section V concludes the paper.

II. Reinforcement Learning for Dynamic Pricing

Reinforcement learning bypasses the need for a physically accurate environment model by directly approximating a score, the Q-value, as a parameterized function of the environment's observed state, $\hat{x}$, and the agent's action, u. This score is extrapolated across all values of the state space, $\mathcal{X}$, and action space, $\mathcal{U}$. The score given by the Q-value represents the reward obtained as a function of the resulting state, $r(x_{t+1})$, and the possible future reward. The Q-function is updated online as the RL agent explores the environment.

Conventionally, RL algorithms maximize the Bellman's operator of the q-value over an infinite time horizon (Eq. (1)). Let the Q-function parameterization (Eq. (2)) be an approximation of the q-value achieved over any remaining timesteps. The parameterization is given by a weight vector, θ_t, and state-action vector, ϕ. In Deep Q-Learning, the state and state-action vectors are replaced by a neural network. The Q-value reflects the critic in the actor-critic agent which scores the state-action pair after it is implemented.

$$\max_{u \in \mathcal{U}} \left(r(x_0) + \gamma \sum_{t=1}^{\infty} q(x_t, u_t) \right) \tag{1}$$

$$\sum_{t=1}^{\infty} q(x_t, u_t) \approx Q_\theta(x_1, u_1) = \theta_t \cdot \phi(x_t, u_t) \tag{2}$$

The aggregator proposed utilizes the Soft Actor Critic (SAC) algorithm for continuous state and action spaces. The SAC algorithm utilizes experience replay, dual Q-networks, and entropy maximization to improve performance over the conventional Q-learning [12]. The implementation of SAC used can be found in the open-source package for RL in Python, Stable Baselines [13].

A. Stochastic Actor Networks

In stochastic actor networks, actions are selected from a Gaussian distribution where the mean and variance are independently parameterized functions of the state vector. Let the weight vectors for the mean, μ, and the variance, σ, be given by θ_μ and θ_σ respectively. The policy is given as:

$$u_t \sim \mathcal{N}\left(\theta_\mu \cdot \phi(x), \theta_\sigma \cdot \phi(x)\right) \tag{3}$$

In contrast to the critic network which optimizes the RL rewards at future timesteps, the actor network attempts to optimize just the current timestep by selecting the presumed best action μ with high probability. In actor-critic networks, the policy and Q-value can be combined into a single *value* term that describes the expected Q-value according to the probability of selecting the action given by the policy. The value is estimated by a third parameterized function.

B. Experience Replay

Experience replay speeds up convergence of the actor and critic networks and is often used in RL algorithms for control. Experience replay samples m memory tuples from a set of observed memory tuples $\mathcal{M}$ and uses them in addition to the current state-action-reward tuple to train the current policy. Let a memory tuple be defined as $(x_j, u_j, r_j, x_{j+1}, u_{j+1})$ where j is an arbitrary previous timestep. Experience replay is used for Q-function, policy, and value function updates; here, we show an example of experience replay for Q-function approximation with the objective of minimizing the Bellman's residual squared error given in (4). Let the subscript θ denote the current parameterization of the respective Q, V, and policy functions.

$$\delta_{Qj} = \mathbb{E}_{(x_j, u_j) \sim \mathcal{M}} \left[\frac{1}{2} \left(Q_\theta(x_j, u_j) - q(x_j, u_j) \right)^2 \right] \tag{4}$$

where $q(x_j, u_j) = r_j + V_\theta(x_{j+1}, u_{j+1})$. The general form of gradient descent for each of the three functions is preserved with m gradient steps per timestep:

$$\theta_Q \leftarrow \theta_Q + \alpha \cdot \nabla_\theta Q_\theta \cdot \delta_{Qj} \tag{5}$$

Similarly, the weight vectors of the value and policy networks are updated to minimize residual error and Kullback–Leibler divergence respectively [12].

C. Dual Q-Networks and Entropy Maximization

Dual Q-Networks increase the stability of the Q-function. The work in [14] shows that the Q-function approximation is positively biased. To prevent divergence, two Q networks are updated alternately and the minimum of the two Q-values is used when calculating the Bellman's error in (4) and other gradient descent steps. Entropy maximization rewards the RL agent for higher exploration values, in contrast to the conventional objective of maximizing the Bellman's operator of the Q-value to a multi-objective problem. The tuning parameter of this secondary objective is known as the temperature constant and is auto-tuned throughout the training period in SAC [12].

D. State Observations and Reward Function Design

The Q-function is extrapolated from observed rewards at past state-action pairs and the expected Q-value. Let the Q-function be given by the deep neural network comprised of state values, current and forecasted aggregate demand, datetime values, observed demand over the last n timesteps, forecasted weather data (GHI and OAT), and the observed maximum demand for the current day. Once an action is selected to maximize the expected Q-value the reward price signal is sent to the MPC responsive homes and carried out. The resulting state is evaluated by the RL agent with the weighted multi-objective reward function in (6), based on the aggregate demand.

$$r(\hat{x}_t) = -\left(\frac{p_{grid}^{(t)}}{\bar{p}_{grid}} \right)^2 + \left(\frac{p_{grid}^{(t)}}{\bar{p}_{grid}} \right) \tag{6}$$

The reward function presented here promotes solutions where the community consumes half of the maximum possible aggregate demand. Normalizing the aggregate observed

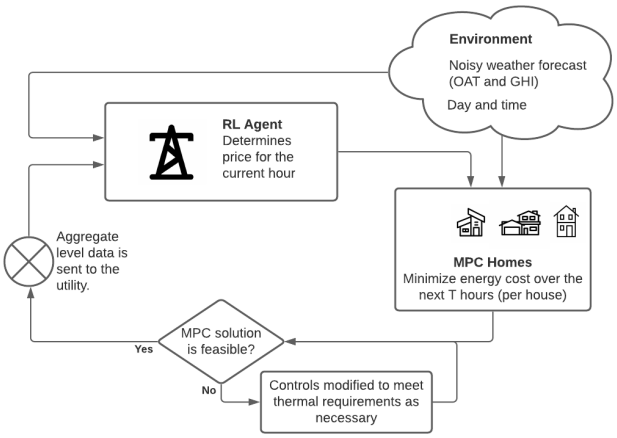

Fig. 1: Aggregator-Home Interaction

demand $p^{(t)}_{grid}$ against the maximum possible demand load allows for the RL agent to be immediately scaled to larger communities without modifying the RL agent.

The RL agent was tuned by increasing the default learning rate by a factor of 10^2. Because the environment is dynamic, a faster learning rate that trades some precision for flexibility over the course of several weeks of simulation time is preferred. Rewards, state observation values, and control actions are normalized during the initialization of the RL agent.

III. Consumer Response and Device Models

One of the major goals of this paper is to compare the use of RL-designed real time pricing (RTP) with other types of dynamic pricing using realistic energy consumption models rather than assuming a general function for overall home energy consumption as in prior work. When aggregated, these home energy models can be considered as a black box h-model of the RL agent's environment. This environment is provided in DRAGG-Gym, an OpenAI Gym interface for price responsive community model described here.

Note that in a black box formulation, input and output state values, x_t and x_{t+1}, are the true state value of the system, in contrast to the state values observed by the RL agent, $\hat{x}_t$. Assume that not all system data can feasibly be collected or accounted for and that $\hat{x}_t$ is a subset of x_t. This allows for the case where every participating building does not have to have the capability to report their consumption in real-time. Here, we only observe the aggregate demand of the community, which can be measured at distribution substations, for instance.

The MPC based home energy management system (HEMS) that each participating building is equipped with is provided in the opensource Python package Distributed Resource AGGregation (DRAGG) [15] and is used to model consumer response to broadcasted RTP signals. The objective function within each HEMS minimizes the cost over time horizon τ for each house i, at every RL timestep. Fig. 1 shows the interaction of the MPC HEMS with the RL agent aggregator. DRAGG utilizes the opensource CVXPY wrapper for convex optimization and the opensource mixed integer linear programming solver GLPK_MI. We consider the linear model of HVAC equipment in single thermal zone houses, domestic water heaters, solar panel arrays, and home battery systems to be a sufficient approximation. Subsystems such as residential HVAC equipment are modelled using the discrete control over the MPC timestep to reflect minimum run times of the home subsystems in each duty cycle. The MPC timestep is independent of the RL aggregator pricing timestep.

Due to each home's thermal mass, HVAC equipment offers the greatest opportunity for load shifting and is considered in the base controllable system for the homes participating in the real time pricing scheme (i.e., every home is assumed to have an electric HVAC system). The linearized thermal model is given in (7) and (8) using an single resistance (R_i) and capacitance (C_i) node. The discrete HVAC operation is given by $S^{(t)}_{cool}, S^{(t)}_{heat}$ which represent the duty cycle state of the heating and cooling systems rated at $P_{cool,i}$ and $P_{heat,i}$, respectively, at time t. During winter months, cooling is turned off, and during summer months, heating is turned off to prevent the cycling of heating and cooling systems.

$$T^{(t)}_{in,i} = T^{(t-1)}_{in,i} + \frac{\left(\frac{T^{(t)}_{OAT,i} - T^{(t-1)}_{in,i}}{R_i} + P^{(t)}_{HVAC,i}\right)\Delta t}{C_i} \tag{7}$$

$$P^{(t)}_{HVAC} = S^{(t)}_{heat,i} P_{heat,i} - S^{(t)}_{cool,i} P_{cool,i} \tag{8}$$

Water heaters also provide load shifting via thermal energy storage. Water heater tanks are modeled with a single resistance ($R_{wh,i}$) and capacitance ($C_{wh,i}$) node as shown in (9) and respond to the time varying water draw profiles ($d^{(t)}_i$) as observed in [16] as a fraction of their total tank capacity $\zeta_{wh,i}$. Water draws are modelled as a single draw in (10) occurring between control timesteps (i.e. time t+0.5). The tank is refilled with tap water ($T_{tap,i}$) at 15°C. The rated capacity of the water heater is given by $P_{wh,i}$.

$$T^{(t)}_{wh,i} = T^{(t-0.5)}_{wh,i} + \frac{\left(\frac{(T^{(t)}_{in,i} - T^{(t-1)}_{wh,i})}{R_{wh,i}} + S^{(t)}_{wh,i} P_{wh,i}\right)\Delta t}{C_{wh,i}} \tag{9}$$

$$T^{(t-0.5)}_{wh,i} = \frac{\zeta_{wh,i} - d^{(t-1)}_i}{\zeta_{wh,i}} T^{(t-1)}_{wh,i} + \frac{d^{(t-1)}_i}{\zeta_{wh,i}} T^{(t-1)}_{tap,i} \tag{10}$$

PV generation is given by the linear model in (11) as a function of solar irradiance ($P^{(t)}_{GHI}$), solar panel area ($A_{pv,i}$), solar panel efficiency ($\eta_{pv,i}$), and the current curtailment ($U^{(t)}_i$). Solar irradiance prediction is outside of the scope of this project and is read in from another source by DRAGG. Solar irradiance is obtained from the National Renewable Energy Laboratory's National Solar Radiation Database [17].

$$P^{(t)}_{pv,i} = P^{(t)}_{GHI} A_{pv,i} \eta_{pv,i} (1 - U^{(t)}_{pv,i}) \tag{11}$$

System parameters for each home within the community were selected from a uniform random distribution with the

ranges provided in Table I. All HVAC system sizing parameters (rated power, resistance, and capacitance) and the water heater power rating and resistance values are based on the values provided in [10]. The tank size and capacitance of the water heater are independently selected based on the Department of Energy guidance of 75 liters per occupant and assuming 3-4 occupants per home [18]. The capacitance is based on the thermal capacitance of the water volume.

	WH	**HVAC**	**PV**
P (kW)	[2.5]	[3.5]	-
R (°C/W)	[18.7, 25.3]	[6.8, 9.2]	-
C (kJ/°C)	[840, 1260]	[4.25, 5.75]	-
Setpoint (°C)	[45.5, 48.5]	[18, 22]	-
Deadband (°C)	[2, 3]	[9, 12]	-
Size	[200, 300] L	-	[20, 32] m^2
Efficiency	1.0	1.0	[0.15, 0.2]
Control	Discrete	Discrete	Continuous

TABLE I: Subsystem Parameters

The interaction of each house i with the grid is given as:

$$P^{(t)}_{grid,i} = P^{(t)}_{load,i} - P^{(t)}_{pv,i} \tag{12}$$

$$P^{(t)}_{load,i} = S^{(t)}_{cool,i} P_{cool,i} + S^{(t)}_{heat,i} P_{heat,i} + S^{(t)}_{wh,i} P_{wh,i} \tag{13}$$

where P_{load} gives the sum of energy consumed by the active systems considered in each house, P_{pv} is the PV generation and P_{grid} is the net demand of the house.

The HEMS optimize a cost-based objective, subject to the aforementioned constraints (depending on what resources each individual HEMS controls). An additional weighting factor, $\beta \leq 1$, discounts the cost of actions that are taken at future timesteps exponentially. The objective, for each home i, is:

$$\min \sum_{t=0}^{\tau} \beta^t c^{(t)}_e P^{(t)}_{grid,i} \tag{14}$$

where $c^{(0)}_e = c^{(0)}_{e,base} + c_{e,reward}$ and $c^{(t)}_e = c^{(t)}_{e,base}\ \forall\ t > 0$.

The HEMS observe uncertainty in their environments from the outdoor air temperature, hot water draw times/sizes, and GHI (for PV-enabled houses). The weather forecast for OAT and GHI include normally distributed noise with a standard deviation of up to 0.5°C and 2.5% of the GHI. HEMS are given a forecast of all water draws within 1 hour of the actual draw and implement an expected value for the water draw averaged over 3 hours.

IV. Simulation Results

The hyperparameters of the RL agent were tuned until the RL agent improved performance of the community demand response when compared to static electricity pricing and/or a TOU price structure. In this paper, we consider the main objective of the grid operator to be load flattening and therefore evaluate performance on the basis of peak demand and load spread (max - min load) over 24 hours. The RL agent itself was evaluated for performance improvement on the basis of observed reward, entropy, policy loss, value loss, and Q-losses.

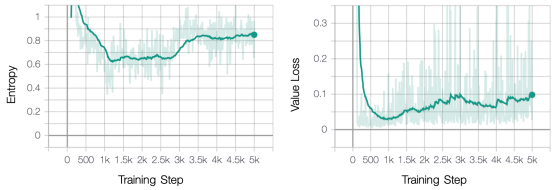

(a) Entropy (b) Value Loss

Fig. 2: RL Agent Validation

Num. Houses, Total	20
Num. Houses, with PV	4
HEMS MPC Horizon	6 hours
HEMS Runtime Interval	10 min
Baseline Electricity Cost	$0.10/kWh
TOU Off Peak Electricity Cost	$0.07/kWh
TOU Shoulder Electricity Cost, (9AM - 9PM)	$0.09/kWh
TOU Peak Electricity Cost, (2PM - 6PM)	$0.13/kWh
RL Agent Min, Max Reward Price	-$0.02, $0.02/kWh

TABLE II: Simulation Parameters

Over 5 epochs of 1000 training steps each the average reward of the RL agent increases to achieve a 2.5% reward improvement over the baseline case. This improvement is also shown by convergence of the entropy and value loss (Fig. 2).

	RL	**Baseline**	**TOU**
Overall Peak Demand (kW)	67.5	69.8	85.6
Average Daily Spread (kW)	16.1	17.4	21.8

TABLE III: P_{grid} Summary

Table III shows a 3.4% reduction from the baseline in the observed maximum demand over a 30 day period. Due to the impact of peak demand in the use of expensive “peaker" power plants and impact on infrastructure upgrades and sizing, reducing peaks in aggregate consumption can be one of the most important drivers for grid planners. In [19], for instance, a 5% reduction of peak demand across the US was estimated to save $3 billion annually. Beyond reducing the peak demand, our results show a 7.5% reduction of the average daily load spread (the difference between the max and min load), which would drastically reduce dependence on fast-ramping generators and needs for shutting off/turning on plants throughout the day.

The improvement of the agent is seen in the environment response, as measured by the daily load spread and the maximum aggregate demand. Out of 30 simulated days, the maximum daily demand is reduced from the baseline on 16 days, and is reduced from either the baseline or TOU on 28 days. Fig. 3 shows maximum daily demand and 12-hour rolling average for the second week of simulation time where the peak demand is reduced or maintained for 5 of the 7 days shown. In the average daily demand profile (Fig. 4), the maximum demand is reduced by 8% and the minimum demand is increased by 13%. Therefore, the RL agent also successfully reduces the average daily load spread. We also see that the community’s total energy consumption is preserved due to the thermal constraints on each HEMS.

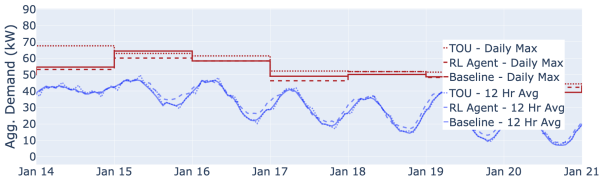

Fig. 3: 12-Hour Moving Average and Daily Max Load

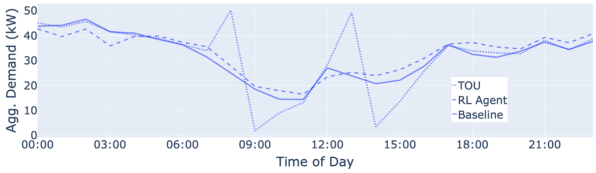

Fig. 4: Daily Load Profile Averaged over 30 Days

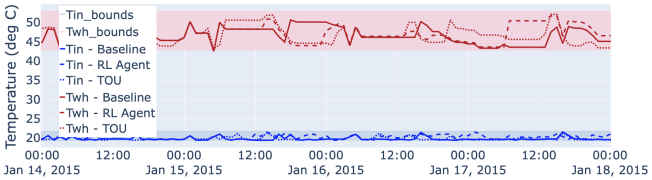

Fig. 5: T_{in} and T_{wh}

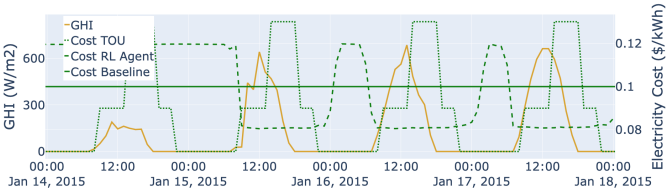

Fig. 6: Real Time Price and GHI

When the RL agent with the performance described above is initialized on a community with 50% more houses and proportionally more PV, the peak demand is reduced by 1% and the average daily spread is maintained. Comparatively, a new RL agent with no prior training maintains the peak demand but increases the average daily spread by 5%. The use of transfer learning, in which one RL agent is applied to a new environment, creates a stable initialization period without developing a full model of the community.

The responses of the MPC controlled HEMS can be validated by tracking the subsystem temperatures within each home (Fig. 5) and individual home demand and generation. Note that the vast majority of timesteps produce a feasible MPC solution except in cases of large water draws. In communities with significant PV penetration, the electricity price signal is strongly correlated with the GHI, as shown in Fig. 6, as the aggregator attempts to incentivize consumption during periods of high renewable production. The RL determined price signal skews positive so in order to encourage enrollment in such a price structure the utility could scale the base cost and the reward price magnitude proportionally, for example.

V. Conclusions

In this paper we have shown that RL can be used to design a real time price signal to successfully flatten load curves in residential communities with minimal impact on consumers, and only with visibility of the community's aggregate demand. We have further shown that an RL agent trained on one community can be applied to another community with similar makeup and retain success in the performance benchmarks indicated. This work would be useful in pre-training RL agents without expensive precision community models that can later be deployed in real communities without long learning periods.

References

[1] M. Ruth, A. Pratt, M. Lunacek, S. Mittal, H. Wu, and W. Jones, "Effects of home energy management systems on distribution utilities and feeders under various market structures," in *Proc. CIRED 23rd International Conference and Exhibition on Electricity Distribution*, 2015.

[2] K. Garifi, K. Baker, D. Christensen, and B. Touri, "Stochastic home energy management systems with varying controllable resources," in *IEEE Power Energy Society General Meeting (PESGM)*, 2019.

[3] Commonwealth Edison, "Real-time hourly prices," tech. rep., https://hourlypricing.comed.com/live-prices/, Last accessed May 2020.

[4] H. Lo, S. Blumsack, P. Hines, and S. Meyn, "Electricity rates for the zero marginal cost grid," *The Electricity Journal*, vol. 32, no. 3, pp. 39 – 43, 2019.

[5] CAISO, "What the duck curve tells us about managing a green grid." https://www.caiso.com/Documents/FlexibleResourcesHelpRenewables_FastFacts.pdf, Last accessed 11/16/2020.

[6] P. L. Denholm, J. Nunemaker, W. J. Cole, and P. J. Gagnon, "The potential for battery energy storage to provide peaking capacity in the united states,"

[7] R. Lu and S. H. Hong, "Incentive-based demand response for smart grid with reinforcement learning and deep neural network," *Applied Energy*, vol. 236, pp. 937 – 949, 2019.

[8] B. Kim, Y. Zhang, M. van der Schaar, and J. Lee, "Dynamic pricing and energy consumption scheduling with reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 7, no. 5, pp. 2187–2198, 2016.

[9] E. Foruzan, L. Soh, and S. Asgarpoor, "Reinforcement learning approach for optimal distributed energy management in a microgrid," *IEEE Transactions on Power Systems*, vol. 33, no. 5, pp. 5749–5758, 2018.

[10] X. Kou, F. Li, J. Dong, M. Starke, J. Munk, T. Kuruganti, and H. Zandi, "A distributed energy management approach for residential demand response," in *3rd Int. Conf. on Smart Grid and Smart Cities*, 2019.

[11] X. Jin, K. Baker, D. Christensen, and S. Isley, "Foresee: A user-centric home energy management system for energy efficiency and demand response," *Applied Energy*, vol. 205, pp. 1583 – 1595, 2017.

[12] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," 2018. [Online] Available: https://arxiv.org/abs/1801.01290.

[13] A. Hill *et al.*, "Stable baselines." https://github.com/hill-a/stable-baselines, 2018.

[14] H. Hasselt, "Double Q-learning," *Advances in Neural Information Processing Systems*, pp. 2613–2621, 2010.

[15] C. Mosiman and A. Pigott, "Distributed resource aggregation (DRAGG)." https://github.com/corymosiman12/dragg, 2020.

[16] "Heat pump water heater model validation study." https://ecotopewebstorage.s3.amazonaws.com/2015_001_1_HPWHModelVal.pdf, 2015. Prepared for the Northwest Energy Efficiency Alliance.

[17] M. Sengupta, Y. Xie, A. Lopez, A. Habte, G. Maclaurin, and J. Shelby, "The national solar radiation data base (NSRDB)," *Renewable and Sustainable Energy Reviews*, vol. 89, pp. 51 – 60, 2018.

[18] U.S. Department of Energy, "Sizing a new hot water heater." https://www.energy.gov/energysaver/water-heating/sizing-new-water-heater, Last accessed 11/5/2020.

[19] A. Faruqui, R. Hledik, S. Newell, and H. Pfeifenberger, "The power of 5 percent," *The Electricity Journal*, vol. 20, no. 8, pp. 68 – 77, 2007.