Using Meta Reinforcement Learning to Bridge the Gap between Simulation and Experiment in Energy Demand Response

Doseok Jang¹, Lucas Spangher¹, Manan Khattar¹, Utkarsha Agwan¹, Costas Spanos¹

1. Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
Berkeley, California, USA

ABSTRACT

Our team is proposing to run a full-scale energy demand response experiment in an office building. Although this is an exciting endeavor which will provide value to the community, collecting training data for the reinforcement learning agent is costly and will be limited. In this work, we apply a meta-learning architecture to warm start the experiment with simulated tasks, to increase sample efficiency. We present results that demonstrate a similar a step up in complexity still corresponds with better learning.

CCS CONCEPTS

• **Hardware** → **Energy distribution**; • **Computer systems organization** → *Neural networks*; • **Theory of computation** → **Reinforcement learning**.

KEYWORDS

prosumer, aggregation, reinforcement learning, microgrid, transactive energy

1 INTRODUCTION

The bridge from simulation to experiment is difficult to cross for both statistical and practical reasons. When considering experiments of phenomena in the energy grid, techniques to help a learned controller retain some information gained during simulation can be very beneficial for reducing these practical considerations.

We focus here on the electrical grid. As the grid decarbonizes, volatile resources like wind and solar will replace on-demand resources like fossil fuels, and there arises a mismatch between generation and demand. Grids that do not prepare for this question will face daunting consequences, from curtailment of resources [22] to voltage instability and physical damage, despite having adequate generative capability. Indeed, these problems will only grow larger as the world moves away from fossil fuels.

Demand response, a strategy in which customers are incentivized to shift their demand for energy resources to parts of the day where generation is plentiful, is seen as a common solution to the problem of generation volatility. Given the lack of needed material infrastructure and cheapness of the incentives, it has several positives above physical energy storage systems.

Building energy is a primary target of demand response, and both the central administration of signals and building-level response has been thoroughly studied in residential and industrial settings ([1], [15], [11], [28], [6].) However, while physical infrastructures [2] of office buildings have been studied for demand response ([7]), there has been no large scale experiment aimed at eliciting a behavioral demand response. The lack of an office-centered study is understandable when we consider that most offices do not have a mechanism to pass energy prices onto workers. If they did, however, not only would a fleet of decentralized batteries – laptops, cell phone chargers, etc. be able to be coordinated to function as a large deferrable resource, but building managers could save money by adapting their buildings' energy usage to a dynamic utility price.

The SinBerBEST collaboration has developed a Social Game [10] that facilitates workers to engage in a competition around energy [9], [3]. Through this framework, a first-of-its-kind experiment has been proposed to implement behavioral demand response within an office building [22]. Prior work has proposed to describe an hourly price-setting controller that learns how to optimize its prices [21]. However, given the costliness of iterations in this experiment and the work that has been put into building a complex simulation environment, warm-starting the experiment's controller with learning from the simulation could prove valuable to its success.

We report a simulated warm-started experiment. We will in Section 2 contextualize the architecture of our reinforcement learning controller within reinforcement learning and meta-learning. In Section 3 we will describe the simulation setup and the way we test the warm-started controller. In Section 4 we will give results. Finally, in Section 5 we will discuss implications of the controller and the future work this entails.

2 BACKGROUND

2.1 Reinforcement Learning

Reinforcement learning (RL) is a type of agent-based machine learning where control of a complex system requires actions that optimize the system [25], i.e. they seek to optimize the expected sum of rewards for actions (a_t) and states (s_t) in a policy parameterized by θ ; i.e., $J(\theta) = \mathbb{E}(\sum_{s_t, a_t \sim p_\pi} [r(s_t, a_t)])$. Often we model a policy as a deep neural network with weights θ that takes states as input and outputs actions. RL is useful in contexts where actions and environments are simple or data is plentiful, with early use cases being optimizing the control of backgammon [26], the cart-pole problem, and Atari [17]. Recently, RL has shown promise in solving more difficult problems like optimizing the control of Go beyond human capability [20]. In addition, much work has been done trying to extend RL to other situations.

This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

e-Energy '21, June 28–July 2, 2021, Virtual Event, Italy

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8333-2/21/06.

<https://doi.org/10.1145/3447555.3466589>

Policy gradient methods are a class of RL algorithms used to train policy networks that suggest actions (*actors*) via critic networks, trust regions, and other methods. A variant of these methods, Proximal Policy Optimization (PPO) algorithms ([19]), optimizes a surrogate loss

(1)

which enables multiple gradient updates to be done on these actors using the same samples. However, even with PPO optimization, several months worth of real-world training data would have to be collected to fully train an hourly price-setting controller [21] in our Social Game. We seek to leverage a detailed simulation with behaviorally reasonable dynamics encoded in a model that can train on both simulated and experimental environments. We thus propose the use of Model Agnostic Meta Learning, a few-shot learning algorithm which trains an RL meta-architecture that learns to warm start a model's parameters as it trains over a distribution of similar tasks. We propose a MAML optimization procedure over a PPO agent architecture.

2.2 Model Agnostic Meta-Learning (MAML)

Model Agnostic Meta-Learning (MAML) ([4]) is a meta-learning algorithm that seeks to learn a neural network weight initialization θ that allows for fast adaptation to different types of tasks. We first define a task as a transition distribution $q_i(x_{t+1}|x_t, a_t)$ and a reward function R_i . MAML can be implemented over any model. In the RL variant, given a set of tasks τ and a distribution of those tasks $p(\tau)$ MAML seeks to find an initialization θ that can achieve as high a reward as possible after K gradient update steps in any task. To maximize reward after K updates, MAML optimizes a policy network, f_ϕ for K gradient update steps starting from the learned initialization θ and computes the following meta loss:

$$L_{\tau_i}(f_\theta) = -\mathbb{E}_{x_t, a_t \sim f_\phi, q_{\tau_i}} \left[\sum_{t=1}^H R_i(x_t, a_t) \right] \quad (2)$$

MAML then optimizes θ by gradient descent on the sum of the meta-losses over all tasks, computing parameters θ that produce good results after just K update steps. Training MAML on PPO could be described as consisting of an inner adaptation step and an outer, meta adaptation step. In the inner step, PPO is trained on a task τ_i , randomly chosen according to $p(\tau)$ for K PPO update steps, starting from the weight initialization MAML is optimizing. In the outer step, the weight initialization is updated according to the training trajectories of PPO that were collected in the inner adaptation steps. These outer and inner adaptation steps are alternated to compute the weight initialization that produce the highest reward after PPO is trained starting from it for K steps, on a task from the task distribution τ . We observe in this paper the performance of a trained MAML+PPO weight initialization on several tasks that are outside of the task distribution τ but are still within our simulated Social Game setting.

2.3 RL in Demand Response

RL has been applied to a number of demand response situations, but almost all work centers on agents that directly schedule resources [14], [12], [29], [27], [18], [5]. RL architectures can vary widely, for example Kofinas et. al. deploy a fuzzy Q-learning multi-agent

Figure 1: Reinforcement Learning Control Flow

that learns to coordinate appliances to increase reliability [8]. In another example, Mbuwir et. al. manage a battery directly using batch Q-learning [16].

However, there are few works where the RL controller is a price setter for demand response in an office. To the best of our knowledge, no one has used MAML to improve the learning of the RL controller over a vanilla controller for this use. We are proposing this agent for use in a novel experiment with demand response within a social game.

3 METHODS

We adapt MAML to the problem of optimizing a price-serving agent for energy demand response. We will briefly explain the baseline PPO architecture and then describe the simulation environment we test this in. Finally, we will then explain the MAML implementation and how we tested it in the environment.

Can an RL agent preemptively estimate the most effective demand response price using historical data and implicitly predicting causal factors? Furthermore, can we pretrain an agent in simulation that can quickly adapt to real-world data? We employ Proximal Policy Optimization (PPO) and Model-Agnostic Meta-Learning (MAML) to train an agent in OfficeLearn to optimize our policy network to adapt quickly to new environments. We use the RLLib [13] implementation of PPO and MAML, and their default neural network architecture of two hidden layers with 256 units each and tanh activations. The reward for the price-setting agent is $-\log(d^t g) - \lambda * \mathbb{1}(d^t g - \hat{d})$, where d is the demand of the person it studies, g is the grid pricing, $\hat{d}$ is minimal baseline energy demand that the grid must meet, and $\lambda * \mathbb{1}(d^t g - \hat{d})$ is a regularizing penalty that is applied if the simulated energy grid does not meet $\hat{d}$. The penalty helps avoid the locally optimal solution of the agent driving down prices indiscriminately without regard for energy supply. For other implementation choices, please see our Github¹ and RLLib.

3.1 Environment

We summarize an OpenAI gym environment built to simulate demand response in office buildings [23]. Each step in the environment is a day, where the agent proposes prices to office workers. Based on the agent's prices, the office workers modify their energy consumption behaviors in order to achieve the lowest cost of energy possible, which the controller then uses to assign points. Each simulated person has a deterministic response to the points that are offered to them. Notably, we employ several different models of simulated

¹Please find our Github here: https://github.com/Aphoh/temp_tc/tree/austin_maml

response, with two levels of complexity: “Deterministic Function” and “Curtail and Shift”. Their descriptions are listed below:

3.1.1 “Deterministic Function” Person. We include three types of deterministic response within one type of agent, with the option of specifying a mixed office composed of all three types.

A “Deterministic Function” Person with **linear response** decreases their energy consumption linearly below an average historical energy consumption baseline. Therefore, if b_t is the historical baseline at time t and p_t are the points given and m is a simulation set points multiplier, the energy demand d at time t is $d_t = b_t - p_t * m$, clipped at ceiling and floor values d_{min} and d_{max} , which are 5% and 95% of a historical energy distribution (i.e. a distribution made from two months worth of historical data from [24].) A “Deterministic Function” Person with **sinusoidal response** is one who responds to points towards the middle of the distribution and not well to low or high points. Therefore, the energy demand d at time t is $d_t = b_t - \sin p_t * m$, clipped at d_{min} and d_{max} .

A “Deterministic Function” Person with **threshold exponential response**, we define an office worker who does not respond until points pass a threshold, at which point they respond exponentially. Therefore, the energy demand d is $d_t = b_t - (\exp p_t * (p_t > 5))$, clipped at d_{min} and d_{max} .

3.1.2 “Curtail And Shift Office Worker”. Office workers need to use electricity to do their work, and may be unable to curtail their load below a minimum threshold, e.g. the power needed to run a PC. They may have the ability to shift their load over a definite time interval, e.g. choosing to charge their laptops ahead of time or at a later time. We model a response function that exhibits these behaviors. We can model the aggregate load of a person (b_t) as a combination of fixed demand (b_t^{fixed}), curtailable demand ($b_t^{curtail}$), and shiftable demand (b_t^{shift}), i.e., $b_t = b_t^{fixed} + b_t^{curtail} + b_t^{shift}$. All of the curtailable demand is curtailed for the $T_{curtail}$ hours (set to 3 hours in practice) with the highest points, and for every hour t the shiftable demand is shifted to the hour within $[t - T_{shift}, t + T_{shift}]$ with the lowest energy price. For example, such an office worker may need to charge their appliances for a total of 1000 Wh throughout the day, 300 of which are for printing documents that could be curtailed, 300 of which are for presenting at a meeting whose time can be shifted, and 400 of which are the minimum required to run a PC. Upon receiving a price signal with high prices from 11am-2pm, this simulated worker would curtail their printing away from the 3 hours with the highest energy price, and schedule their meeting at the hour within $[t - T_{shift}, t + T_{shift}]$ with the lowest energy price. This would decrease their energy usage.

3.2 Simulated Experiment

To test our hypothesis that MAML + PPO will enable faster adaptation to unfamiliar environments like a real-world Social Game, we trained MAML on several models of simulated person response. We then evaluated how quickly PPO, starting from the MAML weight initialization, can learn in an OfficeLearn environment with different models of simulated person response. The training environments had randomized “Deterministic Function” response types and multipliers for how many “points” simulated humans received for reducing energy usage. Though the training environments used

to train MAML had randomized parameters, the validation environments were kept constant to ensure fairness. To ensure an accurate representation of each network’s capabilities, we averaged the results from 5 different test trials and report the mean and standard error for each test. MAML is trained with an ADAM optimizer with learning rate 0.0001, 0.9 β_1 , and 0.999 β_2 , where β_1 refers to the first moment and β_2 refers to the second. MAML sampled approximately eight PPO training trajectories in (parallel) training environments at a time between (sequential) meta-update steps. We trained PPO for 5 steps at a time in MAML’s inner adaptation phase and trained MAML once per trajectory sampling. We trained MAML for up to 200 iterations. For evaluation, we trained PPO for 100 steps from the MAML-learned weight initialization and from a random initialization and compared the two. PPO (in both MAML+PPO and PPO) is trained with the clipped surrogate loss with a clipping parameter of 0.3 and a SGD optimizer with learning rate 0.01. The action and value estimators shared layers in the neural network for our implementation, and the value loss had a weighting of 0.5. We present three different experiments:

3.2.1 Adaptation to “Curtail and Shift” Response. We train MAML on a distribution of environments with three different models of human response to price: “Deterministic Function” with all responses, and see if it can adapt to a different model that we believe may be more representative of real-world behavior: “Curtail and Shift”, where testing was done over 5 different test trials of each.

3.2.2 Adaptation to Threshold Exponential Response. In this experiment, we train MAML on a distribution of environments with two different models of human response to price instead of three: “Deterministic Function” with linear or sinusoidal responses. Each variant environment, along with other parameters, are chosen at random when MAML samples from the task distribution during training. We test this trained model’s ability to adapt to an environment with a different model of human response: “Deterministic Function” Person with threshold exponential response, with all 5 validation runs of MAML+PPO and PPO.

3.2.3 MAML Training Behavior. We observe the performance of MAML + PPO in the *Adaptation to Curtail and Shift Response* task after different numbers of MAML training iterations. When training MAML, we saved checkpoints of the meta-optimized weight initialization at 50, 100, 150, and 200 MAML iterations and observed the performance of PPO starting from each of those warm starts.

4 RESULTS

We will now describe the results obtained from using MAML.

4.1 MAML Results (Compared to baseline)

As can be seen in Fig. 2.A and B, MAML + PPO significantly outperforms baseline PPO from the start to the finish of training. In both adaptation experiments, MAML+PPO converges to a solution with higher reward, and thus lower simulated energy cost than PPO. Note that PPO appears to have plateaued with constant reward. MAML decreases total energy cost by about 40% in both environments compared to PPO by the end of training on average. It takes fewer than 2 PPO steps worth of training data for MAML+PPO to outperform PPO after training for 100 steps, which suggests that the

Figure 2: MAML+PPO Results

A. Performance of MAML+PPO in training on "Deterministic Function", and adapting to a more complex model, "Curtail and Shift" in comparison to PPO. **B.** Performance of MAML+PPO in a lateral shift in complexity, i.e. adapting to a different model of human behavior, "Deterministic Threshold Exponential" in comparison to PPO.

use of MAML with PPO to warm start learning from our simulation was successful. Using MAML allowed us to find a weight initialization that generalized to more complex tasks, thus decreasing the amount of data necessary to train the controller.

As we are interested in both the reward achieved by each algorithm and how quickly the reward is achieved, we chose to plot the training graphs of the validation runs of each checkpoint in Fig. 3 instead of simply plotting the total reward of each validation run. We observe that performance appears to increase with number of MAML training steps up until 200 iterations, where MAML appears to have overfit to a local minimum. Training MAML for too few steps, for example only 50 iterations, appears to increase instability in the PPO training on the MAML weight initialization.

5 DISCUSSION

Meta-learning and K-shot learning are fields of machine learning focused on applying knowledge obtained from one domain to another. This makes techniques like MAML especially helpful in taking advantage of simulations, where data is plentiful, to increase performance in the real world, where data is costly. Whereas other RL tasks like robotic grasping can use highly accurate physics simulations to train a model and successfully employ the same model to real world equivalent tasks, the human response to changes in electricity price is not as well studied as physics. The models of human behavior we employ in simulations are simplistic and likely inaccurate with regard to real human behaviour. However, the novelty of our experiment is in showing how MAML might be

Figure 3: MAML Reward vs Update Steps

Performance of MAML+PPO on "Curtail and Shift" task after different number of iterations of MAML in simulation.

applied to transfer knowledge acquired from training in simulation to accelerating adaptation in the real world, by demonstrating that MAML allows fast adaptation to different models of human responses to energy price. From our results, we are hopeful that MAML+PPO can be applied in real world price controller RL as well. By training MAML+PPO in an environment different from the test environment, we were able to affirm that the technique can be used to adapt training in simulation to a different task in simulation. The jump in complexity between the "Deterministic Function" response the model was trained on and the "Curtail and Shift" response shows that MAML+PPO can be used to adapt to tasks that are outside of its training task distribution and are more complex. We believe that this ability for MAML to close gaps in complexity has a high probability of enabling a MAML optimized model to work as a real world RL price controller.

5.1 Future Research

We hope to integrate MAML with other enhancements that we are testing in the learning controller: namely, a Surprise-Minimizing RL reward regularizer, detailed in a different submission to this same conference, and a multi-agent architecture of TD3. We would like to test this in simulation.

More practically, we will use this enhancement directly in the rollout of our agent for our office-scale demand response experiment. We present in this work a training on a distribution of tasks defined by simpler people and tested on more complex people, but we will essentially train the experiment controller on the superset of the presented training and validation simulations, and test it on real people. This will allow us to confirm our hypothesis that MAML allows for faster adaptation from simulation to the real world when training an RL price controller.

5.2 Acknowledgements

We gratefully acknowledge the input of Manan Khattar, William Arnold, and Tarang Srivastava as working group collaborators who met weekly on their own research thrusts and provided advice (and at times a listening ear.) We give thanks to Andreea Bobu and Peter Hendrickson who are graduate students that gave guidance during this process. Finally, we would like to thank Chelsea Finn for her initial work on MAML and Pieter Abbeel and Sergey Levine for advice throughout the process.

BIBLIOGRAPHY

- [1] Ailin Asadinejad, Alireza Rahimpour, Kevin Tomsovic, Hairong Qi, and Chien-fei Chen. 2018. Evaluation of residential customer elasticity for incentive based demand response programs. *Electric Power Systems Research* 158 (2018), 26–36.
- [2] Hari Prasanna Das, Ioannis Konstantakopoulos, Aummul Baneen Manasawala, Tanya Veeravalli, Huihan Liu, and Costas J Spanos. 2020. Do Occupants in a Building exhibit patterns in Energy Consumption? Analyzing Clusters in Energy Social Games. (2020).
- [3] Hari Prasanna Das, Ioannis C Konstantakopoulos, Aummul Baneen Manasawala, Tanya Veeravalli, Huihan Liu, and Costas J Spanos. 2019. A novel graphical lasso based approach towards segmentation analysis in energy game-theoretic frameworks. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, 1702–1709.
- [4] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*. PMLR, 1126–1135.
- [5] Danilo Fuselli, Francesco De Angelis, Matteo Boaro, Stefano Squartini, Qinglai Wei, Derong Liu, and Francesco Piazza. 2013. Action dependent heuristic dynamic programming for home energy resource scheduling. *International Journal of Electrical Power & Energy Systems* 48 (2013), 148–160. <https://doi.org/10.1016/j.ijepes.2012.11.023>
- [6] Brandon J Johnson, Michael R Starke, Omar A Abdelaziz, Roderick K Jackson, and Leon M Tolbert. 2015. A dynamic simulation tool for estimating demand response potential from residential loads. In *2015 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. IEEE, 1–5.
- [7] Y. Kim. 2018. Optimal Price Based Demand Response of HVAC Systems in Multi-zone Office Buildings Considering Thermal Preferences of Individual Occupants Buildings. *IEEE Transactions on Industrial Informatics* 14, 11 (2018), 5060–5073. <https://doi.org/10.1109/TII.2018.2790429>
- [8] P. Kofinas, A.I. Dounis, and G.A. Vouras. 2018. Fuzzy Q-Learning for multi-agent decentralized energy management in microgrids. *Applied Energy* 219 (2018), 53–67. <https://doi.org/10.1016/j.apenergy.2018.03.017>
- [9] Ioannis C Konstantakopoulos, Andrew R Barkan, Shiyang He, Tanya Veeravalli, Huihan Liu, and Costas Spanos. 2019. A deep learning and gamification approach to improving human-building interaction and energy efficiency in smart infrastructure. *Applied energy* 237 (2019), 810–821.
- [10] Ioannis C Konstantakopoulos, Hari Prasanna Das, Andrew R Barkan, Shiyang He, Tanya Veeravalli, Huihan Liu, Aummul Baneen Manasawala, Yu-Wen Lin, and Costas J Spanos. 2019. Design, benchmarking and explainability analysis of a game-theoretic framework towards energy efficiency in smart infrastructure. *arXiv preprint arXiv:1910.07899* (2019).
- [11] Chaojie Li, Chen Liu, Xinghuo Yu, Ke Deng, Tingwen Huang, and Liangchen Liu. 2018. Integrating Demand Response and Renewable Energy In Wholesale Market.. In *IJCAL* 382–388.
- [12] D. Li and S. K. Jayaweera. 2014. Reinforcement learning aided smart-home decision-making in an interactive smart grid. In *2014 IEEE Green Energy and Systems Conference (IGESC)*. 1–6. <https://doi.org/10.1109/IGESC.2014.7018632>
- [13] Eric Liang, Richard Liaw, Robert Nishihara, Philipp Moritz, Roy Fox, Ken Goldberg, Joseph Gonzalez, Michael Jordan, and Ion Stoica. 2018. RLlib: Abstractions for distributed reinforcement learning. In *International Conference on Machine Learning*. PMLR, 3053–3062.
- [14] Y. Liu, C. Yuen, N. Ul Hassan, S. Huang, R. Yu, and S. Xie. 2015. Electricity Cost Minimization for a Microgrid With Distributed Energy Resource Under Different Information Availability. *IEEE Transactions on Industrial Electronics* 62, 4 (2015), 2571–2583. <https://doi.org/10.1109/TIE.2014.2371780>
- [15] Kai Ma, Guoqiang Hu, and Costas J Spanos. 2015. A cooperative demand response scheme using punishment mechanism and application to industrial refrigerated warehouses. *IEEE Transactions on Industrial Informatics* 11, 6 (2015), 1520–1531.
- [16] Brida V. Mbuwir, Frederik Ruelens, Fred Spiessens, and Geert Deconinck. 2017. Battery Energy Management in a Microgrid Using Batch Reinforcement Learning. *Energies* 10, 11 (2017). <https://www.mdpi.com/1996-1073/10/11/1846>
- [17] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [18] Leo Raju, Sibi Sankar, and R.S. Milton. 2015. Distributed Optimization of Solar Micro-grid Using Multi Agent Reinforcement Learning. *Procedia Computer Science* 46 (2015), 231–239. <https://doi.org/10.1016/j.procs.2015.02.016> Proceedings of the International Conference on Information and Communication Technologies, ICICT 2014, 3-5 December 2014 at Bolgatty Palace & Island Resort, Kochi, India.
- [19] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [20] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529, 7587 (2016), 484–489.
- [21] Lucas Spangher, Akash Gokul, Manan Khattar, Joseph Palakapilly, Utkarsha Agwan, Akaash Tawade, and Costas Spanos. 2020. Augmenting Reinforcement Learning with a Planning Model for Optimizing Energy Demand Response. In *Proceedings of the 1st International Workshop on Reinforcement Learning for Energy Management in Buildings & Cities*. 39–42.
- [22] Lucas Spangher, Akash Gokul, Manan Khattar, Joseph Palakapilly, Akaash Tawade, Adam Bouyamourn, Alex Devonport, and Costas Spanos. 2020. Prospective experiment for reinforcement learning on demand response in a social game framework. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems*. 438–444.
- [23] Lucas Spangher, Akash Gokul, Joseph Palakapilly, Utkarsha Agwan, Manan Khattar, Wann-Jiun Ma, and Costas Spanos. [n.d.]. OfficeLearn: An OpenAI Gym Environment for Reinforcement Learning on Occupant-Level Building's Energy Demand Response. In *Tackling Climate Change with Artificial Intelligence Workshop at NeurIPS, 2020*.
- [24] Lucas Spangher, Akaash Tawade, Alex Devonport, and Costas Spanos. 2019. Engineering vs. Ambient Type Visualizations: Quantifying Effects of Different Data Visualizations on Energy Consumption. In *Proceedings of the 1st ACM International Workshop on Urban Building Energy Sensing, Controls, Big Data Analysis, and Visualization (New York, NY, USA) (UrbSys'19)*. Association for Computing Machinery, New York, NY, USA, 14–22. <https://doi.org/10.1145/3363459.3363527>
- [25] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [26] Gerald Tesauro. 1994. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural computation* 6, 2 (1994), 215–219.
- [27] Q. Wei, D. Liu, and G. Shi. 2015. A novel dual iterative Q-learning method for optimal battery management in smart residential environments. *IEEE Transactions on Industrial Electronics* 62, 4 (2015), 2509–2518. <https://doi.org/10.1109/TIE.2014.2361485>
- [28] Ji Hoon Yoon, Ross Baldick, and Atila Novoselac. 2014. Dynamic demand response controller based on real-time retail price for residential buildings. *IEEE Transactions on Smart Grid* 5, 1 (2014), 121–129.
- [29] Y. Zhang and M. van der Schaar. 2014. Structure-aware stochastic load management in smart grids. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*. 2643–2651. <https://doi.org/10.1109/INFOCOM.2014.6848212>