



CityLearn v2: An OpenAI Gym environment for demand response control benchmarking in grid-interactive communities

Kingsley Nweye
The University of Texas at Austin
Austin, Texas, USA
nweye@utexas.edu

Kathryn Kaspar
Concordia University
Montreal, Quebec, Canada
kathryn.kaspar@mail.concordia.ca

Giacomo Buscemi
Politecnico di Torino
Torino, Italy
giacomo.buscemi@polito.it

Giuseppe Pinto
PassiveLogic
Amsterdam-Centrum
Noord-Holland, Nederland
giuseppe-pinto@polito.it

Han Li
Lawrence Berkeley National
Laboratory
Berkeley, California, USA
hanli@lbl.gov

Tianzhen Hong
Lawrence Berkeley National
Laboratory
Berkeley, California, USA
thong@lbl.gov

Mohamed Ouf
Concordia University
Montreal, Quebec, Canada
mohamed.ouf@concordia.ca

Alfonso Capozzoli
Politecnico di Torino
Torino, Italy
alfonso.capozzoli@polito.it

Zoltan Nagy
The University of Texas at Austin
Austin, Texas, USA
nagy@utexas.edu

ABSTRACT

As more distributed energy resources become part of the demand-side infrastructure, it is important to quantify the energy flexibility they provide, as well as identify the best control strategies to accelerate their real-world adoption. CityLearn provides an environment for benchmarking of simple and advanced control algorithms in virtual grid-interactive communities. The updated CityLearn v2 environment introduced here extends the v1 environment to provide load shedding flexibility through heating ventilation and air conditioning power control coupled with a data-driven temperature dynamics model. The updated environment also includes the functionality to assess the resiliency of control algorithms during power outage events.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

KEYWORDS

reinforcement learning, model predictive, grid resilience, energy management

ACM Reference Format:

Kingsley Nweye, Kathryn Kaspar, Giacomo Buscemi, Giuseppe Pinto, Han Li, Tianzhen Hong, Mohamed Ouf, Alfonso Capozzoli, and Zoltan Nagy. 2023. CityLearn v2: An OpenAI Gym environment for demand response control benchmarking in grid-interactive communities. In *The 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and*

Transportation (BuildSys '23), November 15–16, 2023, Istanbul, Turkey. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3600100.3626257>

1 INTRODUCTION

The shift from fossil-fuel energy sources to renewable energy sources (RESs) and the electrification of end-uses are pathways towards climate change mitigation, but risk creating new challenges for the electricity grid if mismanaged. The intermittency of RESs could undermine grid resilience due to the mismatch between electricity generation and end-use demand [11]. Furthermore, extreme-weather events such as heat waves and winter storms can cause both increase in demand and decrease in supply due to power outages.

While distributed energy resources (DERs) are able to provide load shifting and shedding flexibility [7], it is challenging to control DERs to serve diverse occupant behaviors and coordinate multiple resources in many buildings. Advanced control algorithms such as model predictive control (MPC) [4] and reinforcement learning control (RLC) [6] can effectively manage DERs, by adapting to unique building characteristics while, cooperating towards grid-level energy flexibility and resiliency objectives. As more DERs become part of the demand-side infrastructure, it is important to quantify the energy flexibility that they provide, as well as identify best control strategies to accelerate demand response (DR) program adoption.

CityLearn, is an open-source Gym environment for the easy implementation and benchmarking of simple, e.g., rule-based control (RBC) and advanced control algorithms for DR in grid-interactive communities [12]. It has been applied in voltage regulation [9], transfer learning [8], and meta-reinforcement learning [13] problems. In contrast to its alternatives e.g., ACTB [5] and BOPTEST [3], it supports district-level and multi-agent control and does not require a compute-intensive co-simulation engine by using simplified first-order energy models compared to DOPTEST [2]. Here, we introduce the CityLearn v2 environment that improves on the v1



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

BuildSys '23, November 15–16, 2023, Istanbul, Turkey
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0230-3/23/11.
<https://doi.org/10.1145/3600100.3626257>

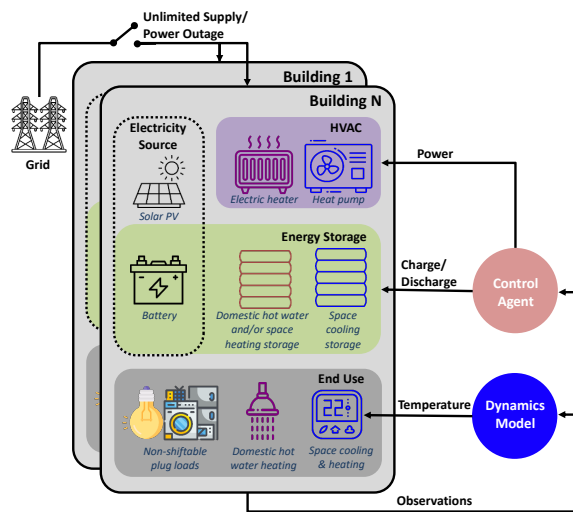


Figure 1: CityLearn v2 environment and control interaction.

environment to provide load shedding flexibility through heating ventilation and air conditioning power control coupled with a data-driven temperature dynamics model. Additionally, our updated environment includes the functionality to assess the resiliency of control algorithms during power outage events.

2 CITYLEARN

The CityLearn environment (Fig. 1) includes simplified energy models of buildings that contain heating ventilation and air conditioning (HVAC) systems (heat pumps and electric heaters) and energy storage systems (ESSs) storage. Each building's space cooling, space heating and domestic hot water (DHW) heating end-use loads are independently satisfied through air-to-water heat pumps. Alternatively, electric heaters in place of heat pumps can be used to satisfy space and DHW heating loads. ESSs are charged by the HVAC system that satisfies the end-use that the stored energy services. All HVAC systems as well as plug loads consume electricity from any of the available electricity sources including the grid, photovoltaic (PV) system, and battery. RBCs, RLC or MPC agent(s) manage load shifting in the buildings by determining how much energy to store or release at any given time. The control architecture is either one agent to many buildings (centralized) or one agent to one building (decentralized) with optional information sharing amongst agents.

With CityLearn v2¹, the agents can also, control the available power from the HVAC system to shed space thermal loads, preheat, or precool the building. The consequence of the controlled power on the indoor dry-bulb temperature, i.e., the building dynamics, is modeled using a long short-term memory (LSTM) surrogate model based on the work by Pinto et al. [10].

CityLearn v2 also provides the functionality to assess the resiliency of control algorithms during power outage events. CityLearn v2 provides a stochastic power outage model based on System Average Interruption Frequency Index (SAIFI) and Customer Average Interruption Duration Index (CAIDI) distribution system reliability

metrics [1] such that it is generalizable by providing custom or location-specific metrics.

3 CONCLUSION

We introduce the CityLearn v2 environment in this short paper that extends the control capabilities of the previous v1 environment as well provides support for power outage simulation. This new environment contributes to the limited tools for district-level and multi-agent control benchmarking and energy flexibility assessment of a diverse set of distributed energy resources.

REFERENCES

- [1] 2012. IEEE Guide for Electric Power Distribution Reliability Indices. <https://doi.org/10.1109/IEEESTD.2012.6209381> Conference Name: IEEE Std 1366-2012 (Revision of IEEE Std 1366-2003).
- [2] Javier Arroyo, Lucas Verleyen, Lucas Bex, Louis Hermans, Muhammad Saeed, Yucun Lu, Joris Depoortere, Attila Bálint, Erik Delarue, Johan Driesen, Geert Deconinck, David Blum, and L. Helsen. 2023. *Prototyping the DOPTTEST Framework for Simulation-Based Testing of System Integration Strategies in Districts*.
- [3] David Blum, Javier Arroyo, Sen Huang, Ján Dragoňa, Filip Jorissen, Harald Taxt Walnum, Yan Chen, Kyle Benne, Draguna Vrabie, Michael Wetter, and Lieve Helsen. 2021. Building optimization testing framework (BOPTTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation* 14, 5 (Sept. 2021), 586–610. <https://doi.org/10.1080/19401493.2021.1986574>
- [4] Ján Dragoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L. Vrabie, and Lieve Helsen. 2020. All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50, May (2020), 190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>
- [5] Thibault Marzullo, Sourav Dey, Nicholas Long, José Leiva Vilaplana, and Gregor Henze. 2022. A high-fidelity building performance simulation test bed for the development and evaluation of advanced controls. *Journal of Building Performance Simulation* 15, 3 (2022), 379–397. <https://doi.org/10.1080/19401493.2022.2058091> Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/19401493.2022.2058091>
- [6] Zoltan Nagy, Gregor Henze, Sourav Dey, Javier Arroyo, Lieve Helsen, Xiangyu Zhang, Bingqing Chen, Kadir Amasyali, Kuldeep Kurte, Ahmed Zamzam, Helia Zandi, Ján Dragoňa, Matias Quintana, Steven McCullogh, June Young Park, Han Li, Tianzhen Hong, Silvio Brandi, Giuseppe Pinto, Alfonso Capozzoli, Draguna Vrabie, Mario Bergés, Kingsley Nweye, Thibault Marzullo, and Andrey Bernstein. 2023. Ten questions concerning reinforcement learning for building energy management. *Building and Environment* 241 (Aug. 2023), 110435. <https://doi.org/10.1016/j.buildenv.2023.110435>
- [7] Monica Neukomm, Valerie Nubbe, and Robert Fares. 2019. Grid-Interactive Efficient Buildings. (April 2019). <https://doi.org/10.2172/1508212>
- [8] Kingsley Nweye, Siva Sankaranarayanan, and Zoltan Nagy. 2023. MERLIN: Multi-agent offline and transfer learning for occupant-centric operation of grid-interactive communities. *Applied Energy* 346 (Sept. 2023), 121323. <https://doi.org/10.1016/j.apenergy.2023.121323>
- [9] Aisling Pigott, Constance Crozier, Kyri Baker, and Zoltan Nagy. 2022. GridLearn: Multiagent reinforcement learning for grid-aware building energy management. *Electric Power Systems Research* 213 (2022), 108521. <https://doi.org/10.1016/j.ejpsr.2022.108521>
- [10] Giuseppe Pinto, Davide Deltetto, and Alfonso Capozzoli. 2021. Data-driven district energy management with surrogate models and deep reinforcement learning. *Applied Energy* 304 (Dec. 2021), 117642. <https://doi.org/10.1016/j.apenergy.2021.117642>
- [11] Mohammed Yekini Suberu, Mohd Wazir Mustafa, and Nouruddeen Bashir. 2014. Energy storage systems for renewable energy power sector integration and mitigation of intermittency. *Renewable and Sustainable Energy Reviews* 35 (2014), 499–514. <https://doi.org/10.1016/j.rser.2014.04.009>
- [12] J.R. José R. Vázquez-Canteli, Jérôme Kämpf, Gregor Henze, and Zoltan Nagy. 2019. CityLearn v1.0: An OpenAI gym environment for demand response with deep reinforcement learning. *BuildSys 2019 - Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation* (2019), 356–357. <https://doi.org/10.1145/3360322.3360998> ISBN: 9781450370059.
- [13] Huiliang Zhang, Di Wu, and Benoit Boulet. 2022. MetaEMS: A Meta Reinforcement Learning-based Control Framework for Building Energy Management System. (2022). <https://doi.org/10.48550/ARXIV.2210.12590>

¹<https://www.citylearn.net>