

A framework for the design of representative neighborhoods for energy flexibility assessment in CityLearn

Original

A framework for the design of representative neighborhoods for energy flexibility assessment in CityLearn / Nweye, Kingsley; Kaspar, Kathryn; Buscemi, Giacomo; Pinto, Giuseppe; Li, Han; Hong, Tianzhen; Ouf, Mohamed; Capozzoli, Alfonso; Nagy, Zoltan. - In: BUILDING SIMULATION CONFERENCE PROCEEDINGS. - ISSN 2522-2708. - ELETTRONICO. - 18:(2023), pp. 1814-1821. (18th Conference of International Building Performance Simulation Association Shangai (China) 4-6 September) [10.26868/25222708.2023.1404].

Availability:

This version is available at: 11583/2995375 since: 2024-12-16T17:17:57Z

Publisher:

IBPSA

Published

DOI:10.26868/25222708.2023.1404

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

A framework for the design of representative neighborhoods for energy flexibility assessment in CityLearn

Kingsley Nweye¹, Kathryn Kaspar², Giacomo Buscemi³, Giuseppe Pinto⁴, Han Li⁵,
Tianzhen Hong⁵, Mohamed Ouf², Alfonso Capozzoli³, Zoltan Nagy¹

¹The University of Texas at Austin, Austin, Texas, USA

²Concordia University, Montreal, Quebec, Canada

³Politecnico di Torino, Torino, Italy

⁴PassiveLogic, Salt Lake City, Utah, USA

⁵Lawrence Berkeley National Laboratory, Berkeley, California, USA

Abstract

The electricity grid in the U.S. continues to undergo system-wide changes due to electrification of end-uses, as well as the adoption of intermittent on-site renewable energy systems. When effectively managed, distributed energy resource in buildings can provide energy flexibility to alleviate grid loads at critical periods. However, it is important to understand the role of geographic, climatic and occupant behavioral differences on the effectiveness of their flexibility. Thus, we provide a framework that uses an open-source U.S. building-stock database with clustering techniques to design representative neighborhoods for distributed energy resource control algorithm benchmarking. We demonstrate an application in three neighborhoods that use reinforcement learning control for energy storage system management and simulation results show up to 42% reduction in peak demand, amongst other key performance indicators.

Highlights

- Design of representative residential neighborhoods.
- Synthetic datasets creation for building control algorithm benchmarking.
- Quantification of energy flexibility for representative neighborhoods.

Introduction

The United States (U.S.) electricity grid is undergoing significant changes due to building and transportation electrification, unpredictable electricity demand, and increased adoption of renewable energy systems (RESs) (U.S. Department of Energy, 2021). Originally, the grid served as a centralized energy provider while buildings remained strictly energy consumers. However, buildings now play an active role in the electricity supply market by integrating on-site RESs and energy storage systems (ESSs). Predicting building energy demand has become more complex due to the proliferation of energy-intensive appliances such as heat pumps, electric heaters, washing machines, dishwashers, and electric vehicle (EV) charging stations in the residential sector. At the urban scale, this leads to

high energy demand during peak hours of 6-9 am and 6-10 pm. Furthermore, extreme-weather events such as heat waves, winter storms, and hurricanes can cause both increases in demand as well as decreases in supply due to power outages.

Distributed energy resources (DERs) including RESs, ESSs, and heat pumps can provide building energy flexibility in response to grid signals, changing weather conditions, or occupant preferences (Jensen et al., 2017). Temporary load shedding, shifting load to off-peak hours, load modulation, and management of renewable energy generation are several ways in which buildings can exhibit their flexibility by providing grid services during demand response (DR) events (Neukomm et al., 2019). It is, however, challenging to coordinate DERs in multiple buildings to ensure efficient and flexible operation. Advanced control algorithms such as model predictive control (MPC) (Drgoña et al., 2020) and reinforcement learning control (RLC) (Wang and Hong, 2020; Pinto et al., 2021) provide a solution to this problem by effectively managing DERs, adapting to unique building characteristics, and cooperating towards achieving grid-level objectives.

As more DERs are integrated into the demand-side infrastructure, quantifying flexibility capabilities of the existing building stock as well as identifying best control strategies to accelerate the design and adoption of DR programs are crucial. Particularly, understanding the impact of geographic, climatic, and occupant behavioral differences on DER effectiveness in providing grid services can inform building design choices and guide policy makers. However, to carry out such an analysis at the urban scale, an inventory of building stock energy models is required. Efforts have been made to provide archetypes for building stock energy modeling, including the recent End-Use Load Profiles (EULP) for the U.S. Building Stock database (Wilson et al., 2022) generated from ResStock (Wilson, 2017) and ComStock (Horsey et al., 2020). This database uses EnergyPlus physics-based simulation models to provide 900,000 synthetic buildings representative of the residential and commercial building stock in the

U.S.. With this comprehensive dataset, residential, commercial, and mixed-use neighborhoods can be created to study the impact of DERs and DR programs.

Physics-based models accurately capture building thermodynamics and are the industry standard for retrofit assessment. However, they require domain knowledge and effort to design controllable systems and integrate standardized advanced control libraries for co-simulation and algorithm benchmarking. Emulators provide various levels of abstractions of the controlled systems (environment), which enable the designer to focus on control aspects with fewer resources spent on designing the environment. Such emulators include Gym-Eplus (Zhang and Lam, 2018), an EnergyPlus co-simulation OpenAI Gym environment that allows for the evaluation of deep RLC algorithms, and BOPTTEST (Blum et al., 2021), which provides high-fidelity energy models for control algorithm benchmarking. Although these emulators take advantage of robust simulation engines such as EnergyPlus and Mod-elica, their dependency on such engines raises the entry level for users. Moreover, they are designed for specific system-level or building-level environments and do not allow for district or neighborhood-level control or objectives.

CityLearn is an open-source OpenAI Gym environment for the easy implementation and benchmarking of reinforcement learning (RL) and MPC algorithms in a demand response setting (Vázquez-Canteli et al., 2019). CityLearn is used to reshape the aggregated curve of electricity demand by controlling the energy storage of a diverse set of buildings in a district, thus unlocking the potential for cooperative measures amongst the available DERs in a neighborhood.

However, to the best of the authors' knowledge, existing literature does not provide a framework that takes advantage of the availability of large building stock database for comprehensive control algorithm benchmarking. The contribution of this work is thus to provide a framework for building stock control algorithm benchmarking where the EULP database is used as a data source for CityLearn environment design. We prescribe procedures to design representative neighborhoods of any size and demonstrate our framework in a case study of three neighborhoods whose buildings have undergone electrification and have been equipped with active storage systems under RLC for flexibility.

Methodology

We describe in this section a framework for the design of representative neighborhood datasets from the EULP database as well as other data sources for whole building stock control and performance quantification in CityLearn. Figure 1 provides an overview of this framework split into three phases: 1) neighborhood design & data collection, 2) load simulation & dataset preparation and 3) control simulation & reporting. We elaborate on these phases in the following subsections.

Phase 1: Neighborhood Design & Data Collection

In the neighborhood design phase, the designer pre-filters the EULP database for buildings whose metadata match certain criteria such as archetype (including building use type and age), location, and equipment types. A designer may also choose to define the building count in the neighborhood. Building selection to meet the specified count can be done through random selection (White et al., 2021) or data-driven methods (Wilson et al., 2019). The metadata, occupancy, load schedules, and energy models for the selected buildings are then stored in a central database for easy retrieval and manipulation. Relevant weather files for energy simulations can be sourced from weather data repositories such as Wilcox and Marion (2008).

Optionally, the variance in building loads and indoor environmental conditions amongst neighborhood buildings can be improved by replacing generalized thermostat setpoint schedules with real-world thermostat setpoints from buildings in the same region using the ecobee Donate Your Data dataset (Luo and Hong, 2022). A similar approach taken for building selection, i.e. random selection or clustering, can be applied to map the real-world thermostat setpoints to selected buildings.

Phase 2: Load Simulation & Dataset Preparation

In the building load simulation & dataset preparation phase, the collected data are used to run EnergyPlus simulations to obtain ideal space cooling and heating loads. Other data determined from these simulations are domestic hot water (DHW) loads, lighting and plug loads.

Finally, designer input for DER availability including heat pumps, electric heaters, thermal storage, batteries, and photovoltaic (PV) systems as well as the system sizing for the buildings is defined in Phase 2. The resulting system specifications, and simulated ideal loads are then utilized to create a virtual representation of the intended neighborhood in CityLearn.

Phase 3: Control Simulation & Reporting

In Phase 3, the designer selects a control algorithm such as rule-based control (RBC), MPC, or RLC to manage the DERs in a control simulation. The control agent receives observations from the environment and in return predicts appropriate actions that improve defined objectives. Post-simulation evaluation of control performance is achieved by the user-selected key performance indicators (KPIs) that quantify energy flexibility, environmental impact, or occupant comfort.

Application

Here we provide background information on a case study where the framework in the Methodology has been applied to develop three neighborhoods whose buildings have been equipped with active storage systems under RLC for flexibility. To reproduce this application case study, we refer the reader to the project repository¹. We

¹https://github.com/intelligent-environments-lab/DOE_XStock/tree/bs2023

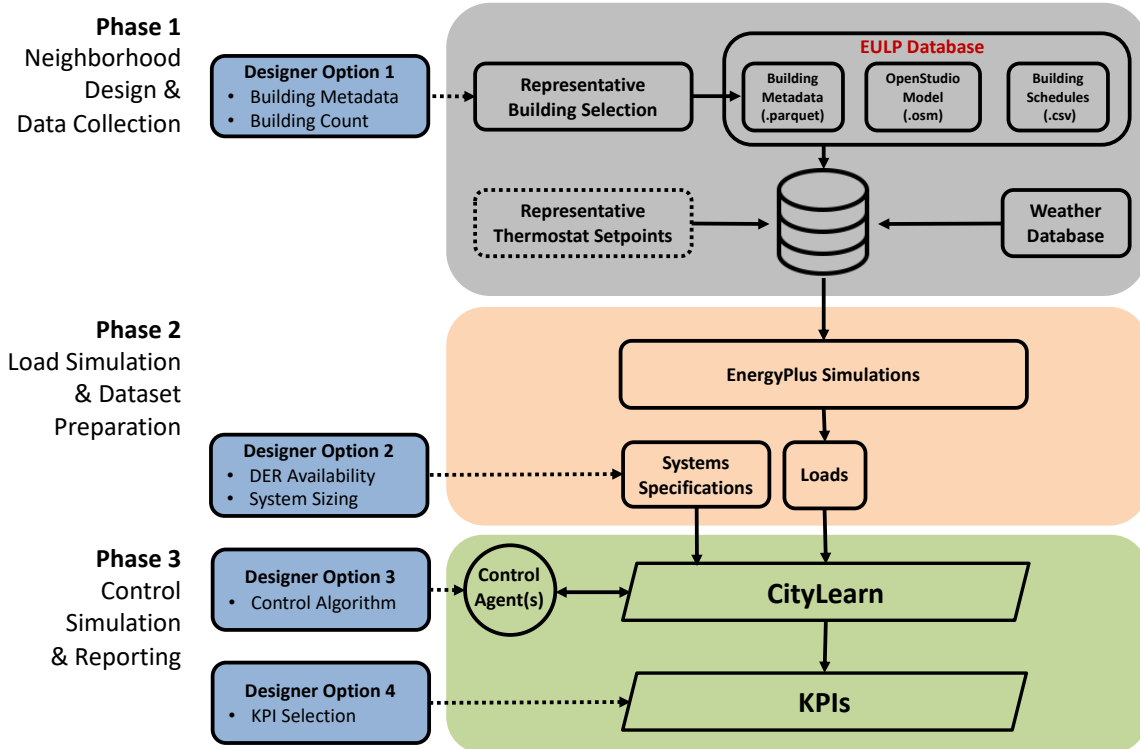


Figure 1: Methodology overview.

utilize the *resstock-amy2018-2021-release-1* version of the EULP dataset which is its first release of residential buildings stock models in the 2021 calendar year that were simulated using actual meteorological year (AMY) 2018 weather data. We are interested in developing CityLearn input data for three 100-building neighborhoods in Alameda Co., CA (CA), Travis Co., TX (TX), and Chittenden Co., VT (VT) that have contrasting climatic conditions and thus different building envelope properties and operational schedules (see Table 1). For these three counties, we query the EULP database for occupied single-family detached buildings that were constructed between the 1940s and 2010s. There are 962, 985, and 117 buildings in CA, TX, and VT that match this query.

We also utilize AMY (Wilson et al., 2022; Lamprecht, 2023) augmented with TMY3 (Wilcox and Marion, 2008) (for missing or unavailable variables) weather data for the 2018 calendar year from weather stations defined in the EULP database for the three neighborhoods. All data are stored in a local SQLite database. Our analysis period is January-March, 2018 (heating season) for CA and VT and June-August (cooling season) for TX, where the first two months are used to train the control agents and the last month is used to test the trained agents. The selection of analysis season for each neighborhood is guided by the greater of their cooling degree days (CDD) and heating degree days (HDD) as reported in Table 1, where greater HDD indicates heating dominance thus a heating analysis season and greater CDD indicate cooling dominance thus a cooling analysis season.

Table 1: Neighborhood summary where degree days are reported for 2018 AMY weather with respect to 18.3 °C base temperature.

County	Climate Zone	CDD	HDD
Alameda Co., CA	3C (Marine)	11	1590
Travis Co., TX	2A (Hot-Humid)	1821	93
Chittenden Co., VT	6A (Cold)	513	3858

Representative Building and Thermostat Setpoints Selection

For each neighborhood, we select a subset of 100 representative buildings from the total number of buildings that match our initial query and assign real-world setpoint schedules to each building. We follow the procedure outlined in Figure 2 that uses a combination of data-driven methods and random selection to identify representative buildings and thermostat setpoint schedules.

Representative building groups are determined by clustering metadata fields including building year of construction, orientation, occupant count, infiltration rate, ceiling, slab and wall insulation, window-to-wall-ratio (WWR) and energy use intensity (EUI) using the k-means algorithm (Hartigan and Wong, 1979) as shown in Figure 2a. We transform textual metadata to numerical values after which min-max normalization is applied to all fields to aid clustering. Using the Silhouette Index (Rousseeuw, 1987) we determine the appropriate number of building clusters to be six in each neighborhood, where the building vintage, orientation, slab and wall insulation are most distinct across clusters. Given the distribution of buildings across

clusters for each neighborhood, a desired number of buildings is randomly sampled using a weighted approach from all clusters where the proportion of buildings contributed by each cluster, p_i is defined in Equation (1) such that C_i is the number of buildings in cluster i and n is the number of clusters. The metadata, occupancy and load schedules, and energy models for these selected buildings are stored in the database.

$$p_i = \frac{C_i}{\sum_{i=0}^{n-1} C_i} \quad (1)$$

We identify representative setpoint schedules for selected buildings from the ecobee Donate Your Data dataset (Luo and Hong, 2022) that are sourced from the same or nearby locations. We first designate TX as cooling dominant and CA and VT as heating dominant as described previously (see Table 1) and determine the average daily setpoint schedule for each house in the region. We then follow the methodology outlined in Panchabikesan et al. (2021) to apply Dynamic Time Warping (DTW) k-shape clustering on the average daily cooling setpoint or heating setpoint schedules from at least 1000 buildings in each region for cooling or heating dominant neighborhoods respectively. K-shape clustering compares the shape of time series data and groups together similar patterns in shape (Paparrizos and Gravano, 2015). For our application, we group buildings by similar setpoint preferences by comparing the setpoint curves throughout the day. We select the cluster count that maximizes the cumulative score of three cluster validity indices (CVIs): Dunn Index (Dunn, 1973), Davies-Bouldin Index (Davies and Bouldin, 1979), and Silhouette Index (Rousseeuw, 1987). For all three neighborhoods, two clusters are identified. Finally Equation (1) is used to select the same number of buildings as in the neighborhoods and the actual setpoint schedule from an ecobee building is randomly assigned to an EULP building.

Load Simulation & Dataset Preparation

The collected data are used to run EnergyPlus simulations to obtain building ideal space cooling and heating loads. We first translate the OpenStudio energy models to EnergyPlus models then run two simulations on each building's model. The first simulation uses the as-provided model as a reference point for subsequent simulations. A second simulation where the mechanical heating ventilation and air conditioning (HVAC) systems are replaced with ideal load systems is then run and validated using the results from the first simulation. Given the multiplicity of energy models we use, we follow a scripting approach for using the openstudio² and eppy³ Python packages for energy model manipulation and simulation.

In validating the ideal load simulations, we find that only 73/100 and 43/100 selected buildings in CA and VT have $\geq 95\%$ of total timesteps where their indoor temperature

²<https://github.com/NREL/OpenStudio>

³<https://github.com/santoshphilip/eppy>

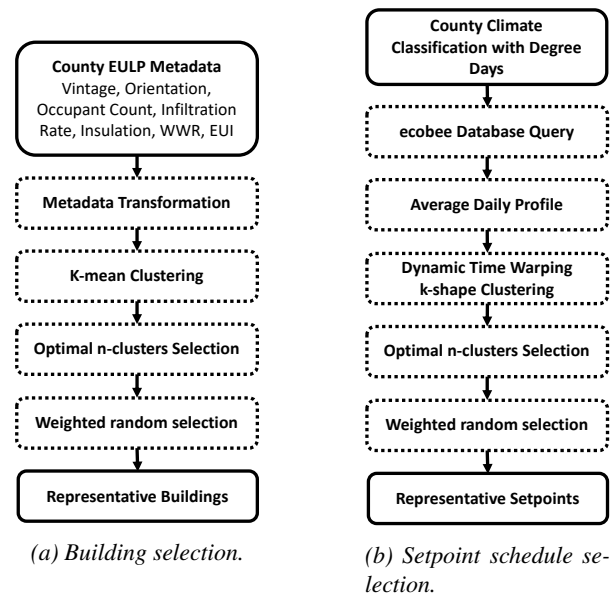


Figure 2: Neighborhood Design.

meets the setpoint for a $\pm 0.5^\circ\text{C}$ comfort band. The simulation errors in the other buildings are a result of one or more of the following: (a) errors in converting OpenStudio OSM file to EnergyPlus IDF file; (b) errors in the EnergyPlus IDF file when replacing the mechanical HVAC system with ideal load system as a result of broken object inter-relationships between HVAC objects that could not be resolved; (c) ecobee thermostat setpoint schedule being outside the bounds of the ideal load system supply temperature limits. These observations highlight some of the challenges that arise from taking a scripting approach instead of the OpenStudio graphical user interface to manipulate a large number of energy models. We proceed with the validated 73, 100, and 43 buildings in CA, TX, and VT respectively for the remainder of this work.

Control Problem

We utilize the three neighborhoods to assess the differences in energy flexibility caused by unique climatic conditions and envelope construction properties. The control problem we solve is the management of electrical storage (battery) and thermal storage for DHW heating in the neighborhoods for load shifting, which we elaborate on in the following subsections.

Environment

Figure 3 is an overview of the systems and interactions within the neighborhood's CityLearn environment. Each building in the environment has an air-to-water heat pump to satisfy space thermal loads, an electric heater sized to meet the peak hourly load for DHW heating, a battery, and DHW storage. All batteries are sized at 13.5kWh capacity and 2.5kW nominal power with 90% round-trip efficiency to replicate a standard home battery while the hot water storage tank is sized for the maximum daily load. Electric devices and plug loads consume electricity from any of the

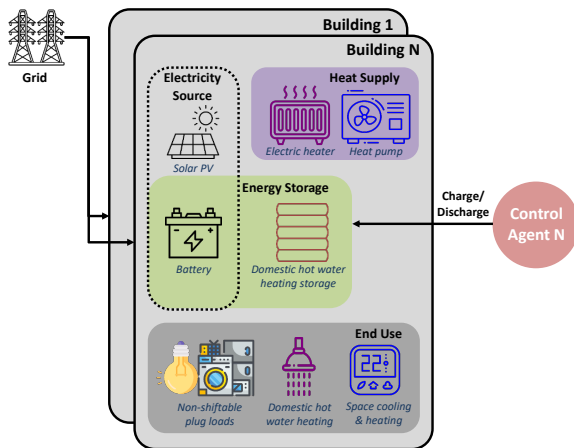


Figure 3: CityLearn environment and control agent interaction.

available electricity sources including the grid, a 10kW PV system, and battery. Excess solar generation is wasted as there is no energy flow between buildings.

Reinforcement Learning Controller

An independent controller manages energy storage in each building by determining how much energy to store or release at each timestep. Each building's battery and DHW storage are controlled by an independent soft actor-critic (SAC) controller. SAC is a model-free, off-policy RL algorithm as it is able to reuse experience and learn from fewer samples. SAC is based on three key elements: an actor-critic architecture, off-policy updates, and entropy maximization for efficient exploration and stable training. It learns three different functions: the actor (policy), the critic (soft Q-function), and the value function. For more details about SAC, we refer the reader to Haarnoja et al. (2018). We make use of the Stable Baselines3 v1.7.0 (Raffin et al., 2021) implementation of the SAC algorithm with its default hyperparameters.

We train the controller for 15 epochs on the initial two months in either the winter or summer season as applicable to the neighborhood and test its performance on the third month.

Observation and Action Space Design

The agent observation space is identical in all buildings and is made up of district-level temporal (day-of-week, hour) and weather observations (dry-bulb temperature and direct solar irradiance, and their six-hour and 12-hour forecasts), building-specific states (solar generation, battery and DHW storage state of charge (SoC), plug and net loads). The observations are transformed to aid the learning process by applying cyclical transformation, one-hot encoding, and min-max normalization on periodic, discrete, and continuous observations respectively.

The action space for each building's controller is a two-dimensional vector of the proportion of 1) battery and 2)

DHW storage capacities to be charged or discharged that are bound between $[-1.0, 1.0]$.

Reward Design

The reward function (Equation (2)) is designed to minimize net electricity consumption, E , at each timestep. It encourages net-zero energy use by penalizing grid load satisfaction when there is energy in the battery or DHW storage as well as penalizing net export when these storage devices are not fully charged through the penalty term, p defined in Equation (3). There is no penalty nor reward given when both storage systems are fully charged during net export to the grid. Whereas, the penalty is maximized when there is net import from the grid when the storage systems are charged to capacity.

$$r = -p \cdot |E| \quad (2)$$

$$p = 2 + \frac{E}{|E|} \cdot (\text{SoC}_{\text{Battery}} + \text{SoC}_{\text{DHW storage}}) \quad (3)$$

Key Performance Indicators

We evaluate the controllers' performance in the building energy management using five KPIs that are to be minimized: electricity consumption, average daily peak, ramping, peak demand, and (1 - Load Factor). Average daily peak, ramping, peak demand and (1 - load factor) are neighborhood-level KPIs that are calculated using the aggregated district-level net electricity consumption. Electricity consumption is a building-level KPI that is calculated using the building-level electricity consumption and is reported at the neighborhood level as the average of the building-level values. We refer the reader to (Nweye et al., 2023) for the formulation of these KPIs.

Control Simulation & Reporting

We show the distribution of the electricity consumption KPI in Figure 4 for all buildings in each neighborhood for the one-month test period. All buildings in CA and TX are able to reduce their electricity consumption from the grid compared to the baseline by taking advantage of load shifting provided by their battery and DHW storage. On average, buildings in CA and TX neighborhoods are able to reduce energy drawn from the grid by 19% and 33% respectively, whereas only 14 of 43 buildings in VT performed better than the baseline, and on average, the RL controller only maintains the same electricity consumption as the baseline.

We summarize for the one-month test period the neighborhood-level KPIs in Figure 5 where the TX neighborhood is able to provide better grid services compared to the others. Ramping is most improved on average in the three neighborhoods. CA and TX are able to reduce peaks indicated by average daily peak and peak demand but are unchanged in VT. 1 - Load factor is reduced by 2% and 12% in VT and TX but increases by 3% in CA.




Figure 4: Distribution of net electricity consumption compared to the baseline in all buildings.




Figure 5: Neighborhood-level KPI summary compared to the baseline scenario of no storage.

We show the neighborhood-level net electricity consumption profiles for the cases with and without storage as well as the average $SoC_{battery} \pm$ one standard deviation for the last 14 days of testing in Figure 6. The profiles in CA and TX show the ability of the independently controlled building storage systems to take advantage of the daytime solar generation for charging and release the stored energy in the evening, thus reducing the aggregated grid load during this peak period. In the case of the CA neighborhood, there is a morning and an evening peak. The controllers are however only able to learn to discharge stored energy to reduce the evening peak thus there is still dependence on the grid in meeting the earlier peak. The VT profiles for the cases with and without storage are very similar as the controllers do not learn the load shifting task as in the case of the other two neighborhoods. The $SoC_{battery}$ distribution in VT shows irregular charge/discharge pattern and underutilized battery capacity for load shifting, which could be attributed to sub-optimal controller hyperparameters and reward function.

Discussion

The principal objective of this study is to provide a framework for building stock control algorithm benchmarking. Our case study uses an RL controller to coordinate the charging and discharging of batteries and DHW storages to best provide building energy flexibility to the grid, and the results shown are promising for achieving neighborhood-level DR. The four neighborhood-level KPIs reflect how well the RL controller is able to flatten the demand curve of the neighborhood, and it is shown that across the three neighborhoods the most significant reductions on average occurred in ramping. Ramping represents the change in energy demand from one hour to the




Figure 6: Comparison between net electricity consumption with and without storage at the neighborhood-level, as well as, the average $SoC_{battery} \pm$ one standard deviation.

next for each building, and thus is minimized when the demand from the grid is as smooth as possible throughout the day. With the proliferation of PV generation, ramping is becoming an increasingly important variable as the decreasing production of energy from PVs in the late afternoon leads into the period of high demand in the evening (Fattaheian-Dehkordi et al., 2022), resulting in the so-called duck curve (Denholm et al., 2008). Average reductions in ramping of 19% for CA, 35% for TX, and 2% for VT show that the RL controller can significantly reduce ramping in CA and TX, though only slightly in VT.

Looking at the different regions more specifically, TX shows significant improvements ranging from 12% to 42% reductions in (1 - Load Factor) and average daily peak energy demand respectively, thus highlighting the potential services these buildings could provide for the grid during periods of critical demand. Ke et al. (2016) estimated that during periods of extreme heat waves, the average energy consumption of a building could increase by approximately 30% and thus could be sustained by the grid if the buildings had adequate storage and renewable energy production as proposed in this study. However, fairness amongst buildings is not yet addressed in our control problem and remains an area for future work. For example, one building achieved the highest reductions in average daily peak demand at 47% while another building saw 0% reductions in average daily peak demand for TX. Likewise in CA, the changes in average daily peak demand seen by two different buildings were -24% and +2% respectively, meaning at least one building of the neighborhood actually saw an increase in the average peak demand. As reductions in grid energy consumption would translate directly to reductions in energy cost, this discrepancy remains an area to be explored in future work to ensure fairness in remunerations for DR services.

While the results in CA and TX are encouraging, the RL

controller is not able to effectively manage the battery and DHW storage to have an impact on the KPIs in the case of VT. Several modeling choices could explain the discrepancy. Firstly, the RL controller is trained on 15 epochs for the three neighborhoods. However, the VT building energy consumption data showed large changes in consumption and solar production throughout the two-month epoch, and thus further epochs of training could enhance the performance. Likewise, the SAC algorithm hyperparameters are fixed for all buildings and all neighborhoods; further hyperparameter tuning could yield better results. Similarly, the same reward function is used across all three neighborhoods. The results from VT show that the RL controller had not yet learned to exploit all solar production during the afternoon, thus a neighborhood-tailored reward function could better incentivize the RL controller to exploit all solar production during the day and use the excess to charge the ESSs. The fixed PV system and battery sizing for all buildings are current limitations that could be addressed in future work as some buildings may be disadvantaged in available flexible capacity compared to their fixed loads.

Conclusion

This paper presents a methodology for designing representative neighborhoods for energy modeling and simulation using open-source data in the U.S. Due to the availability of this data and the state of computing technology, we have shown that it is now possible to do large-scale simulations of energy flexibility for battery and DHW storage. Using a RL controller, reductions in neighborhood-level KPIs were on average 9.69% in CA, 31.2% for TX, and less than 1% in VT thus showing that while promising results were seen for TX and CA, further work remains to achieve enhanced energy flexibility in VT. As demand response and battery storage are estimated to provide approximately 25% of global energy flexibility requirements by 2030 (and increasing to 50% by 2050) (IEA, 2022), there is a vital need for developing tools like the RL controller shown in this study to maximize the effectiveness of various DERs.

References

Blum, D., J. Arroyo, S. Huang, J. Drgoña, F. Jorissen, H. T. Walnum, Y. Chen, K. Benne, D. Vrabie, M. Wetter, and L. Helsen (2021, 9). Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation* 14(5), 586–610.

Davies, D. L. and D. W. Bouldin (1979, Apr). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-1*(2), 224–227.

Denholm, P., R. Margolis, and J. Milford (2008, 2). Production cost modeling for high levels of photovoltaics penetration.

Drgoña, J., J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50(May), 190–232.

Dunn, J. (1973). A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters. *Journal of Cybernetics* 3(3), 32–57.

Fattaheian-Dehkordi, S., A. Abbaspour, M. Fotuhi-Firuzabad, and M. Lehtonen (2022). A new management framework for mitigating intense ramping in distribution systems. *Energy* 254, 124100.

Haarnoja, T., A. Zhou, P. Abbeel, and S. Levine (2018, 1). Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *ICML*.

Hartigan, J. A. and M. A. Wong (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28(1), 100–108.

Horseý, H., A. Parker, A. Farthing, M. Dahlhausen, M. Praprost, C. Bianchi, J. Robertson, S. Horowitz, L. Zhang, U. O. of Energy Efficiency, R. Energy, U. L. D. Research, and Development (2020, 10). Comstock™.

International Energy Association (2022). *Demand Response*.

Jensen, S. Ø., A. Marszal-Pomianowska, R. Lollini, W. Pasut, A. Knotzer, P. Engelmann, A. Stafford, and G. Reynders (2017). Iea ebc annex 67 energy flexible buildings. *Energy and Buildings* 155, 25–34.

Ke, X., D. Wu, J. Rice, M. Kintner-Meyer, and N. Lu (2016). Quantifying impacts of heat waves on power grid operation. *Applied Energy* 183, 504–512.

Lamprecht, C. S. (2023). Meteostat Python.

Luo, N. and T. Hong (2022, 3). Ecobee donate your data 1,000 homes in 2017.

Neukomm, M., V. Nubbe, and R. Fares (2019, 4). Grid-interactive efficient buildings.

Nweye, K., S. Sankaranarayanan, and Z. Nagy (2023). Merlin: Multi-agent offline and transfer learning for occupant-centric energy flexible operation of grid-interactive communities using smart meter data and citylearn.

Panchabikesan, K., M. Ouf, U. Eicker, G. Newsham, and H. Knudsen (2021, Sep). Investigating thermostat setpoint preferences in canadian households. In *Proceedings of the 17th IBPSA Conference*, Bruges, Belgium.

- Paparrizos, J. and L. Gravano (2015, May). k-shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, Melbourne Victoria Australia, pp. 1855–1870. ACM.
- Pinto, G., M. S. Piscitelli, J. R. Vázquez-Canteli, Z. Nagy, and A. Capozzoli (2021). Coordinated energy management for a cluster of buildings through deep reinforcement learning. *Energy* 229.
- Raffin, A., A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann (2021). Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research* 22(268), 1–8.
- Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* 20, 53–65.
- U.S. Department of Energy (2021). *Next-Generation Grid Technologies*.
- Vázquez-Canteli, J. J. R., J. Kämpf, G. Henze, and Z. 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*, 356–357.
- Wang, Z. and T. Hong (2020). Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy* 269(April), 115036.
- White, P. R., J. D. Rhodes, E. J. H. Wilson, and M. E. Webber (2021, 9). Quantifying the impact of residential space heating electrification on the texas electric grid. *Applied Energy* 298(C).
- National Renewable Energy Lab.(NREL), Golden, CO (United States) (2008). *Users manual for tmy3 data sets (revised)*.
- Wilson, E. J. (2017, 9). Resstock - targeting energy and cost savings for u.s. homes.
- Wilson, E. J., C. B. Harris, J. J. Robertson, and J. Agan (2019). Evaluating energy efficiency potential in low-income households: A flexible and granular approach. *Energy Policy* 129, 710–737.
- Wilson, E. J. H., A. Parker, A. Fontanini, E. Present, J. L. Reyna, R. Adhikari, C. Bianchi, C. CaraDonna, M. Dahlhausen, J. Kim, A. LeBar, L. Liu, M. Praprost, L. Zhang, P. DeWitt, N. Merket, A. Speake, T. Hong, H. Li, N. Mims Frick, Z. Wang, A. Blair, H. Horsey, D. Roberts, K. Trenbath, O. Adekanye, E. Bonnema, R. El Kontar, J. Gonzalez, S. Horowitz, D. Jones, R. T. Muehleisen, S. Plathotam, M. Reynolds, J. Robertson, K. Sayers, and Q. Li (2022, 3). End-use load profiles for the u.s. building stock: Methodology and results of model calibration, validation, and uncertainty quantification.
- Zhang, Z. and K. P. Lam (2018). Practical implementation and evaluation of deep reinforcement learning control for a radiant heating system. In *Proceedings of the 5th Conference on Systems for Built Environments*, BuildSys '18, New York, NY, USA, pp. 148–157.