

Estimating cooling demand flexibility in a district energy system using temperature set point changes from selected buildings

Ryan C. Triolo^{a,*}, Ram Rajagopal^a, Frank A. Wolak^b, Jacques A. de Chalendar^c

^a*Department of Civil and Environmental Engineering, Stanford University, Stanford, 94305, California, United States*

^b*Department of Economics, Stanford University, Stanford, 94305, California, United States*

^c*Department of Energy Resources Engineering, Stanford University, Stanford, 94305, California, United States*

Abstract

Energy demand flexibility from buildings remains a largely untapped resource in electric power systems, in spite of its potential to be a low-cost substitute for investments in rarely used generation capacity. We develop a general methodology to estimate the aggregate cooling demand response for a large number of co-located buildings to thermostat temperature set point increases using empirically estimated building-level demand reductions in a subset of these buildings. For this subset of buildings, estimates were previously computed from real-world experimental data. The response of each remaining building is estimated as a different weighted sum of the empirical estimates, where the weights depend on observable characteristics of the buildings. We apply our method to a district energy system at a university campus that is roughly equivalent to a city of 30,000 people. Cooling is produced at a central energy facility with electric heat pumps and distributed to 124 commercial buildings through a chilled water loop. The response of six of these buildings to 1.1°C (2°F) daily temperature set point adjustments was previously estimated. Our methodology provides estimates for all 124 buildings and an estimate of the campus-wide demand response potential by leveraging a dataset including both structural (e.g. age, square footage) and operational (cooling loads and types of building operation) features for the full set of buildings. We estimate a 13.47% reduction in the campus energy system capacity needs under a 1.1°C daily set point increase in all campus buildings during the 10 highest system demand days in 2020. On the highest demand day of 2020, we find that our predicted demand reduction could provide services equivalent to those provided by a lithium-ion battery with \$4.6 - \$8.0 million installation cost at current prices and a storage capacity of 35.6 - 52.6 MWh.

Keywords: Demand response, Building feature clustering, District energy systems, Load control, Demand-side flexibility

*Corresponding author: rtriolo@stanford.edu

1. Introduction

1.1. Motivation and context

Residential and commercial buildings accounted for 29% of total final energy consumption in the United States in 2020 [1]. Increased deployment of renewable generation technologies coupled with retirements of conventional, dispatchable generation resources has led to challenges managing the real-time balance of supply and demand [2]. Active management of energy consumption in buildings—and in commercial buildings in particular—holds potential to provide significant flexibility in electric grids with a high share of intermittent renewable generation [3].

Commercial buildings in California accounted for 34.0% of the state’s electricity demand in 2020 [4]. According to the 2006 California Commercial End Use Survey, heating ventilation and air conditioning (HVAC) accounts for 29% of commercial building electricity use, which implies 9.9% of total in-state electricity consumption in 2020. Moreover, because this is an annual average, the share of consumption devoted to air conditioning is considerably higher during peak demand hours of hot summer days.

Thermal demand response through control of HVAC system thermostat temperature set points has the potential to provide sizeable energy demand reductions with relatively small impacts on the heating and cooling services delivered to building occupants [5, 6]. Dynamic thermostat set point control can also reduce the need for future HVAC capacity investments.

District heating and cooling systems are of particular interest for harnessing demand-side energy flexibility because of their large energy demands. They consist of groups of buildings with their heating and cooling needs served through a central energy facility (CEF). Types of locations that may have heating and cooling provided by this type of system include university campuses, medical centers, corporate campuses, and buildings in urban areas [7].

The capital investment required to upgrade heating and cooling capacity for district energy systems is substantial. Moreover, where systems are designed to meet peak loads that are a large multiple of average loads, this implies that a high proportion of capacity is used infrequently. Dynamic thermostat set point changes can increase the capacity utilization rate of a district energy system while at the same time can also provide real-time energy and operating reserves to the electric system operator.

Deriving estimates of district energy system demand reductions from temperature set point increases with an acceptable degree of accuracy is challenging. Physics-based building-level models require a large amount of input data and the demand response estimates obtained can be highly sensitive to modeling assumptions. A statistical model of the relationship between building-level energy use and the temperature set point recovers an estimate based on actual use of the building, but implementing these set point changes and measuring their impacts on building-level energy use is expensive. Consequently, leveraging estimation results from a small subset of buildings where set point increases are implemented can be a low-cost way to

estimate the demand reduction from set point changes in all buildings in a district energy system.

We propose a methodology to extend set point increase estimation results from a small number of buildings under a variety of summer month weather conditions to the level of the district energy system using observed structural and operational characteristics of all buildings in the system. We focus on a district energy system that serves a university campus in Northern California. In this system, the CEF produces heated and chilled water to serve the HVAC demands of 124 buildings through hot and cold water loops. The system includes cold water chillers, hot water generators, heat recovery chillers, and hot and cold water storage. See de Chalendar, et al. [8] for a detailed description of this district energy system and the optimization problem of the system operator. Methods described here also will be applicable to other similar settings where district energy systems are deployed to serve groups of buildings [9], and to estimate the demand flexibility in larger building stocks more generally. In this study we focus only on managing cooling demand of the system, and do not address heating capacity requirements. This is because the cooling capacity is the binding constraint in this system and experimental data was only available for cooling demand response to set point changes.

Figure 1 displays the daily cooling load duration curve for this campus energy system for 2020 (with 2018 and 2019 as dashed lines for reference). This load duration curve displays the daily cooling demand of the system arranged from the highest to lowest day of the year. The horizontal axis of the plot labeled “Day of the year” indicates the ranking of each daily cooling load in the associated year (e.g. day 1 is the highest demand day of the year, day 2 is the second highest, and so on). This curve provides an indication of how much potential there is for demand flexibility to save on CEF capacity. In this study we define “demand flexibility” as the change in demand that can be realized as the result of a particular load control action. If the CEF is designed with the capacity necessary to serve cooling demand on all days of the year with no forced curtailment, this amounts to a “design for the worst day” approach. For a system with the load duration curve pictured in Figure 1, this would lead to a sizeable amount of capacity being used very infrequently. The cross-hatched area in the figure denotes the amount of unused installed capacity if a system were designed to meet the maximum demand of this year.

Figure 1 also shows a steep downward sloping curve in the top left region of the plot, indicating an opportunity for cooling demand reduction on a small number of days to save on installed capacity. For example, for a district energy system designed to meet the daily cooling demands for 2020, 10% of the capacity would only be needed on five days of the year, while 20% of the capacity would only be needed fifteen days of the year. This implies that load flexibility on the highest demand days can greatly reduce the need for installed capacity.

In this study, we assess the load flexibility potential of a 1.1°C (2°F) air conditioning set point increase on high cooling demand days for this district energy system. To do this, we devise a methodology based on building feature clustering to estimate the likely demand response of all buildings on the campus using the




Figure 1: 2020 daily cold water duration curve (2018 and 2019 indicated as dashed lines)

average measured demand response for a small set of buildings on the campus to a 1.1°C set point change under a variety of summer month weather conditions. We then use the resulting demand response estimates for all buildings to construct a counterfactual annual cooling load duration curve that reflects the impact of a 1.1°C set point change in all buildings on the campus during the ten highest demand days of the year. We find a predicted annual peak demand reduction of more than ten percent. This result is encouraging for the use of building load flexibility as a lower cost alternative to managing demand peaks relative to batteries or other energy storage devices.

1.2. Review of related work

Estimating demand response potential of residential and commercial buildings has been the topic of a number of studies in recent years. In contrast to the current study, the majority of research in this area leverages simulation through physics-based models, rather than experimental methods. Previous research summarized here has focused on i) estimation and simulation of building-specific demand response potential under thermostat set point adjustments, ii) characteristics of commercial buildings that are important determinants of flexibility, iii) estimating demand response potential in large building stocks, and iv) clustering methods for characterizing building stocks.

A large number of studies use the EnergyPlus modeling software to simulate building energy use and demand response [10, 11, 12]. For example, Li et al. [13] use EnergyPlus and regression models to investigate the demand response potential of office buildings in Beijing, China. They find an average 10.1% decrease in energy consumption during working hours under a 1°C set point increase. Liu et al. [14] compare EnergyPlus simulation results with data from 12 commercial buildings participating in a demand response program in California. The authors find much higher variation of energy reductions across multiple DR events in field data compared to simulation.

Other studies have used EnergyPlus to test the efficacy of less computationally intensive statistical methods. Yin et al. [6] investigated how piece-wise linear modeling can be a computationally efficient method to estimate HVAC demand response potential at the hourly level comparing model performance to a large database of EnergyPlus simulation data. Other physics-based modeling tools have been used, such as eQuest, used by Cai et al. [15] to estimate HVAC demand response potential in a single three story building in the Chicago area. The study found only a 1.6% average daily energy savings with a 5°F (2.8°C) set point increase.

Examples of demand flexibility assessments that employ experimental methods using thermostat set point adjustments include Aduda et al. [16] where the authors find that an average-sized office building in the Netherlands reduced peak cooling demand by up to 25% for 20 minutes of operation. An experimental study on commercial buildings in Northern California found that chiller electricity consumption can be reduced by approximately 33% over a four hour demand response event while maintaining building interior temperatures within a defined comfort range [17]. Yin et al. [18] explores pre-cooling strategies for peak demand reduction in eleven office buildings in California, comparing simulation results to field tests. The authors find that a strategy with pre-cooling can reduce peak electrical demand by 15-30%.

The results of studies measuring HVAC demand flexibility indicate a high degree of heterogeneity in demand response potential. Demand flexibility in commercial buildings is likely to depend on a large number of physical parameters and other non-static factors. A variety of studies have focused on the identification of the characteristics that are most important for determining HVAC demand flexibility potential.

Certain characteristics of commercial buildings that have been found to be important for determining building demand flexibility include surface area to volume ratio, window area to external surface area ratio, insulation of walls and ceiling, thermal mass, and occupancy details [6, 18]. Khalilnejad et al. [19] found that occupancy and higher peak load relative to base load of buildings were key determinants of energy use flexibility with increased HVAC thermostat set points. Chen et al. [20] found that thermal mass and occupant behavior are important determinants of building energy flexibility. A study by Ghahramani et al. [21] investigated the effect of various structural and operational building characteristics on HVAC demand response potential through thermostat set point adjustment in commercial office buildings.

Aggregation of demand response across many buildings has the potential to unlock a large amount of demand flexibility [22]. A variety of techniques have been used in the literature to estimate flexibility at a large scale. Hoyt et al. [23] use EnergyPlus to demonstrate the large HVAC load savings that can be achieved in a variety of climate zones. Stinner et al. [24] propose a methodology to measure flexibility in small groups of buildings. However, the building cluster assumed identical building characteristics and only included four buildings.

Simple physics-based models can be used in a bottom-up approach to estimate flexibility in building stocks. Hedegaard et al. [25] use a reduced-order resistance-capacitance model for the space heating demand

response of 159 homes in Aarhus, Denmark. The authors use a statistical model to update parameters important for determining demand response potential including window to floor ratio, infiltration rate, insulation value (U-value), occupant density, and thermal capacity of construction elements.

An alternative to physics-based models is to use building features, as in Reynders et al. [26], where thermal storage and heating demand flexibility are estimated for the Belgian building stock. The authors use factors such as age of buildings, renovation status, size of buildings, building structure, and heating system type to estimate flexibility potential.

Other work has focused on demonstrating the importance of incorporating aggregations of electric loads, and thermostatically controlled loads in particular, in power system operation to achieve reliable operation of power grids under high renewable penetration [27]. The benefits and challenges of demand side management (DSM) generally in the UK context is discussed in [28]. The author points to low utilization of generation and network infrastructure as an opportunity for DSM (including demand response) to provide efficiency gains in electricity systems. Callaway [29] investigates the dynamic response of aggregations of thermostatically controlled loads over short time scales to smooth output from wind generators. The author finds that small thermostat adjustments can provide significant balancing of short term fluctuations in wind power output with minimal impact on building temperatures.

Clustering approaches have been used to characterize local, regional, and national building stocks for simulation and estimation of energy usage performance and flexibility. Some clustering approaches have taken building characteristics as features upon which to cluster while others have used timeseries data of energy usage profiles or HVAC system operational state.

A review of demand response potential estimation in building clusters shows that classification and clustering of buildings based on building characteristics and occupant behavior have been used to estimate the potential at a large scale [30]. The review suggests that accurate estimation of demand response potential requires models that account for diversity among buildings, where diversity may be in energy inputs, end-use equipment, building types, building thermal characteristics, and occupancy details.

Jones et al. [31] describes a software tool (the Energy and Environmental Prediction Model) intended for sustainable development planning at the urban scale. The clustering technique uses features of buildings including location, building dimensions, building age, built form (e.g. number of floors and window to wall ratio), and additional assumptions (e.g. U-values, heating and cooling equipment types) to group buildings with similar expected energy use patterns. Although computer hardware and software capabilities have advanced considerably since the publication of this work, the methodology in this study demonstrates that there exists a precedent for using feature clustering methods in estimating building energy usage and performance, and provides evidence in academic literature as to building features that are indicative of energy usage patterns.

A number of studies are based on K -means clustering. Gao and Malkawi [32] employ this method to

predict building energy performance and find an improvement from the results of their proposed clustering algorithm over the EPA’s Energy Star approach for energy benchmarking of commercial buildings. Zhan et al. [33] employ a two-step K -means clustering process to categorize buildings using daily load profiles and conduct energy benchmarking. The authors quantify operational dissimilarity between buildings as the Euclidean distance between feature vectors. The method is also used by Petteue et al. [34] to find representative residential buildings to estimate energy use flexibility in large building stocks. The authors use time series data on electricity demand and air conditioning system status (on/off).

1.3. Research contribution

This paper makes three main contributions. First, a methodology is developed to estimate cooling demand flexibility of a large number of co-located buildings based on their observable characteristics and measured daily cooling energy demand reduction from a temperature set point increase in a small number of buildings. Second, the methodology is applied to a real-world district energy system serving 124 buildings. Flexibility estimates were available from previously published work for six of these buildings, our method allows us to now provide demand flexibility estimates for all 124 buildings and for the combined district energy system. Finally, the results of this study provide insights into district energy system capacity planning and the potential role that thermal demand response from commercial buildings can play in future electricity systems.

2. Methodology

We estimate the demand flexibility potential in a commercial building stock where air conditioning demand is served by a district energy system. In a first step, thermostat set point adjustment experiments were conducted in a subset of the buildings to empirically measure the achievable cooling demand reductions under a 1.1°C (2°F) daily set point increase. Second, we constructed a dataset of building features for all buildings in the building stock. Feature data were selected that capture the most important building characteristics that determine cooling demand flexibility. In a third step, we develop and implement a methodology to extend the empirical demand flexibility estimates to estimate the flexibility potential of unobserved buildings. Finally, using these building-level estimates we construct an estimate potential load reduction for the entire building stock.

In the following subsections we describe the methodology implemented for the first three steps here. The first step was described in detail in a previous study [35]. In Subsection 2.1 we provide a brief summary of that work. Subsections 2.2 and 2.3 describe the methodology of the second and third steps here. The system level estimate results of the fourth step are discussed in Section 4. In Subsection 4.1 we implement a leave-one-out (LOO) validation process to test the accuracy of the methodology and choose appropriate model parameters.

2.1. Estimating cooling demand flexibility of the observed buildings

During the summer of 2021 temperature set point adjustment experiments were conducted in six buildings in the district energy system under study. All buildings had a business as usual (BAU) upper bound thermostat set point of 74°F (23.3°C). On treatment days the upper bound of the set point was increased to 76°F (24.4°C) for the entire day. For a detailed discussion of the experimental setting and results, see de Chalendar et al. [35].

In that study, the following regression model was estimated to predict the daily cooling load reduction for each building under the 2°F (1.1°C) set point increase for the day:

$$\log y_t = \beta I_{SP,t} + \theta_0 T_t + \theta_1 I_{W,t} + \theta_2 + \epsilon_t. \tag{1}$$

Here, observations were at the daily level, indexed by t . y_t is the building cooling load for day t , measured in joules (J). $I_{SP,t}$ is an indicator variable equal to 1 if the buildings set point was increased on day t and 0 otherwise. T_t is the daily mean outdoor temperature (F), measured and recorded onsite. $I_{W,t}$ is an indicator variable equal to 1 if day t is a weekend day and 0 otherwise. For buildings that were determined to not have a changing weekend and weekday schedule, the variable $I_{W,t}$ was excluded from the regression. The parameters β , θ_0 , θ_1 , and θ_2 are estimated for each of the six buildings by ordinary least squares. Estimates of the model parameters are indicated by $\hat{\beta}$, $\hat{\theta}_0$, $\hat{\theta}_1$, and $\hat{\theta}_2$, respectively. ϵ_t is a random error term with mean zero. For more details on alternative statistical models considered and the performance of this model see [35].

The assumed data generating process and regression structure imply that the value of $\hat{\beta}$ provides an estimate of the mean daily cooling load reduction, as a percent of BAU load, controlling for daily average temperature and a weekend indicator. In particular, the estimated daily percent load reduction is given by: $100 \cdot (1 - \exp(\hat{\beta}))$. Values of $\hat{\beta}$, standard error (S.E.) of $\hat{\beta}$, and the associated estimated percentage demand reduction for each building are shown in Table 1. In this work, we seek to obtain approximations of $\hat{\beta}$ for out-of-sample buildings for which experimental data is not available.

Building	$\hat{\beta}$	S.E.	Estimated load reduction
BLDG-1	-0.24	0.048	21.3%
BLDG-2	-0.33	0.028	28.1%
BLDG-3	-0.14	0.031	13.1%
BLDG-4	-0.23	0.022	20.5%
BLDG-5	-0.044	0.017	4.3%
BLDG-6	-0.034	0.0092	3.3%

Table 1: Estimated coefficients of the temperature set point treatment for in-sample buildings.

2.2. Feature Selection

From the values of $\hat{\beta}$ reported in Table 1, we observe that there is a high degree of heterogeneity in demand response potential among the observed buildings, with predicted daily cooling load reductions ranging from 3.3% to 28.1% for a 1.1°C (2°F) set point change. This result implies that heterogeneity in demand responses is likely to be present in the 118 remaining buildings in the campus energy system under study. The experimental buildings differed from each other in a variety of aspects that contribute to heterogeneity in demand response potential. Some buildings had a high proportion of floor space taken up by offices, while others have higher proportions of classrooms, laboratories, or library facilities. Buildings also varied in their energy intensity and ratio of daily maximum demand to daily average demand. In addition, there was variety in building size, year of construction, building materials, and other characteristics.

In order to extend these experimentally derived demand response results to the entire system, we create a metric of similarity or dissimilarity of out-of-sample buildings relative to the buildings in the experimental data set. We hypothesize that an out-of-sample building that is more similar to a certain in-sample building according to the metric, will tend to have a demand response potential closer to that building than another building to which it is less similar. This similarity metric is based upon a set of feature data, observable for all buildings, that is related to building energy demand and potential for flexibility.

We use a vector of features that is intended to capture the most important building characteristics that determine cooling load flexibility potential. We then use these features to quantify the similarity of out-of-sample buildings to the observed buildings. A review of relevant literature was undertaken to inform the selection of building characteristics (Subsection 1.2). Data features were then chosen that can capture these characteristics to the greatest extent possible given the availability of data.

Nine features were collected for each of the 124 buildings in the system. The nine features are: (1) “Area”: total floor space of the building (m²); (2) “LAB”: proportion of floor space accounted for by laboratories; (3) “CLS”: proportion of floor space accounted for by classrooms; (4) “LIB”: proportion of floor space accounted for by library facilities; (5) “OFF”: proportion of floor space accounted for by offices; (6) “Intensity”: cooling demand intensity (MJ/m²/day); (7) “Year”: year of building construction completion; (8) “Peak”: metric of daily load increase relative to daily average load; and (9) “SA-Vol”: surface area to volume ratio.¹

2.3. Estimating similarity scores

To quantify the similarity of unobserved buildings to those in the experimental dataset we have developed a “similarity score” that assigns an estimate of $\hat{\beta}_i$ for all remaining buildings on the campus, which ultimately

¹ “Peak” is defined as the average of $(a_k - b_k)/c_k$ across all summer weekdays indexed by k . Here a_k : maximum hourly load on day k ; b_k : minimum hourly load on day k ; and c_k : average daily load on day k .

allows us to estimate demand response under a 1.1°C set point change at the campus level. The calculation of this similarity score draws upon the logic of K -means clustering in the machine learning literature. K -means clustering is an unsupervised learning approach to partition a set of data into sets of similar observations based on individual feature vectors [36]. The premise is that a non-experimental building that is closer, measured by Euclidean distance in the feature space, to one in-sample building is expected to behave more like that building than a building further away.

The computation of similarity scores appears in Equations (2)-(4).

$$\gamma_{i,s} = \|\mathbf{x}_i - \mathbf{x}_s\| \quad (2)$$

$$\pi_{i,s} = \frac{\min(\gamma_{i,\cdot}^\alpha)}{\gamma_{i,s}^\alpha} \quad (3)$$

$$\pi_{i,s}^* = \frac{\pi_{i,s}}{\sum_{s \in \mathcal{S}} \pi_{i,s}} \quad (4)$$

Let \mathcal{S} be the set of observed, in-sample buildings and let \mathcal{I} be the set of all buildings in the system where $\mathcal{S} \subset \mathcal{I}$. Let $\hat{\beta}_s$ for all $s \in \mathcal{S}$ represent the values of $\hat{\beta}$ for the in-sample buildings, where point estimates and asymptotic distributions were estimated via Equation (1). Let $\hat{\beta}_i$ for all $i \in \mathcal{I} \setminus \mathcal{S}$ represent the unknown value of $\hat{\beta}$ for each out-of-sample building. \mathbf{x}_i is an m -element vector of building feature data (where m is the number of features) for building i . The feature data is normalized to have zero mean and standard deviation 1 across all buildings in the dataset. Feature data used in this analysis are discussed in greater detail in Section 3.

In (2) for a given out-of-sample building, i , the Euclidean distance in the feature space is computed between that building and each in-sample building, s , and denoted by $\gamma_{i,s}$. Values $\pi_{i,s}$ are computed in (3), note that a value of 1 is assigned to the element associated with the closest in-sample building while values of less than or equal to 1 are assigned to the elements associated with the other buildings. Finally, in (4) the vector $\pi_{i,\cdot}$ is normalized such that $\sum_s \pi_{i,s}^* = 1$. Here, α is an important tuning parameter that serves as a penalty on distance from building i to building s , the properties of which are discussed in greater detail below.

Once a similarity score vector has been computed we then compute for all out-of-sample buildings the following estimate of $\hat{\beta}$, which is a weighted average of the in-sample estimated responses, with relative weighting given by the vector $\pi_{i,\cdot}^*$:

$$\hat{\beta}_i = \sum_{s \in \mathcal{S}} \pi_{i,s}^* \hat{\beta}_s \quad (5)$$

The estimate $\hat{\beta}_i$ has the following properties with respect to the tuning parameter α :

1. $\alpha = 0$ results in a simple average of $\hat{\beta}_s \forall s \in \mathcal{S}$.

2. $\alpha = 1$ results in a weighting of $\hat{\beta}_s \forall s \in \mathcal{S}$ inversely proportional to Euclidean distance in the feature space.
3. As $\alpha \rightarrow \infty$, $\pi_{i,s}^* \rightarrow 1$ for the building s that is closest in the feature space to building i measured by Euclidean distance, while $\pi_{i,s}^* \rightarrow 0$ for all others (i.e. the estimated response of building i tends to exactly that of the building to which it is closest).

Finally, we may apply weighting to the relative importance of each feature by replacing Equation (2) with Equation (6):

$$\gamma_{i,s} = \|\Lambda(\mathbf{x}_i - \mathbf{x}_s)\| \quad (6)$$

Here, Λ is a diagonal $m \times m$ matrix. Let $\lambda = \text{diag}(\Lambda)$ where λ is an m -element vector of weights associated with each feature in the vector \mathbf{x}_i . Larger, and smaller, values of λ allow for features to have more, or less, influence in the similarity calculation, respectively. This weighting may be applied if the analyst has reason to believe that certain features may be more important than others in determining demand response potential. Note that using the unweighted version in Equation (2) implicitly assumes that each feature is equally important in determining demand response potential.

3. Overview of data

The data used in this analysis were the following: i) daily cooling demand served from the CEF (2017-2021, inclusive); ii) daily cold water demand for each of the 124 buildings in the district energy system for the same time horizon; and iii) a dataset of features associated with the buildings in the energy system.

3.1. Daily system load served

The daily total system cooling load served is a key metric of interest. This is the total amount of energy delivered via chilled water to serve the cooling demand of the district energy system. Here, daily values are considered because during hot summer days cooling load in afternoon hours leads to the greatest demand on the system while water can be chilled overnight to refill the cold water storage tanks. Thus a daily cycle is the most appropriate for determining the maximum capability of the system to serve cooling demand.

Data observations are collected at the hourly level and then summarized at the daily level. Over the five year period five hourly observations were unavailable and filled using linear interpolation. Six days were known to have some amount of curtailed load, and counterfactual non-curtailed loads were approximated using a linear model described in the appendix.

3.2. Building level demand data

Hourly cooling demand was collected for all 124 buildings in the study for all hours of the years 2017-2021, inclusive. Data were summarized at the daily level. Procedures for dealing with missing data are described in Appendix A. Table 2 provides a summary of the building level cooling load data over the five years.

Statistic	2017	2018	2019	2020	2021
Mean	17.54 (43.02)	16.80 (38.79)	18.93 (47.55)	19.35 (53.88)	17.56 (48.00)
P10	8.86 (28.24)	7.85 (22.52)	8.99 (25.87)	9.13 (29.40)	8.79 (27.50)
Median	15.86 (41.02)	15.75 (38.23)	16.59 (43.63)	17.46 (51.39)	15.47 (44.29)
P90	28.19 (61.51)	27.53 (58.23)	32.02 (76.01)	31.90 (82.36)	28.80 (73.78)

Table 2: Summary of building level daily cold water demand data (GJ/day). Values indicate mean (standard deviation) of the specified statistic across the 124 buildings for the specified year. “P10” and “P90” indicate the 0.1 and 0.9 quantiles, respectively.

3.3. Building feature data

Feature data were collected for all in-sample and out-of-sample buildings in the district energy system. Feature values for the in-sample buildings and a summary of the building feature data are displayed in Table 3. The building feature data were normalized to have mean zero and standard deviation of one before computing similarity scores as described in Subsection 2.3.²

Figure 2 displays the correlation matrix of the building feature data. The associations here show that in our dataset buildings with high surface area to volume ratios also tended to have higher cooling energy intensity. Buildings with high surface area to volume ratios also tended to be smaller buildings, measured by floor space. Buildings with high summertime peaking tended to have lower cooling energy intensity. Buildings with a high proportion of office space tended to be older, and buildings with a high proportion of offices tended to have less research laboratory space.

4. Results

In this section we present the results of the estimation procedure with feature weighting as described in Subsection 2.3. First, we discuss the process used for selection of model parameter value α and weighting vector λ . Secondly, we provide an overview of the estimates of demand response potential for buildings

²Surface area to volume ratio was approximated using number of floors and floor area by story. We assume the building footprint to have equal length and width with footprint area defined as the average of floor area by story and an assumed story height of 4.27 m (14 feet).

Building/ Statistic	Area (m ²)	LAB (prop.)	CLS (prop.)	LIB (prop.)	OFF (prop.)	Intensity (MJ/m ² /day)	Year	Peak (ratio)	SA-Vol (ratio)
Feature values for observed buildings									
BLDG-1	13,427	0	0	0.005	0.267	0.544	2000	2.27	0.057
BLDG-2	2,640	0.118	0	0	0.423	0.418	1893	2.23	0.105
BLDG-3	15,768	0.061	0.083	0.307	0.133	0.228	1966	2.02	0.048
BLDG-4	9,780	0.202	0.036	0	0.314	0.899	1998	1.63	0.060
BLDG-5	6,975	0.196	0.030	0	0.361	4.267	1965	1.05	0.061
BLDG-6	7,108	0.297	0	0	0.305	1.912	1963	0.70	0.061
Summary statistics for all building in system ($N = 124$)									
Mean	8,312	0.134	0.034	0.037	0.265	2.456	1969.0	1.57	0.087
P10	1,143	0.000	0.000	0.000	0.022	0.481	1901	0.63	0.054
Median	5,280	0.014	0.000	0.000	0.230	1.355	1979	1.53	0.074
P90	15,920	0.439	0.083	0.051	0.543	5.115	2013	2.37	0.142

Table 3: Feature values for observed buildings and summary statistics of feature values for all buildings in the system ($N=124$). LAB: research laboratory, CLS: classroom, LIB: library facilities, OFF: office, SA-Vol: surface area to volume ratio, prop.: proportion.

in the district energy system. Finally, we discuss the simulation of a counterfactual demand response program under global temperature set point adjustments across all buildings in the system over the data time horizon (2017-2021, inclusive) with emphasis on demand reduction for the highest demand year, 2020. All calculations presented were implemented in the Python programming language (version 3.8.13) with the following packages: Numpy (version 1.21.1), Pandas (version 1.3.1), and Scikit-learn (version 0.24.2) [37].

4.1. Leave-one-out validation of parameter values

The α parameter in the above calculation is an important tuning parameter that penalizes the distance of an out-of-sample building to the in-sample buildings. To determine the values of α and λ that are most appropriate for this study we have implemented a LOO validation process. However, because we only have six in-sample buildings upon which we can train our data, we must be sensitive to the possibility of overfitting to our observed data. In this subsection we discuss how we select a vector λ , and then select our preferred value of α .

To implement the LOO validation process, for each building $s \in \mathcal{S}$ we remove this building from the set \mathcal{S} and compute the predicted $\hat{\beta}$ for that building for a range of values of α using the proposed estimation algorithm and for a set of candidate values of λ . The first candidate value of λ is “uniform weights” where each value in the vector is 1, corresponding to the unweighted formulation in Equation (2). An additional




Figure 2: Correlation matrix of feature data (LAB: research laboratory, CLS: classroom, LIB: library facilities, OFF: office, SA-Vol: surface area to volume ratio)

10 candidate values of λ were tested as “weighted” vectors where weights were chosen based on assumed importance of the features. The set of candidate weight vectors was developed by the authors and is intended to capture the relative importance of features in determining flexibility. The specific values are qualitatively based upon a review of relevant literature and researcher judgement based on familiarity with the field. See Appendix B for the specific values used in the candidate weight vectors and additional discussion.

Figure 3 displays the results of the tuning procedure. Note that the names (1 to 10) of weight vectors are simply labels and do not indicate ordering. For each value of alpha the error is defined as the difference between the actual value of $\hat{\beta}$ for the building that was estimated with experimental data, and the value of $\hat{\beta}$ that was computed with the similarity score vector computed as in Subsection 2.3. For a given value of α the square of this error is computed and averaged across the six buildings, and the square root is taken, resulting in the root mean squared error (RMSE) pictured in Figure 3. In the figure, the RMSE for uniform weights and the selected weights are shown in bold.

For all values of α , all candidate weight vectors had lower RMSE than the uniform weights vector. For values of α less than 4 we found similar RMSE in the LOO procedure across the ten candidate weight vectors with “weights 5” resulting in the lowest RMSE for these values. For high values of α (greater than 31) “weights 5” also had the lowest RMSE. This weighting vector achieved the lowest RMSE of 0.0688 at $\alpha = \infty$. As discussed in Subsection 2.3, $\alpha = \infty$ results in a matching algorithm that assigns one of the six observed $\hat{\beta}_s$ values to each of the unobserved buildings. However, in the building stock, we have high heterogeneity across features (building vintage and size, for example), and some buildings may be further away (as measured by Euclidian distance) in the feature space than those included in the observed set. Thus a high penalty on distance in the feature space is likely to overfit to the small set of observed buildings. For these reasons we expect that the actual response will be better represented by some weighted average of the

$\hat{\beta}_s$ given by a lower value of α . See Appendix C for more discussion on how the selection of α affects the estimated demand response potential values, and see Appendix D for a sensitivity analysis that tests how the choice of α affects the system-wide flexibility result.

The vector “weights 1” resulted in lowest RMSE for values of α from 4 to 31, and resulted in minimum RMSE at $\alpha = 14$. In Figure 3 we can see the lower RMSE of the “weights 1” vector over this range. Because this vector achieved the lowest RMSE for this range of values of α this was the weighting vector selected to be used for the remainder of the analysis. The values of the selected weights are displayed in Table 4. The uniform weights vector resulted in minimum RMSE at $\alpha = 6$.

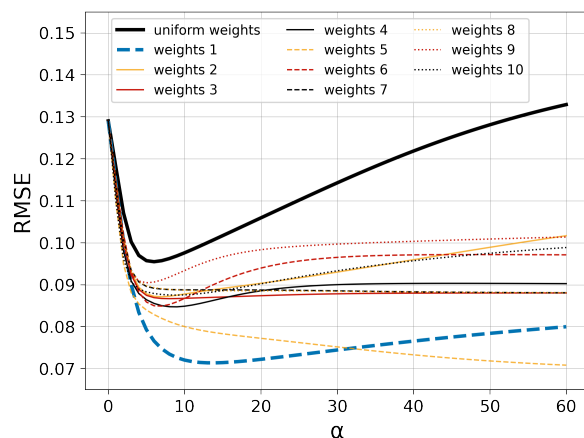


Figure 3: Root mean squared error (RMSE) versus values of parameter α (“uniform weights” and the selected weight vector, “weights 1”, pictured in bold)

Area	LAB	CLS	LIB	OFF	Intensity	Year	Peak	SA-Vol
1.0	0.5	0.5	0.5	0.5	2.0	1.0	1.5	1.0

Table 4: Values of λ for the selected weighting vector (LAB: research laboratory, CLS: classroom, LIB: library facilities, OFF: office, SA-Vol: surface area to volume ratio)

For the minimum error values of α we present the predicted versus actual values of $\hat{\beta}$ in Figure 4. In the figure we observe that the selected weighted prediction outperforms the uniform weight predicted values for four of the six buildings. The only building for which the weighted prediction does significantly worse is BLDG-3. Note, that in this formulation the prediction is a weighted average of the other buildings. Here, across the feature vector BLDG-3 was more similar to the high response buildings than the low response buildings and resulted in this error. For BLDG-2, we observe that the estimation procedure results in a considerable under-estimation of flexibility, however, because the estimate is a weighted average of the five remaining buildings, the estimate can not be closer than the value of the second largest response building. However, the estimate does improve under the weighted estimation process.

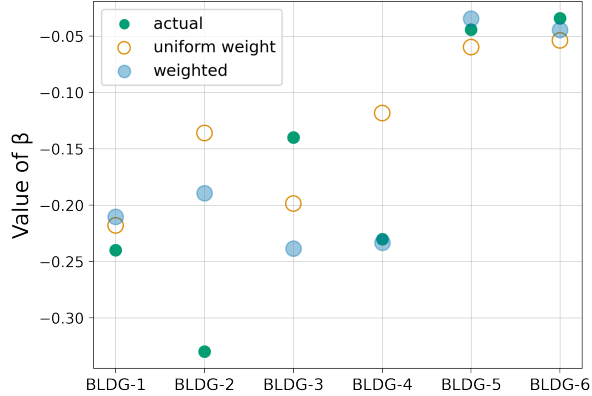


Figure 4: Predicted versus actual values of $\hat{\beta}$ (uniform weight: $\alpha = 6$; weighted: $\alpha = 14$)

The results of the LOO tuning procedure suggest that the similarity score procedure is likely to provide reasonable estimates of demand flexibility in out-of-sample buildings. These results also suggest that the weighted score is likely to outperform the unweighted version, however, the high α value corresponding to minimum RMSE applies a very high penalty on distance in the feature space, closely corresponding to a matching algorithm. Due to the small number of in-sample buildings and the risk of overfitting to the observed sample, we use $\alpha = 3$ for the analysis in this paper with weighting vector 1. In Appendix C we provide more details about the choice of this α value, and in Appendix D we provide a sensitivity analysis to investigate how the choice of parameter value α and weight vector λ affect our results.

4.2. Estimates of building-level response potential

Values of $\hat{\beta}$ were estimated for all 118 unobserved buildings in the system using $\alpha = 3$, and the weighting vector λ as defined in Subsection 4.1. The corresponding estimated load reduction (%) under a 1.1°C (2°F) daily set point increase was computed for each building in the system.

Figure 5 displays a histogram of estimated demand response potential for each of the 124 buildings in the system, including the six observed buildings. The average estimated response across all buildings in the system (including observed experimental buildings) was 16.1% with a standard deviation of 5.3%, and a median response of 16.0%. Half of all estimated responses fell between 12.6% and 19.3%, while 90% of all estimated responses fell between 7.0% and 25.2%. The minimum and maximum estimated response values were those of the minimum and maximum responding buildings in the observed sample with values of 3.3% and 28.1%, respectively.

Table 5 reports the pairwise correlations between the estimated demand response and building features. The high correlation between the “Peak” feature and the estimated demand response potential indicates that for buildings that had large peak to trough load relative to average load on summer weekdays, we estimated higher demand response potential, with a relatively high correlation value (0.796). This result




Figure 5: Histogram of estimated building demand response potential for $\alpha = 3$

aligns with the expectation that buildings that tend to have greater afternoon increases in cooling demand will have more potential to reduce that demand with a thermostat set point increase. In addition, buildings with a high proportion of office space tended to have higher estimated demand response potential but with a relatively weaker linear association. Meanwhile, buildings with higher energy intensity and buildings constructed more recently tended to have lower estimated demand response potential. Note that both of these features were negatively correlated with the “Peak” value (see Figure 2).

Area	LAB	CLS	LIB	OFF	Intensity	Year	Peak	SA-Vol
-0.123	-0.123	0.077	-0.087	0.345	-0.278	-0.364	0.796	0.118

Table 5: Correlation of building characteristics and estimated demand response ($\alpha = 3$) (LAB: research laboratory, CLS: classroom, LIB: library facilities, OFF: office, SA-Vol: surface area to volume ratio)

4.3. System-wide demand response potential

Using the demand response estimates presented in Subsection 4.2 we can estimate the load reduction potential of a counterfactual demand response program in the district energy system. We assume that the load control action under this demand response program is a system-wide 1.1°C thermostat set point increase on the 10 highest demand days of the year. We then construct a modified load duration curve to visualize maximum load before and after the demand response program. The pre- and post-demand response load duration curves for 2020 are displayed in Figure 6 for the top 100 load days of that year.³ Note that although the experiments discussed in Subsection 2.1 were conducted in 2021, we focus on the 2020 calendar year in this section. This is because 2020 experienced the highest daily cooling demand over the time horizon and is thus most relevant for design decisions relative to system installed capacity. In this analysis, we assume

³The highest demand days in 2020 are of particular importance in this analysis. Note that none of the days in 2020 experienced known curtailments where load served was estimated as described in Subsection 3.1.

that the 10 highest demand days of the year are known, and these are the days in which demand response is implemented. In practice, the highest demand days must be forecasted in advance, but we do not address the issue of demand forecasting in this work.




Figure 6: 2020 actual and counterfactual load duration curves, top 100 days, with 95% confidence interval pictured (BAU: “business as usual”). “Min. Effect” and “Max. Effect” indicate the load duration curve that would result if all unobserved buildings have the demand response potential of the building with the smallest (BLDG-6) and largest (BLDG-2) estimated response, respectively. The top panel shows the counterfactual demand in the same order as actual demand. The bottom panel shows the counterfactual demand reordered from highest to lowest.

The 2020 calendar year experienced the highest daily cooling load over the time horizon reaching 5.588 TJ (441.3 kiloton-hours). Under the demand response program, estimates indicate that maximum daily load would be reduced by 13.47% to 4.835 TJ, with a 95% confidence interval of (11.70%, 15.24%). We estimate the asymptotic distribution of the counterfactual load using the delta method, as described in Appendix E. This confidence interval treats the values of π_{i_s} as fixed constants, capturing the asymptotic variance of the estimates of $\hat{\beta}_s \forall s \in \mathcal{S}$.

In the top panel of Figure 6 we see the effect that the demand response program has on the 2020 demand curve.⁴ At the left of the plot, where the orange line is shifted downward, load has been reduced due to the

⁴In the bottom panel we have reordered the counterfactual load duration curve to show how the load duration curve changes

load control action. The percent reduction for each day will differ slightly due to variability in the proportion of load each building accounts for on that day. For days that were not in the top 10 load days, the load duration curve remains unchanged. The shaded area in the plot denotes the 95% confidence interval for each day. The dashed lines denote the modified load duration curve if all unobserved buildings were to have the estimated response of the least and most responsive building in the observed sample, labeled as minimum and maximum effect, respectively.

An important point illustrated in the load duration curves in the top panel of Figure 6 is that the steepness of the curve on the top days enables a large reduction in maximum daily demand with a small number of demand response days. Observe that the daily load on the highest demand day that is *not* a demand response day is lower than the estimated demand on the highest demand response day after the load reduction. For 2020, we have pictured 10 days of demand response, but this still would have been true with only 8 days of demand response. This means that if the objective is to reduce peak daily load, the amount of load reduced is the full amount reduced on the highest BAU demand day. If the load duration curve were not as steep, with a similar amount of load reduction, daily load on the highest non-demand response day may be higher than all demand response days. In this case the maximum daily load reduction would be the difference between the maximum BAU load, and maximum daily load of the non-demand response days.

The bottom panel of Figure 6 shows what the load duration curve would look like after demand response, ordered from highest to lowest load. In Section 1 it was noted that the 2020 BAU daily load duration curve suggested that 10% of system capacity would be used only on five days of the year, and 20% of capacity needed only on 15 days of the year.⁵ In contrast, if this demand response program were implemented and the system were designed for the post-DR load duration curve, 5.6% of system capacity would be used on only five days, while 11.4% of capacity would be necessary for the top 15 days.

In Appendix D we provide a sensitivity analysis that includes results for selected values of α and results for maximum daily load reduction for all years in the time horizon. The estimated reduction in peak daily load in 2020 ranged from 9.98% ($\alpha = \infty$) to 15.46% ($\alpha = 0$). Estimates for peak daily load reduction varied little across different years for a given value of α . See Appendix D for more details.

5. Discussion

5.1. Demand response program design requires system level estimates

With cooling demand growth over time due to expansion of the system and/or increasing temperatures as a result of climate change, district HVAC system operators may be faced with a decision to invest in

under this program.

⁵This calculation is under the assumption that the system was designed precisely for the maximum load day of the year.

costly capacity upgrades or implement a demand response program to defer these upgrades. In the short-run, because of the impossibility of a capacity increase, some systems may be faced with the possibility of forcibly curtailing the provision of cooling services when capacity limits are exceeded, and a demand response program may be a preferable alternative. However, implementing thermal demand response in a district energy system can be challenging. In addition, even if thermostat set point changes are small, deviating from BAU set points can impose some degree of inconvenience on building occupants. The benefits of such a program can be substantial, and to quantify these benefits to justify its implementation, a reliable estimate of the magnitude of system-wide demand reductions is needed.

We propose a methodology to compute such estimates based on a series of set point changes in a small subset of buildings and a set of building features for all buildings in the energy system. In many district energy systems, such as the one in this study, there is a high degree of heterogeneity across buildings in the system. This heterogeneity in structural characteristics and operational features results in significant heterogeneity across buildings in the magnitude of the demand response for the same thermostat temperature set point increase. In the observed buildings, the demand response potential estimates ranged from 3.3% to 28.1% cooling demand reduction under a 1.1°C (2°F) daily cooling set point adjustment. This large difference across buildings in demand response potential makes it important to determine if the unobserved buildings have high or low flexibility potential. Leveraging building feature data to extend the experimental results to the entire building stock leverages the heterogeneity in building characteristics to estimate savings in a large, heterogeneous group of buildings.

The methodology proposed here is a data-driven alternative to physics-based models and provides certain advantages. First, we measure actual, realized energy use changes in the building. Second, physical models require a large amount of information about all buildings to reliably estimate the demand reduction for a given building, while our methodology makes use of available data without the need for assumptions about physical model parameters.

The methodology proposed here, may also be used to target demand response implementation under a budget constraint that does not allow for implementation of demand response across the entire system. If a budget constraint allows for implementation in only a subset of buildings, the methodology can be used to identify the high demand flexibility buildings in the system.

5.2. Load duration curves inform demand response potential to reduce system capacity needs

In many systems, adjusting thermostats on a small number of days can lead to a large reduction in capacity needs. Inter-annual variability in cooling demand will mean that the highest demand years will require the most days of load flexibility while lower demand years require less, and some years needing no intervention at all, to maintain the daily cooling demand below an available capacity limit.

By examining annual daily cooling load duration curves we can see how load control can reduce capacity

needs in the system under study. Figure 7 displays daily cooling load duration curves (top 150 days) for 2017-2021. Indicated in the plot are lines for the minimum installed capacity necessary to serve all load both with and without a demand response program, that includes the load control action of a system-wide 1.1°C set point adjustment. Over this time horizon, 2020 was the year with highest daily demand with a peak daily demand of 5.588 TJ. This peak cooling demand day over the time horizon determines the minimum capacity needed to reliably serve cooling demand. Results of this study suggest that implementing a 1.1°C set point increase on that day would reduce cooling demand by 0.753 TJ to 4.835 TJ. This result implies that the minimum system capacity required is 15.6% higher without a demand response program when compared to the reduced cooling demand under the program.




Figure 7: Load duration curves 2017-2021 with estimated peak reduction.

The load duration curves pictured in Figure 7 show how inter-annual variability and a small number of extreme demand days drive capacity needs over the time horizon. In addition, the figure shows that achieving a large reduction in capacity needs would require implementing thermostat set point adjustments on a very small number of days over the five year period. The load duration curve for 2020 exhibits a steeply declining slope moving to the right from the vertical axis suggesting that only eight demand response days would need to be called to reduce demand on all days of that year below the 4.835 TJ threshold. We may also observe that with only five days of demand response in 2019, two days in 2017, and no required demand response in 2018 or 2021, the system could have maintained maximum daily load under this level. This suggests that with only 15 demand response days over the five year period from 2017-2021, the system could have maintained installed capacity needs below the lower 4.835 TJ level. With installed capacity costs for such district energy system facilities measured in the millions of dollars this implies a large financial benefit for an infrequent load control action.

The authors note that the estimates here assume perfect foresight and under real-world implementation

a larger number of demand response days may be required due to forward-looking uncertainty. The analysis also does not account for load growth over time or changing climate patterns. However, the result demonstrates that an infrequent load control action with low impact to occupants can have a very meaningful impact on system capacity needs.

5.3. Leveraging demand-side flexibility in electricity systems with high renewable deployment

In California, high deployment of solar generation has led to electricity supply challenges in the early evening hours on high demand days. When solar generation comes offline while electricity demand is still high, a large amount of ramping capacity, supplied mostly by natural gas peaking units, is needed to balance supply and demand. Limited dispatchable capacity led to forced load curtailment during the summer of 2020 [2]. Mobilizing demand-side flexibility through intelligent HVAC control can be a valuable resource as the California Independent System Operator (CAISO) and other electricity system operator territories continue to experience higher penetration of renewable generating technologies.

In the system under study, days in which the local district energy system is experiencing high cooling demand conditions correspond closely to the days in which the bulk grid is most in need of additional capacity or demand reductions. These conditions tend to occur on hot summer days when air conditioning demand is high. Figure 8 displays daily average day ahead Locational Marginal Price (LMP) in the CAISO market (24 hour average of DLAP-PGAE) and daily CAISO peak net load for all days of 2020.⁶ CAISO net load is defined as system load less wind and solar generation. The points in red indicate the top ten cooling load days for the district energy system. We observe in the figure that these highest demand days are days in which the CAISO peak net load is very high and prices are high. If a demand response program were designed such that thermal demand response days were called on the ten highest demand days in the local system, not only would the system be reducing installed capacity requirements, but the daily reduction in electricity demand would be providing needed demand reduction to the CAISO system.

The university campus in this study reaches peak hourly electricity demands of over 50 MW during the summer, with nearly 20 MW of this accounted for by the CEF chillers on high cooling demand days. Reduced demand for chilled water on high demand days enables flexibility in cold water chiller operational schedules. The on-site cold water storage provides thermal energy storage enabling intertemporal shifting of electricity demands and the ability to greatly reduce system electricity demand when it is most valuable.

In recent years, deployment of lithium ion batteries for stationary storage has increased rapidly, and costs of battery storage have fallen significantly. Battery energy storage systems are likely to be an important component of a future grid with high deployment of renewable energy technologies. However, despite falling costs, it remains extremely costly to provide all needed demand side flexibility with battery systems.

⁶“DLAP-PGAE” is the Pacific Gas and Electric default load aggregation point.




Figure 8: 2020 Day ahead LMP versus daily CAISO peak net load with top 10 campus cooling load days in red (LMP is 24-hour average of DLAP-PGAE; net load is defined as system load less wind and solar generation). Data source: CAISO

In this study, the 2020 load control simulation indicated a daily load reduction of 0.753 TJ on the highest demand day for the 24 hour period. The CEF in this system includes both conventional chillers and heat recovery chillers. The overall efficiency of the CEF ranges from 47.39 - 69.90 kWh per GJ.⁷ This suggests an approximate electrical energy reduction of 35.68 - 52.63 MWh for the day. As of this writing, costs of lithium ion batteries were approximately \$130 - \$140/kWh (hardware costs only) [38]. We also consider a range of discharge efficiencies from 92% - 100%. For this range of assumptions, the system would require a battery with capital investment cost of \$4.6 to \$8.0 million USD to provide an equivalent amount of demand response. The authors note that this is a very rough estimated range, simply to provide an order of magnitude comparison. For example, this figure does not include balance of system costs, operating costs, or cost of charging. It also does not consider the additional services that a dispatchable energy storage system can provide. However, in comparison to a battery system, the costs of implementing thermal demand response in commercial buildings can be very low. This high-level estimate suggests that achieving demand-side flexibility in electricity systems needed to enable large scale renewable deployment at lowest cost should include a strong emphasis on mobilizing thermal demand response programs.

6. Conclusion

Thermal demand response in commercial buildings holds potential to provide large amounts of needed demand-side electricity demand flexibility as electricity grids in the United States and around the world

⁷The efficiency range of the CEF stated here is based on both the engineering design documentation of the CEF, and an estimate based on observed operational data.

decommission conventional fossil fuel generators in favor of renewable energy resources. Moreover, in many district energy systems where capital expenditure for installed capacity is considerable, a large proportion of installed capacity is used on only a small number of days. Demand flexibility through temperature set point adjustments can be implemented where even relatively small and infrequent temperature set point adjustments can translate into large savings while simultaneously providing valuable demand flexibility.

Methods are needed to estimate the specific savings that can be achieved in systems, in particular where high heterogeneity exists among buildings in the system under consideration. Such estimates are required to enable relevant stakeholders to balance potential benefits against costs of implementation. In this work, we have demonstrated a methodology by which experimental measurements of load flexibility can be extended to the level of a district energy system. The methods presented here are not specific to district energy systems and may be implemented for large building stocks to estimate demand response in any well defined group of buildings. The methodology can be generalized to any quantitative or categorical building features that may be available in other settings, and with an increased number of buildings in the observed set, the estimates of demand response potential will increase in accuracy.

Under a range of assumptions, we find that the load control action of a system-wide 1.1°C (2°F) thermostat set point increase results in an estimated reduction in peak daily cooling load for the highest demand day of the year of at least 10%. Implementing this load control action would accrue savings through reduction of investment on installed capacity at the district energy system's CEF, reduction of energy costs, and potential reduction of demand charges. Moreover, we have shown that the highest demand days in the district energy system tend to correspond to the days in which generation capacity is most scarce in the bulk grid. Such a demand response program would then also reduce demand during the times in which this capacity is most needed on the grid.

The methodology presented in this study estimates cooling demand flexibility in a large number of unobserved buildings based on observable features and results of real-world flexibility experiments in a small number of observed buildings. This can be a valuable complement to simulations of building energy behavior using physics-based models. An interesting area of study that we leave for future work is to integrate building simulation tools and empirical, data-driven methods. Details of district energy system operation under load control actions, such as investigating the efficiency impacts of dynamic set point adjustments, will also be an interesting area of research.

In this study we have not investigated the effect of set point adjustments on occupant comfort and this will be a valuable area of future research. Moreover, here we have limited our attention to a 1.1°C thermostat set point increase universally implemented across all buildings in the system. In future work we plan to investigate additional load control actions including set point changes of greater or lesser magnitude and heterogeneous set point changes that differ by building. These additional actions are likely to enable deeper load reductions while minimizing the costs imposed on building occupants.

Acknowledgements

The authors would like to thank TotalEnergies for their funding of this research, and Stanford Land, Buildings and Real Estate, for their operational support.

Appendix A. Methods for data cleaning and completion of raw building data

In this study we use real-world energy usage data over five years (2017-2021, inclusive) from two data sources: 1) cold water delivered to the system by the CEF of the district energy system and 2) the cold water consumption measured at each of 138 meters at 124 buildings in the system. The data for both sources are in units of ton-hours and recorded at the hourly level. A number of data points were missing or unreliable due to recording errors or other events. In this section we describe the processes that were used to verify the validity of available data and estimate data points that were unavailable.

For the CEF the dataset was nearly complete. In Section 3.1 it was noted that only five hourly data points were missing for the entire time horizon and filled using linear interpolation of neighboring datapoints. A sum of hourly data for each day was computed to find the daily cold water served for each day. For six days where some unknown amount of load curtailment occurred, a linear model was used to estimate counterfactual load:

$$z_t = \nu_0 + \nu_1 T_t + \eta_t \tag{A.1}$$

where z_t is the daily CEF load (ton-hours) on day t , T_t is the daily average temperature, and η_t is an error term with mean zero. The model was fit with all available data points (excluding curtailment days), and the estimated parameters were used to compute expected load for those days. Log-linear and quadratic models were also tested. The linear model was selected to complete the dataset because it was both the most accurate and the most conservative of the models.

For the 138 meters at the 124 buildings served by the cold water loop, hourly cooling energy delivered is recorded at the point of delivery to the building. Here the difference in cold water temperature in the send and receive pipes and flow rate are measured to compute cooling energy delivered, recorded in units of ton-hours. A number of data points were not available due to data write failure for unknown reasons. Any data points that were in excess of three times the P99 value were assumed to be erroneous and discarded. For data points not available, wherever three or less consecutive hours were identified the values were linearly interpolated. For buildings where cold water service curtailments were known to have occurred, or temperature set point increases were implemented, these data points were discarded.

For the unavailable data points, a random forest model was implemented to estimate missing values at the hourly level. The features used to fit the random forest model were CEF hourly cooling demand served, hourly temperature (measured onsite), year indicator variables, month indicator variables, hour

indicator variables, and weekday/weekend indicator variables. In the completed dataset 4.3% of total load was accounted for by estimated data points over the study time horizon. These hourly data were then summarized at the daily level for a complete dataset of daily cooling load for each building.

As noted above, the CEF cooling demand served dataset was nearly complete. In order to ensure that the building level data conformed to the CEF level data a scaling multiplier was applied to the daily building data such that the sum of building data would precisely match the CEF data. The same scaling factor was applied to all buildings with one scaling factor per day. The mean scaling factor was 0.969, with standard deviation of 0.032. Data in units of ton-hours were converted to joules in the main text of this study.

Appendix B. Selection of weighting vector λ

In Section 4.1 we discuss the results of the tuning procedure by which we select the weighting vector λ . Two possible approaches to choosing the weighting vector are: i) optimize for weights that minimize the validation RMSE, or ii) test a set of candidate weights and choose the candidate vector that minimizes the validation RMSE. We have chosen the latter approach due to the concern of overfitting to the small sample of six observed buildings. We developed a set of 10 candidate weight vectors and chose the best performing (in terms of RMSE) as the preferred weight vector.

The candidate weighting vectors are displayed below in Table B.6. The features “LAB”, “CLS”, “LIB”, and “OFF” all were indicators of the occupancy, equipment, and types of operations in the buildings and thus were all equally down-weighted across the candidate vectors. The review of literature suggested that “Intensity” and “Peak” were relatively important features in predicting cooling demand flexibility and thus received higher weights across most candidate vectors [6, 18, 19, 20, 21]. Candidate vector 1 achieved the lowest RMSE for moderate values of α (Figure 3) and was thus the chosen weighting vector.

Appendix C. Selection of α parameter value

The value of the tuning parameter α applies a penalty on distance in the feature space when computing estimated values of $\hat{\beta}$ for the unobserved buildings. Under the tuning procedure described in Section 4.1 we found that a value of $\alpha = 14$ achieved the lowest RMSE for the chosen weighting vector. However, due to concern that a high α value will lead to overfitting to the small number of observed buildings we selected $\alpha = 3$ as our preferred parameter value.

Figures C.9 - C.16 display histograms of estimated daily building cooling demand flexibility (%) for selected values of α under the chosen weighting vector. In Figure C.9 we observe that the parameter value $\alpha = 0$ results in a simple average of observed buildings assigned to all unobserved buildings. As the value of α increases there is an increased penalty on distance in the feature space when computing the estimates. At the limit when α approaches infinity, the calculation becomes a matching algorithm where each unobserved

Candidate Index	Area (m ²)	LAB (prop.)	CLS (prop.)	LIB (prop.)	OFF (prop.)	Intensity (MJ/m ² /day)	Year	Peak (ratio)	SA-Vol (ratio)
1	1	0.5	0.5	0.5	0.5	2	1	1.5	1
2	1	0.5	0.5	0.5	0.5	1	1	1	1
3	1	0.5	0.5	0.5	0.5	1.5	1.5	1.5	1.5
4	1	0.5	0.5	0.5	0.5	1.5	1.5	2	1
5	1.5	0.5	0.5	0.5	0.5	1	1	2	1.5
6	1	0.25	0.25	0.25	0.25	2	1	1.5	1
7	1	0.25	0.25	0.25	0.25	1	1	1	1
8	1	0.25	0.25	0.25	0.25	1.5	1.5	1.5	1.5
9	1	0.25	0.25	0.25	0.25	1.5	1.5	2	1
10	1.5	0.25	0.25	0.25	0.25	1	1	2	1.5

Table B.6: Definitions of candidate feature weighting vectors. LAB: research laboratory, CLS: classroom, LIB: library facilities, OFF: office, SA-Vol: surface area to volume ratio, prop.: proportion.

building is assigned the flexibility estimate of the building to which it is the closest in the feature space. This is shown in Figure C.16.

Figure C.15 displays a histogram for $\alpha = 14$. Although this value achieved the lowest RMSE, in the figure we observe that the relatively high tuning parameter led to a strongly multi-modal distribution of estimated effects. Due to concern that this may lead to overfitting to the small number of observed buildings, we chose to select the parameter value that resulted in a nearly bell-shaped distribution of estimated flexibility.

Parameter values $\alpha = 2$ and $\alpha = 3$ (Figure C.11 and Figure C.12) were the values for which the estimated flexibility values most closely approximated a bell-shaped distribution. Parameter value $\alpha = 2$ did not result in any estimated flexibility values in the range of 4.4% to 8.7%, and in addition $\alpha = 3$ had a lower RMSE in the LOO tuning process. Therefore we selected $\alpha = 3$ as our preferred tuning parameter value. We investigate the effect the choice of tuning parameter value has on our results in Appendix D.

Appendix D. Sensitivity Analysis

A sensitivity analysis was conducted to investigate the degree to which the choice of the parameter value α and weight vector λ affected our results. To test the sensitivity to the selected α value, we estimated system-wide peak load reduction under a range of selected parameter values, and for each year of the time horizon. All results here use the selected vector of feature weights as described in Section 4.1. The results are displayed in Table D.7.

Across all selected parameter values and years, the highest estimated percentage reduction was 15.47%,

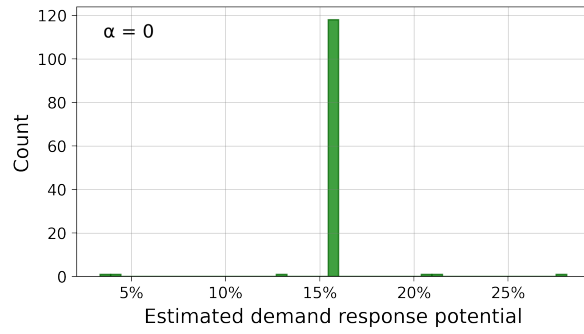


Figure C.9: Distribution of estimated demand response potential for $\alpha = 0$




Figure C.10: Distribution of estimated demand response potential for $\alpha = 1$




Figure C.11: Distribution of estimated demand response potential for $\alpha = 2$




Figure C.12: Distribution of estimated demand response potential for $\alpha = 3$




Figure C.13: Distribution of estimated demand response potential for $\alpha = 4$

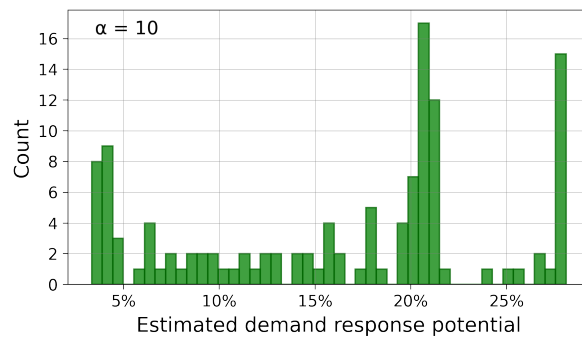


Figure C.14: Distribution of estimated demand response potential for $\alpha = 10$




Figure C.15: Distribution of estimated demand response potential for $\alpha = 14$




Figure C.16: Distribution of estimated demand response potential for $\alpha = \infty$

while the lowest was 9.97%. The highest percentage reductions were estimated for $\alpha = 0$, where the estimated flexibility was simply an average of the observed building effects. The system-wide estimated flexibility decreased with higher α parameter values. The lowest estimates were for $\alpha = \infty$. In Figure C.16 observe that more buildings had estimated flexibility of greater than 15% than had estimated flexibility of less than 15%. This implies that although less buildings had low estimated flexibility, this smaller number of buildings accounted for a larger share of cooling demand. Also note that across all parameter values the estimated peak day load reduction varied little in percentage terms across years. The 95% confidence intervals were slightly less than $\pm 2\%$ across all estimates.

Alpha	Year	Peak Daily Load	Peak Daily Load with DR (C.I.)	Percent Reduction (C.I.)
0	2017	4.921	4.161 (4.066, 4.256)	15.43% (13.51%, 17.36%)
0	2018	3.938	3.330 (3.255, 3.406)	15.44% (13.52%, 17.36%)
0	2019	5.321	4.498 (4.396, 4.600)	15.46% (13.54%, 17.38%)
0	2020	5.588	4.724 (4.616, 4.831)	15.46% (13.54%, 17.38%)
0	2021	4.488	3.794 (3.708, 3.880)	15.47% (13.55%, 17.39%)
3	2017	4.921	4.260 (4.173, 4.346)	13.43% (11.67%, 15.19%)
3	2018	3.938	3.396 (3.326, 3.466)	13.77% (11.99%, 15.56%)
3	2019	5.321	4.595 (4.501, 4.690)	13.63% (11.85%, 15.41%)
3	2020	5.588	4.835 (4.736, 4.934)	13.47% (11.70%, 15.24%)
3	2021	4.488	3.881 (3.801, 3.960)	13.54% (11.77%, 15.32%)
14	2017	4.921	4.339 (4.248, 4.431)	11.81% (9.96%, 13.66%)
14	2018	3.938	3.460 (3.389, 3.532)	12.14% (10.31%, 13.96%)
14	2019	5.321	4.690 (4.594, 4.784)	11.86% (10.08%, 13.65%)
14	2020	5.588	4.939 (4.839, 5.039)	11.61% (9.82%, 13.40%)
14	2021	4.488	3.965 (3.885, 4.046)	11.65% (9.85%, 13.44%)
∞	2017	4.921	4.378 (4.282, 4.474)	11.02% (9.07%, 12.97%)
∞	2018	3.938	3.516 (3.442, 3.591)	10.72% (8.83%, 12.61%)
∞	2019	5.321	4.772 (4.671, 4.873)	10.32% (8.42%, 12.21%)
∞	2020	5.588	5.030 (4.923, 5.137)	9.98% (8.06%, 11.89%)
∞	2021	4.488	4.041 (3.954, 4.127)	9.97% (8.05%, 11.90%)

Table D.7: Estimated peak daily demand (TJ) and percent reduction by year for selected values of α (C.I.: confidence interval).

To test the sensitivity of our results to the selected weight vector (λ), we compute the estimated system-wide demand reduction for the peak demand day of 2020 using all candidate weight vectors. The results are displayed in Table D.8. Here, all estimates use a parameter value of $\alpha = 3$. Note that candidate weight

vector 1 corresponds to our selected weight vector and our central result. The estimate is relatively stable under different weight vector assumptions ranging from 13.30% to 14.03%.

Candidate Weight Vector Index	Percent Demand Reduction
Unweighted	13.90%
1	13.47%
2	13.98%
3	13.67%
4	13.30%
5	13.63%
6	13.42%
7	14.03%
8	13.70%
9	13.31%
10	13.59%

Table D.8: Estimated 2020 peak daily demand reduction (%) for all candidate weight vectors (λ) for $\alpha = 3$.

Appendix E. Counterfactual demand variance calculation

In this section we describe the method by which we compute the asymptotic variance of system load under the counterfactual demand response program. Where D is the counterfactual load on a certain day, t (with t subscript omitted) as a function of the vector $\hat{\psi}$ we have the following expression:

$$D(\hat{\psi}) = \sum_{i \in \mathcal{I}} d_i \exp(\hat{\psi}_i) \quad (\text{E.1})$$

where d_i is the BAU load on day t for building $i \in \mathcal{I}$, where \mathcal{I} is the set of all buildings in the system. Here, $\hat{\psi}_i$ is the weighted sum of the coefficients of the six in-sample buildings, i.e.:

$$\hat{\psi}_i = \sum_{s \in \mathcal{S}} \pi_{i,s} \hat{\beta}_s \quad (\text{E.2})$$

where $\hat{\beta}_s$ is the estimate of the coefficient of building $s \in \mathcal{S}$ where \mathcal{S} is the set of in-sample buildings. Then, $\pi_{i,s}$ is the weighting of each coefficient $\hat{\beta}_s$ in the estimation of $\hat{\psi}_i$. Under the assumption that the weighting terms are non-stochastic and the estimated building coefficients, $\hat{\beta}_s$ are independent, we have:

$$\sqrt{T}(\hat{\beta}_s - \beta_s) \rightarrow N(0, \text{AVar}(\hat{\beta}_s)) \quad (\text{E.3})$$

where $\text{AVar}(\hat{\beta}_s)$ is the asymptotic variance of estimator $\hat{\beta}_s$. Therefore, the asymptotic distribution of $\hat{\psi}_i$:

$$\sqrt{T}(\hat{\psi}_i - \psi_i) \rightarrow N(0, \sum_{s \in \mathcal{S}} \pi_{i,s}^2 \text{AVar}(\hat{\beta}_s)) \quad (\text{E.4})$$

where ψ_i is the linear combination of the true β_s .

Let us define Ω as the variance-covariance matrix of the vector $\hat{\psi}$. The $\hat{\psi}_i$ are linear combinations of the $\hat{\beta}_s$ and thus for $i \neq j$, we derive the asymptotic covariance of $\hat{\psi}_i$ and $\hat{\psi}_j$ as:

$$\omega_{i,j} = \sum_{s \in \mathcal{S}} \pi_{i,s} \pi_{j,s} \text{AVar}(\hat{\beta}_s) \quad (\text{E.5})$$

Then, using the delta method we can express the asymptotic distribution of $D(\hat{\psi})$ as:

$$\sqrt{T}(D(\hat{\psi}) - D(\psi)) \rightarrow N(0, \nabla D(\psi)^T \Omega \nabla D(\psi)) \quad (\text{E.6})$$

where each element of the gradient vector $\nabla D(\psi)$ is given by:

$$\nabla D(\psi)_i = d_i \exp(\psi_i) \quad (\text{E.7})$$

which is computed using $\hat{\psi}_i$ as an estimate of ψ_i .

References

- [1] EIA, Monthly Energy Review, April 2021 (2021).
- [2] F. A. Wolak, Long-term resource adequacy in wholesale electricity markets with significant intermittent renewables, *Environmental and Energy Policy and the Economy* 3 (1) (2022) 155–220.
- [3] J. Barton, S. Huang, D. Infield, M. Leach, D. Ogunkunle, J. Torriti, M. Thomson, The evolution of electricity demand and the role for demand side participation, in buildings and transport, *Energy Policy* 52 (2013) 85–102.
- [4] California Energy Commission, Energy reports: Electricity consumption by entity, <http://www.ecdms.energy.ca.gov/elecbyuti1.aspx>, Last accessed on 2022-09-06 (2022).
- [5] K. Mařík, J. Rojíček, P. Stluka, J. Vass, Advanced HVAC control: Theory vs. reality, *IFAC Proceedings Volumes* 44 (1) (2011) 3108–3113.
- [6] R. Yin, E. C. Kara, Y. Li, N. DeForest, K. Wang, T. Yong, M. Stadler, Quantifying flexibility of commercial and residential loads for demand response using setpoint changes, *Applied Energy* 177 (2016) 149–164.
- [7] ASHRAE, District Cooling Guide (2013).
- [8] J. A. de Chalendar, P. W. Glynn, S. M. Benson, City-scale decarbonization experiments with integrated energy systems, *Energy & Environmental Science* 12 (5) (2019) 1695–1707.
- [9] G. Campos, Y. Liu, D. Schmidt, J. Yonkoski, D. Colvin, D. M. Trombly, N. H. El-Farra, A. Palazoglu, Optimal real-time dispatching of chillers and thermal storage tank in a university campus central plant, *Applied Energy* 300 (2021) 117389.
- [10] D. B. Crawley, L. K. Lawrie, C. O. Pedersen, F. C. Winkelmann, *EnergyPlus: Energy simulation program*, *ASHRAE journal* 42 (4) (2000) 49–56.
- [11] N. A. Campbell, P. E. Phelan, M. Peinado-Guerrero, J. R. Villalobos, Improved air-conditioning demand response of connected communities over individually optimized buildings, *Energies* 14 (18) (2021) 5926.
- [12] J. H. Yoon, R. Baldick, A. Novoselac, Dynamic demand response controller based on real-time retail price for residential buildings, *IEEE Transactions on Smart Grid* 5 (1) (2014) 121–129.

- [13] Y. Li, J. Zuo, T. Qian, X. Guo, Demand response potential estimation for commercial buildings, in: 2018 China International Conference on Electricity Distribution (CICED), IEEE, 2018, pp. 2999–3003.
- [14] J. Liu, R. Yin, M. Pritoni, M. A. Piette, M. Neukomm, Developing and evaluating metrics for demand flexibility in buildings: Comparing simulations and field data, Tech. rep., Lawrence Berkeley National Laboratory (2021).
- [15] M. Cai, S. Ramdasalli, M. Pipattanasomporn, S. Rahman, A. Malekpour, S. R. Kothandaraman, Impact of HVAC set point adjustment on energy savings and peak load reductions in buildings, in: 2018 IEEE International Smart Cities Conference (ISC2), IEEE, 2018, pp. 1–6.
- [16] K. Aduda, T. Labeodan, W. Zeiler, G. Boxem, Y. Zhao, Demand side flexibility: Potentials and building performance implications, *Sustainable cities and society* 22 (2016) 146–163.
- [17] P. Xu, P. Haves, M. A. Piette, L. Zagreus, Demand shifting with thermal mass in large commercial buildings: Field tests, simulation and audits, Tech. rep., Lawrence Berkeley National Laboratory (2005).
- [18] R. Yin, P. Xu, M. A. Piette, S. Kiliccote, Study on Auto-DR and pre-cooling of commercial buildings with thermal mass in California, *Energy and Buildings* 42 (7) (2010) 967–975.
- [19] A. Khalilnejad, R. H. French, A. R. Abramson, Evaluation of cooling setpoint setback savings in commercial buildings using electricity and exterior temperature time series data, *Energy* 233 (2021) 121117.
- [20] Y. Chen, Z. Chen, P. Xu, W. Li, H. Sha, Z. Yang, G. Li, C. Hu, Quantification of electricity flexibility in demand response: Office building case study, *Energy* 188 (2019) 116054. doi:<https://doi.org/10.1016/j.energy.2019.116054>.
- [21] A. Ghahramani, K. Zhang, K. Dutta, Z. Yang, B. Becerik-Gerber, Energy savings from temperature setpoints and dead-band: Quantifying the influence of building and system properties on savings, *Applied Energy* 165 (2016) 930–942.
- [22] H. C. Gils, Assessment of the theoretical demand response potential in Europe, *Energy* 67 (2014) 1–18.
- [23] T. Hoyt, E. Arens, H. Zhang, Extending air temperature setpoints: Simulated energy savings and design considerations for new and retrofit buildings, *Building and Environment* 88 (2015) 89–96.
- [24] S. Stinner, K. Huchtemann, D. Müller, Quantifying the operational flexibility of building energy systems with thermal energy storages, *Applied Energy* 181 (2016) 140–154.
- [25] R. E. Hedegaard, M. H. Kristensen, T. H. Pedersen, A. Brun, S. Petersen, Bottom-up modelling methodology for urban-scale analysis of residential space heating demand response, *Applied Energy* 242 (2019) 181–204.
- [26] G. Reynders, J. Diriken, D. Saelens, Generic characterization method for energy flexibility: Applied to structural thermal storage in residential buildings, *Applied energy* 198 (2017) 192–202.
- [27] D. S. Callaway, I. A. Hiskens, Achieving controllability of electric loads, *Proceedings of the IEEE* 99 (1) (2010) 184–199.
- [28] G. Strbac, Demand side management: Benefits and challenges, *Energy Policy* 36 (12) (2008) 4419–4426.
- [29] D. S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, *Energy Conversion and Management* 50 (5) (2009) 1389–1400.
- [30] S. Goy, D. Finn, Estimating demand response potential in building clusters, *Energy Procedia* 78 (2015) 3391–3396, 6th International Building Physics Conference, IBPC 2015. doi:<https://doi.org/10.1016/j.egypro.2015.11.756>.
- [31] P. J. Jones, S. Lannon, J. Williams, Modelling building energy use at urban scale, in: *Proceedings of the Seventh International IBPSA Conference, Rio de Janeiro, Brazil, Citeseer*, 2001, pp. 175–180.
- [32] X. Gao, A. Malkawi, A new methodology for building energy performance benchmarking: An approach based on intelligent clustering algorithm, *Energy and Buildings* 84 (2014) 607–616.
- [33] S. Zhan, Z. Liu, A. Chong, D. Yan, Building categorization revisited: A clustering-based approach to using smart meter data for building energy benchmarking, *Applied energy* 269 (2020) 114920.
- [34] D. Patteuw, G. P. Henze, A. Arteconi, C. D. Corbin, L. Helsen, Clustering a building stock towards representative buildings in the context of air-conditioning electricity demand flexibility, *Journal of Building Performance Simulation* 12 (1) (2019) 56–67.

- [35] J. A. de Chalendar, C. McMahon, L. F. Valenzuela, P. W. Glynn, S. M. Benson, Unlocking demand response in commercial buildings: Empirical response of commercial buildings to daily cooling set point adjustments, *Energy and Buildings* 278 (112599) (2023).
- [36] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, New York, 2001.
- [37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* 12 (2011) 2825–2830.
- [38] BloombergNEF, Battery pack prices fall to an average of \$132/kwh, but rising commodity prices start to bite, <https://about.bnef.com/blog/battery-pack-prices-fall-to-an-average-of-132-kwh-but-rising-commodity-prices-start-to-bite/>, Last accessed on 2022-09-23 (2021).