Counterfactual Analysis: A Case Study on Impact of External Events on Building Energy Consumption

Carolina Minami Oguchi[1], Dipak Ghosal[1,2], Alex Sim[2], and Kesheng Wu[2]

[1]University of California, Davis, CA, U.S.A.
[2]Lawrence Berkeley National Laboratory, Berkeley, CA, U.S.A.

***Abstract*—Energy consumption in buildings accounts for a significant portion of the global energy use. Consequently, understanding building energy use is important. Data over the past decade show that the energy intensity (Joules/sqft) of commercial buildings has decreased. While some of the improvements (decrease in energy use) are easily measurable such as the use of more energy efficient lighting, impact of other modifications such as changes to the operation of the HVAC system or changes in the usage pattern of the building potentially due to external events are difficult to quantify. Simply comparing energy consumption prior and post change is not accurate as energy use is impacted by many factors including external weather conditions. In this paper, we present a case study to quantify the impact of external events on the energy consumption of a medium-sized office building. We adopt an approach based on counterfactual analysis. Towards this end, we first build two models based on Linear Regression and k-Nearest Neighbors to predict the daily energy use given different input features related to the weather. We determine the statistical features of the weather that are most predictive of energy use. We then use the models to determine a counterfactual baseline and thereby to accurately estimate the impact of the events. The results of the counterfactual analysis provide new insights on the impact of the events on energy consumption. The update to the building cooling system resulted in more energy savings than direct yearly comparison reveals. On the other hand, the tests of a MPC-based controller for the HVAC system saved less energy than determined by the direct yearly comparison. Finally, the results show that there no gains in terms of energy savings due to remote work during the COVID-19 pandemic. An increase in airflow setting in the HVAC system corroborates this finding and further validates the underlying model and the counterfactual analyses.**

***Index Terms*—Building Energy Usage, Counterfactual Analysis, Linear Regression, k-Nearest Neighbor, Weather Conditions**

I. Introduction

The data released by the Energy Information Administration (EIA) in 2015 showed that energy consumed by residential and commercial buildings accounts for 40% of total energy consumption in the United States [1]. The data is similar for the European Union (EU) whereas globally it is 32% and increasing [2]. Consequently, with the perspective of reducing environmental and financial costs, there are significant efforts to study building energy consumption and identify methods to optimize energy usage.

Commercial buildings can be broadly categorized into four types [3] namely hotels, offices, retail, and mixed developments. They all have different characteristics to fulfill the purposes of the facilities. They can be further categorized by the size of the building. The energy usage in a building is determined by many factors, such as ambient weather conditions, building structure and characteristics, the operation of the HVAC systems, lighting, occupancy, and behavior of the occupants. As such optimizing the energy usage of a building is a complex task that must satisfy various constraints. For example, the Occupational Safety and Health Administration (OSHA) Regulations require that building operations maintain the temperature and other factors in the office building within a certain range to ensure a workable environment for the employees [4]. Setting temperature set-points for the HVAC system to minimize energy consumption while meeting the constraints is complex due to inherent physical characteristics of the building including the latent heat [5].

From a building operation perspective, it is important to estimate the impact on building energy usage due to building updates or external events that impact the building occupancy and/or usage pattern. An accurate estimate can guide building operations and/or subsequent updates. However, accurately estimating the impact is challenging. Simply looking at the difference in energy usage prior and post event is not accurate as the energy usage is impacted by various other factors including the ambient weather conditions. In this paper, we present a case study of Building 59 (aka Wang Hall) in Lawrence Berkeley National Laboratory to accurately estimate the energy use impacts due to updates on the building's HVAC system and due to external events including COVID-19 pandemic, wildfires, and others, which all took place between 2018 and 2020.

Even as Building 59 is a relatively new building (operational in 2015), there is curated data on building energy usage only since early 2018 [6], [7]. As such there is no significant historical data to estimate the distribution of the energy use prior to the events. Furthermore, the consecutive events of the HVAC update and changes to occupancy due to remote work due to COVID-19 make it hard to determine the individual impacts. We adopt a counterfactual analysis approach. Specifically, we first develop a model to predict the energy use of the building given the ambient conditions. For a given event, we then use the model to counterfactually estimate the energy usage if the event had not occurred. Comparing this estimate with the actual usage due to the event gives an accurate estimate of the

979-8-3503-2445-7/23/$31.00 ©2023 IEEE

benefit (or loss) due to the event.

A model used for this counterfactual study is sometimes called a baseline [8]. Due to the large number of variables that affect building energy uses, there are considerable challenges in building an effective baseline model [8], [9]. For example, we might need to compare hourly energy use during the same season or month from different years and most time-series modeling approaches are not able to provide such fine-granularity prediction across an extended period of time. Furthermore, the daily cycle of temperature changes and delayed response from HVAC system (due to heat capacity of the building) are challenging for effective modeling. Fortunately, in this study, we could use aggregated electricity usage. By choosing to model the total daily electricity usage, we are able to circumvent most of these challenges [9]. This allows us to treat weather information and electricity usage from each day as an independent data record. A number of published studies suggested the total electricity usage of a building to be a linear function of the average outdoor temperature [9]. In this work, our dataset contains a number of additional measures of ambient weather conditions. Our main objective in model creation is to develop ways to make use of a subset set of non-co-linear weather features and show that these features could be used to improve modeling accuracy.

The key contributions of this paper are the following:

1) We develop two models based on Linear Regression and k-Nearest Neighbors to predict the daily building energy use given the ambient weather conditions. We determine the subset of ambient weather features that are most predictive of energy usage.
2) We design a counterfactual analysis to determine the true gain (or loss) in energy use due to HVAC System Update, Remote Work during the first year of the COVID-19 pandemic, Wildfires, and Model Predictive Control (MPC) Testing.
3) The results show that the HVAC System Update resulted in more gains in energy savings than the direct yearly comparisons reveal. The result also shows that the benefits in terms of energy savings due to Remote Work during COVID-19 were not significant. Increased airflow settings in the HVAC system corroborate this finding. Finally, the energy savings due to MPC Testings were less than the direct comparison revealed.

This paper is organized as follows. Motivation, research goals, and approach are described in Section II. The data we use for this paper is explained in Section III. Section IV describes methods and experimental settings. Results are discussed in Section V. Section VI outlines related works. The paper is concluded in Section VII.

II. Motivation

The data released by EIA [1] in 2018 (based on the Commercial Buildings Energy Consumption Survey (CBECS)) show that even though the total floor space in commercial buildings has increased, the energy consumption has not. The comparison is based on the 2012 CBECS data. This implies that the energy consumption per square foot (energy intensity) has decreased which is attributed to increased building energy efficiency. This is due to improvements in building operations, materials, and design, as well as heating, cooling, and lighting technologies. For example, the use of highly efficient LED lighting in commercial buildings has grown from 9% in 2012 to 44% in 2018. While the impact of energy-efficient LED lighting on the overall building energy consumption is easily quantified, that is often not the case for other updates to the building or changes in the energy use due to external events.

The target of this study is Building 59 which is an office building at Lawrence Berkeley National Laboratory located in Berkeley, California, U.S.A. [7]. Building 59 is a medium-sized four-floor building with mechanical systems and super-computer facilities on the first and second floors, and office floors on the third and fourth floors. The dataset is of the office portion of the building [7]. The office has the North wing and the South wing. Underflow Air Distribution and HVAC system are used for both heating and cooling in the building, and automated Logic WebCTRL building management system is used for the control. The climate around the area is mild, where the temperature varies between around $5°C$ to $25°C$. September is the hottest month, and January is the coldest month on average.

The daily energy use depends on many factors including ambient weather conditions, HVAC operational control, and occupancy. Changes in building energy consumption occur when changes are made to the operation and control of the HVAC system for example, changing to an MPC-based controller to control the HVAC system. Building energy consumption is also impacted by the usage pattern of the building or external events such as wildfires that may impact the ambient weather and/or the operation of the building. The goal of this study is to accurately quantify the impact of these events on the building energy use.

A. *Events of Interest*

Over the period for which the data was collected, there were four major events [7] which are described below.

1) **HVAC System Update:** The update of Building 59 took place in 2019 winter, with the aim to reduce energy usage. The update includes, but is not limited to, installing the screen for the elevator shaft. Two other changes that were aimed primarily at the National Energy Research Scientific Computing Center (NERSC) supercomputer facility and had a secondary impact including a) a new Heat pump water heater that harvested heat from the datacenter's Cooling Water loop and b) A Cooling Water Booster Pump which provides the extra needed loop pressure to get acceptable Cooling Water flow up to the 4 roof-top units (RTUs).
2) **Remote Work:** Like most workplaces, a majority of office workers started remote work in late February and early March of 2020 as a result of the COVID-19 pandemic. As the building hosts NERSC which is an important user facility and also some of the key engineers

managing the ESnet, a small number of people took turns to be in the office. In order to mitigate any exposure to COVID-19, precautions were taken including minimizing the number of occupants in the building at any time and increasing the airflow.

3) **Wildfires:** The 2018 wildfire season was the deadliest and most destructive wildfire season on record in California. In November 2018, strong winds aggravated conditions in another round of large, destructive fires across the state. This new batch of wildfires included the Woolsey Fire and the Camp Fire. These fires significantly impacted the external environment with many days of thick haze and the building occupancy as many worked remotely.
4) **MPC Controller Testing:** In late 2020 over a period of approximately 2 months a new MPC-based controller for the HVAC system was tested [10].

B. Research Goal and Approach

The goal of this research is to accurately estimate the impacts on the building's energy consumption due to the above four events. The naive estimate by taking the difference between the energy usage pre- and post/during- event is not accurate. This is because the energy usage is dependent on other parameters including ambient weather conditions and the occupancy which vary pre- and post- event. In this paper, we adopt counterfactual analysis (aka hindcast) to estimate the energy impact of the events. Given the data prior to the event, we build a model that can predict the building energy use given the ambient weather conditions. We then use the model to estimate (counterfactually) what the energy use would have been had the event not occurred. This gives a more accurate baseline to estimate the impact of the event.

III. Building 59 Data

Building 59 is an all-electrical building. The data is based on 300 sensors [7]. The dataset includes energy use data, ambient weather data, indoor environmental data, HVAC operational data, and occupancy data. Sampling frequencies of data range from 1 minute to 15 minutes. There are some missing portions of data due to different reasons including measurement errors. Missing data is filled using one of the three methods: simple linear interpolation for smaller gaps, k-nearest neighbors algorithm for small gaps, and matrix factorization for large gaps. It is collected, processed, and cleaned by Na et al. [7].

Figure 1 shows the yearly EUI (Energy Use Intensity) for the three years for which the data was collected: 2018, 2019, and 2020 [7]. The components namely HVAC, lighting, and miscellaneous (MELs) are also shown. MELs include, but not limited to, plug loads and elevators. As the data shows there are significant yearly differences both at the aggregate level as well as at the component level. Specifically, while the HVAC system is the largest consumer, the biggest yearly difference was in MELs.

The daily energy use over a month (July 2018) is shown in Figure 2. Again the data shows that the HVAC system is the main consumer of energy accounting for 73.84% of the total energy use. The lighting and MELs account for 4.43% and 21.72% of the total, respectively. Figure 3 shows the maximum, minimum, and mean of the daily outdoor temperature and outdoor relative humidity during July 2018.

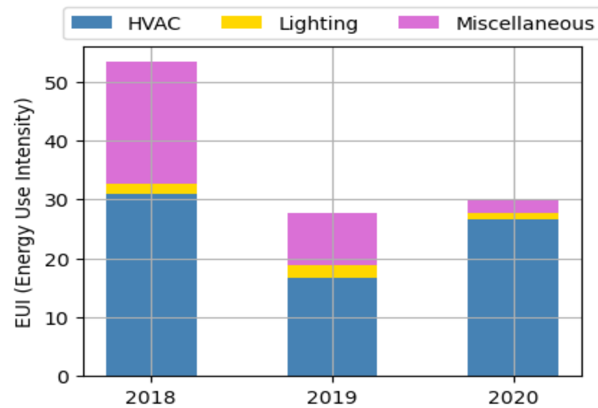

Fig. 1: The EUI (Energy Use Intensity) of the Office Portion of Building 59 in 2018, 2019, and 2020.

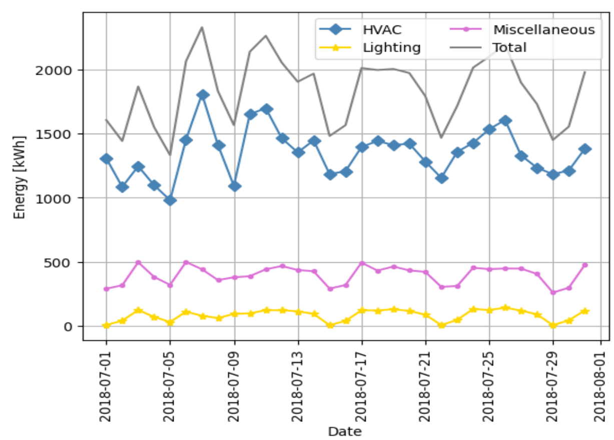

Fig. 2: The daily energy use for July 2018 and the key components. HVAC is the main consumer.

IV. Methods and Experiments

We conduct a counterfactual analysis to accurately determine the impact of different events on the energy use. The overview of this method is shown in Figure 4. An important point to note here is that ambient weather conditions change and we want to make a comparison of energy use before and after the event with that change in consideration. To perform this, we train a machine learning (ML) model to predict the energy use based on ambient weather conditions. We then estimate the counterfactual energy use with the post-event ambient weather conditions as input. This give us the energy that would have been used had it not been for the event and provide an accurate baseline to estimate the impact of the event.

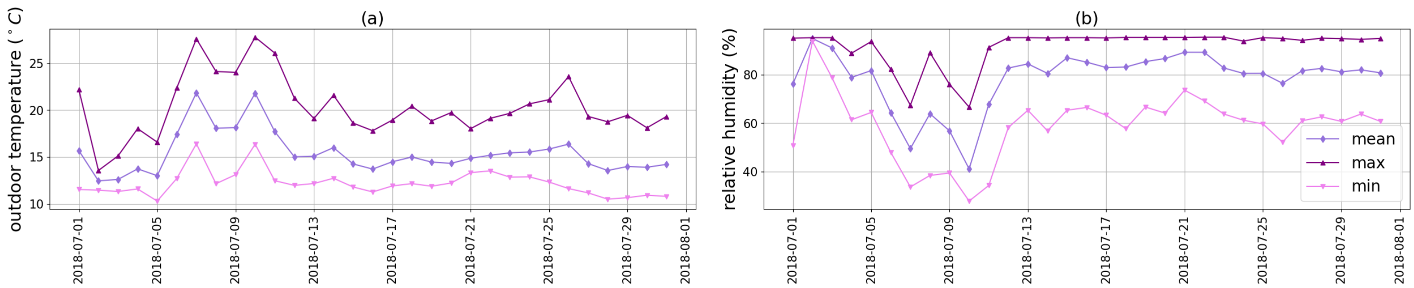

Fig. 3: Outdoor weather conditions for July of 2018; a) maximum, minimum, and mean of daily outdoor temperature, b) maximum, minimum, and mean of daily outdoor relative humidity

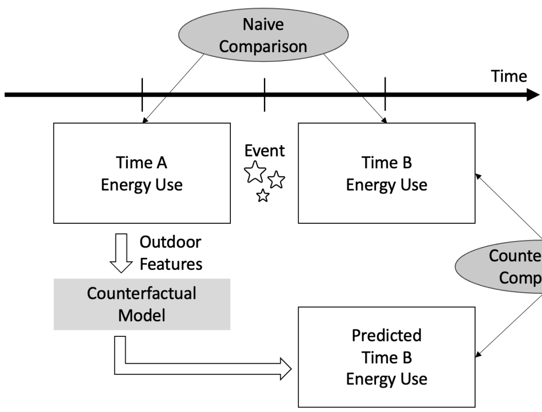

Fig. 4: The illustration of our method. We use the data from time period A to train the model. We use the model and the weather conditions of time period B to make a counterfactual prediction of energy use for time period B. The actual energy used in time period B is compared with the predicted energy use to estimate the change in the energy use. The naive comparison is comparing the energy use in time periods A and B.

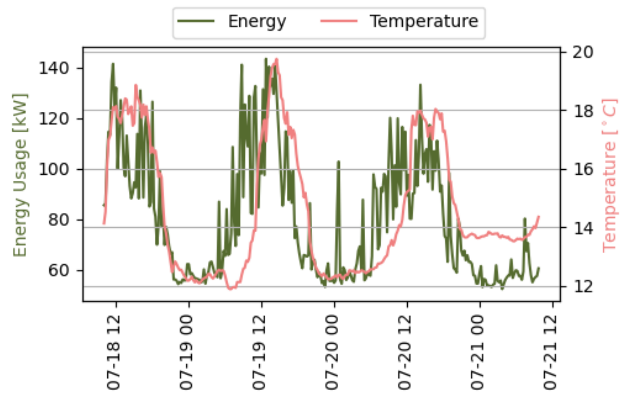

Fig. 5: The energy use and the outdoor temperature during a week in July 2018. The time lag in the external temperature and energy use particularly visible in final diurnal cycle is due to latent heat capacity of the building.

A. Calculating Energy Use

To calculate the energy use we use the sampled values of instantaneous power (in kW) drawn by the various subsystems, sampled every 15 minutes. The data is available separately for the North and South wings of the building. As mentioned before, the various subsystems are categorized into three broad groups - HVAC System, Lighting, and Miscellaneous Electrical (MELs). For any subsystem, the energy use per hour (in kWh) is given by

$$E[kWh] = P[kW] \times T[h] \tag{1}$$

where $E[kWh]$ is energy, $P[kW]$ is power, and $T[h]$ is time in hours. In our case, T is $\frac{1}{4}$ as it is measured every 15 minutes. The total energy use per day E is then given by

$$\begin{aligned} E_{total} = \sum_{i=1}^{96} (E_{HVAC}^{North}(i) + E_{HVAC}^{South}(i) \\ + 2 \times E_{Lighting}^{South}(i) \\ + E_{Miscellaneous}^{North}(i) + E_{Miscellaneous}^{South}(i)) \end{aligned} \tag{2}$$

where i is the sample index with 96 samples per day since the sampling frequency is once every 15 minutes. The lighting data is only provided for the South wing. As suggested in [7], it is reasonable to assume that the energy use due to lighting is very similar for both wings. So, the energy use of the South wing lighting is multiplied by a factor of 2 in Eq. 2.

B. Latent Heat

Buildings are heated or cooled in response to the indoor temperature change. While the indoor temperature is affected by the outdoor temperature, there is a time lag as a result of the latent heat capacity of the building. Consequently, there is a time lag in the response of the HVAC system with respect to the changes in outdoor temperature. As the HVAC system is one of the main consumer of energy, there is a time lag between the energy use and the highs and lows of the outdoor temperature. This effect of latent heat is shown in Fig. 5. While the above discussion is only based on the temperature, in reality how the latent heat is dissipated is also impacted by other ambient weather conditions including relative humidity and solar radiation. In this paper, we consider the energy use at the granularity of a day. As a result, the effect of latent heat does not need to be one of the features in our analysis.

C. Models

We considered two models - one based on Linear Regression (LR) and the other based on k-Nearest Neighbors (kNN) [11]. In kNN, k is the number of most closest points in the data over which the average is taken, and it is a hyper-parameter. In this paper we use Euclidean distance to measure the closeness. We selected LR and kNN for this study because they are capable of learning with a small amount of data. Since we use daily energy use, we are limited to 30 or 31 data points. More complex models such as Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM), and their variants require a lot more data for training [11].

The explanatory/independent variable of the model is a set of statistical features of daily ambient weather conditions and the response/dependent variable is the daily energy use. We picked ambient weather conditions as the explanatory/independent variables because they are the only features that are independent of the HVAC operational control. The data provides HVAC operational data [7], but these are not independent and are impacted by the closed loop control algorithm of the HVAC system. We consider the outdoor temperature, dew point temperature, relative humidity, and solar radiation for the ambient weather conditions. The 96 data points for each day are normalized so that each feature will have equal weights in the model. We consider the daily minimum, maximum, and average of each feature.

The data for the training set and evaluation set are chosen depending on the type of event. The first category includes events that change the behavior of energy consumption permanently or more precisely for the entire duration for which the data is available. HVAC System Update and Remote Work belong to this category. For these events, to analyze the stable part of the data (i.e., after some time past the event), we picked the summer months, as these events happened in winter and spring. Specifically, since the HVAC System Update took place in late 2018, the training set is 31 data points from July 2018 and the evaluation set is 31 data points from July 2019. Similarly for Remote Work due to COVID-19 which started in early 2020, the training set is 31 data points from July 2019 and the evaluation set is 31 data points from July 2020.

The second category include events that change the energy use of the building temporarily only during the time period during which the events occurred. Events in the category are the Wildfires and MPC Testings. For these events, the evaluation set consists of n data points, where n is the number of days over which the event occurred. The training set is 30 data points corresponding to immediate 30 days prior to when the event started. The days used as the training set and the evaluation set for each event are summarized in Table I.

As shown in Figure 4, we also do a naive comparison for each event. For the HVAC System Update, we compare the energy use of 31 days in July 2018 to the energy use of 31 days in July 2019. For the Remote Work, we compare the energy use of 31 days in July 2019 to the energy use of 31 days in July 2020. For the other events, we compare the energy use of n days during the event, and n days immediately prior to when the event started.

Event	Training Set	Evaluation Set
HVAC System Update	2018/07/01 - 2018/07/31	2019/07/01 - 2019/07/31
Remote Work	2019/07/01 - 2019/07/31	2020/07/01 - 2020/07/31
First Wildfire	2018/10/12 - 2018/11/11	2018/11/12 - 2018/11/20
Second Wildfire	2020/07/24 - 2020/08/23	2020/08/24 - 2020/09/06
First MPC Testing	2020/09/19 - 2020/10/19	2020/10/20 - 2020/10/27
Second MPC Testing	2020/10/02 - 2020/11/01	2020/11/02 - 2020/11/06
Third MPC Testing	2020/10/13 - 2020/11/12	2020/11/13 - 2020/11/19
Fourth MPC Testing	2020/11/03 - 2020/12/03	2020/12/04 - 2020/12/14

TABLE I: Details of the data that was used for the test set and the evaluation set for each event.

D. Feature Selection, Parameter Selection, and Method Validation

We select features and a parameter and validate the method using the data from June and July in 2019 and 2020. We chose to use June and July because the weather is similar in those two months. We did not consider August because there was a wildfire in August 2020. We also did not pick May and June because the weather is less similar than June and July.

The ambient weather features available in the dataset are air temperature($°C$), air dew point temperature($°C$), air relative humidity($\%$), and solar radiation(W/m^2). For all those four features, we calculated the minimum, maximum, and mean values per day. For LR, we searched for one best features among the 12 features described above taking into account the co-linearity among the features. We used the Variance Inflation Factor (VIF) to determine the co-linearity among the features. The result showed high co-linearity among all the features except for minimum solar radiation. This is unsurprising because the minimum solar radiation is zero (0) at night. As such the solar radiation is not a good feature for the model. For kNN, we searched the best feature combinations among the following four: 1) minimum, maximum, and mean of all outdoor environment features, 2) minimum, maximum, and mean of temperature, 3) mean of all features, and 4) means of temperature and humidity. The reason for choosing the temperature and humidity as a combination is that they had the highest correlation to energy consumption. We varied the hyper-parameter k from 2 to 7. Therefore, there are 12 features to choose from for LR, and there are 24 combinations of features (4 combinations of features and 6 options for k) to choose from for kNN. Among all those combinations, we pick the best one each for LR and kNN and use it for all the analyses.

The steps for feature and hyper-parameter selection are as follows. First, we obtain the daily outdoor environmental data and the daily energy use data of June and July 2019, which is 61 data points in total (30 days in June and 31 days in July). We shuffle the data. Then we use the first half of the shuffled data as a training set and the remaining half as a validation set. We train the model using the first set and then use the model to make a prediction on the second set. We compute

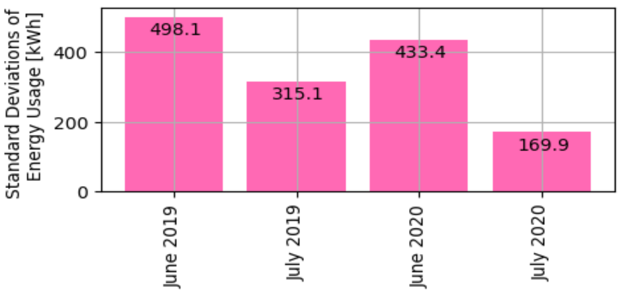

Fig. 6: The standard deviations of the daily energy use for the months used in the training and evaluation sets.

the Root Mean Square Error (RMSE) between the predicted and the actual values on the second set of data points. We do the same for the data of June and July of 2020, and obtain the RMSE for each combination. Then, we find a model that has the smallest average RMSE considering both 2019 and 2020 for each model and employ that combination for our analyses.

For LR, the results of feature selection are listed in Table II. The feature selected for analysis is the maximum outdoor temperature, as it has the smallest average RMSE among all the features considered. For kNN, the results of feature combination selection and hyper-parameter selection are listed in Table III. We observe that the smallest average RMSE is achieved with with $k = 3$ and features consisting of the mean temperature and mean relative humidity. So, we use this combination for all the analyses in the next section.

We performed validation of the method by comparing two similar sets of data which are created from the two months of the same year, which is the same as the sets described above. We consider the method to be effective when the difference between the prediction and the actual energy use is smaller than the standard deviation of the data itself. The standard deviations of daily energy use of June and July in 2019 and 2020 are shown in Figure 6. We observe that the range of RMSEs in Table II and Table III is similar or are smaller than the standard deviations of the daily energy use in Figure 6. We thus conclude that the method is validated and can be used for our analysis.

V. Results and Discussions

In this section, we present the results of our analysis of the events outlined in Section II. As noted in Section IV we will use the following metrics:

- **Average Naive Difference**: We take the daily difference between energy use prior to the event with the energy use post event. We then take the average over the number of days. A positive (negative) value indicates energy saving (loss).
- **Average Counterfactual Difference**: We take the daily difference between energy use predicted by the counterfactual model with the energy use post event. We then take the average over the number of days. Again, a positive (negative) value indicates energy saving (loss).

The results are shown in Table IV. We discuss these results with the actual and predicted energy usage in the following paragraphs.

a) HVAC System Update: Figures 7 and 8, show the actual energy use and counterfactually predicted energy use using LR and kNN, respectively for the case of the HVAC System Update. The results show that the counterfactual prediction is always larger than the actual, implying that there were savings in energy use due to the HVAC System Update. The Average Counterfactual Difference for the HVAC System

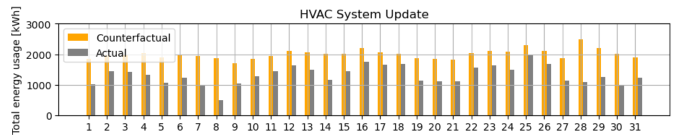

Fig. 7: Comparison Actual energy use and Counterfactual predicted energy use by LR after HVAC System Update.

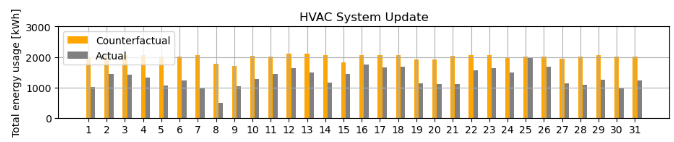

Fig. 8: Comparison Actual energy use and Counterfactual predicted energy use by kNN after HVAC System Update.

Update in Table IV shows the actual energy saving is more than that determined by the Average Naive Difference. It is also noteworthy that both the models, LR and kNN, show similar gains.

b) Remote Work: Figures 9 and 10 show the actual energy use and counterfactually predicted energy use using LR and kNN, respectively for the case of Remote Work. The results show that the predicted energy use is mostly lower than the actual energy use. This implies that there was a loss, i.e., more energy was consumed. The results of the Average Counterfactual Difference for Remote Work in Table IV show that the actual energy loss was significantly more than the loss determined by the Average Naive Difference. The result that the energy consumption was higher during Remote Work due to COVID-19 is counter-intuitive considering that the occupancy of the building was reduced to a very few number of employees. The reason is that the operation of the HVAC system such as air flow rate and fan speed was higher during the COVID-19 Remote Work period than it was before. This is also observed in Fig.11. This was likely done to minimize COVID-19 infections.

c) Wildfires: Figures 12 and 13, show the actual energy use and counterfactually predicted energy use using LR and kNN, respectively for the first Wildfire. The trend for both LR and kNN is the same. The number of days that the counterfactual prediction was higher and the number of days

Feature	June-July 2019 RMSE	June-July 2020 RMSE	Average of 2019 and 2020
temperature min	319.3	301.0	310.1
temperature mean	253.2	205.1	229.1
temperature max	236.8	191.8	**214.3**
dew point temperature min	416.7	316.3	366.5
dew point temperature mean	441.7	343.1	392.4
dew point temperature max	458.2	367.0	412.6
relative humidity min	302.5	220.5	261.5
relative humidity mean	316.9	239.5	278.2
relative humidity max	357.9	277.7	317.8
solar radiation min	433.8	331.0	382.4
solar radiation mean	384.7	259.8	322.3
solar radiation max	419.8	329.6	374.7

TABLE II: The results of feature selection for LR. The feature maximum temperature has the lowest RMSE.

Experimental setting	k	June-July 2019 RMSE	June-July 2020 RMSE	Average of 2019 and 2020
min max mean; temperature	k=2	309.5	227.1	268.3
	k=3	288.8	218.9	253.8
	k=4	298.6	223.1	260.9
	k=5	293.7	212.4	253.0
	k=6	297.9	208.9	253.4
	k=7	303.4	216.9	260.2
min max mean; all features	k=2	296.0	269.7	282.9
	k=3	288.1	255.8	272.0
	k=4	279.6	244.4	262.0
	k=5	288.6	239.9	264.3
	k=6	311.4	253.8	282.6
	k=7	321.6	259.3	290.4
mean; all features	k=2	316.9	219.0	267.9
	k=3	299.7	217.4	258.6
	k=4	300.2	224.3	262.2
	k=5	298.1	235.9	267.0
	k=6	315.6	241.9	278.7
	k=7	319.3	250.1	284.7
mean; humidity, temperature	k=2	274.6	249.1	261.9
	k=3	291.5	211.7	**251.6**
	k=4	307.0	231.8	269.4
	k=5	308.0	234.5	271.3
	k=6	316.8	239.2	278.0
	k=7	315.2	253.1	284.1

TABLE III: The results of feature and parameter selection for kNN. The parameter $k = 3$ and feature mean of humidity and temperature has the lowest average RMSE.

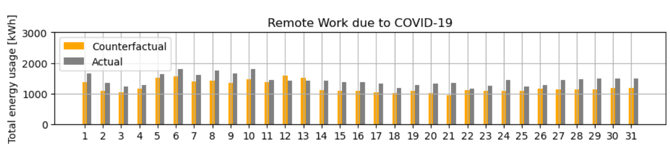

Fig. 9: Comparison Actual energy use and Counterfactual predicted energy usage by LR during Remote Work.

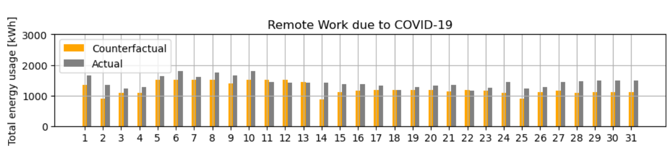

Fig. 10: Comparison Actual energy usage and Counterfactual predicted energy usage by kNN during Remote Work.

that the actual energy use was higher were about the same in the first wildfire. For the second wildfire there were more days when the actual energy use was higher. Comparing the Average Counterfactual Difference and the Average Naive Difference in Table IV, we find that the results are inconclusive. The first wildfire has a positive average difference and both LR and kNN give less savings than the Average Naive Difference. The second wildfire shows markedly different results - gains in the Average Naive Difference but loss with the Average Counterfactual Difference for both the models. Also, the magnitude of the difference is much smaller compared to the other events. Overall, no conclusions can be drawn from the results.

d) MPC Testings: Figures 14 and 15 show the actual energy use and counterfactually predicted energy use using LR and kNN, respectively for the first MPC Testings. We also see that the results are consistent for the two models in all MPC Testing. We observe that the counterfactual prediction is mostly higher than the actual energy use implying energy gains

Event	Average Counterfactual Difference		Average Naive Difference
	LR	kNN	Naive Comparison
HVAC System Update	684.3	682.0	513.7
Remote Work	-222.0	-211.7	-112.7
First Wildfire	64.8	70.6	160.2
Second Wildfire	-242.8	-174.5	333.7
First MPC Testing	612.8	796.9	1086.9
Second MPC Testing	371.9	388.9	438.5
Third MPC Testing	285.1	243.0	465.3
Fourth MPC Testing	474.7	488.5	679.1

TABLE IV: The Average Counterfactual Difference (Predicted Counterfactual Energy Use − Actual Energy Use) and the Average Naive Difference (Actual Energy Use before event − Actual Energy Use after or during Event) for each event. The Building Update and the Remote Work were averaged over 30 days. For the other events, the average was taken over n days, where n is the number of days in July (31) for Building Update and Remote Work, and the number of days during the event for Wildfires and MPC Testings.

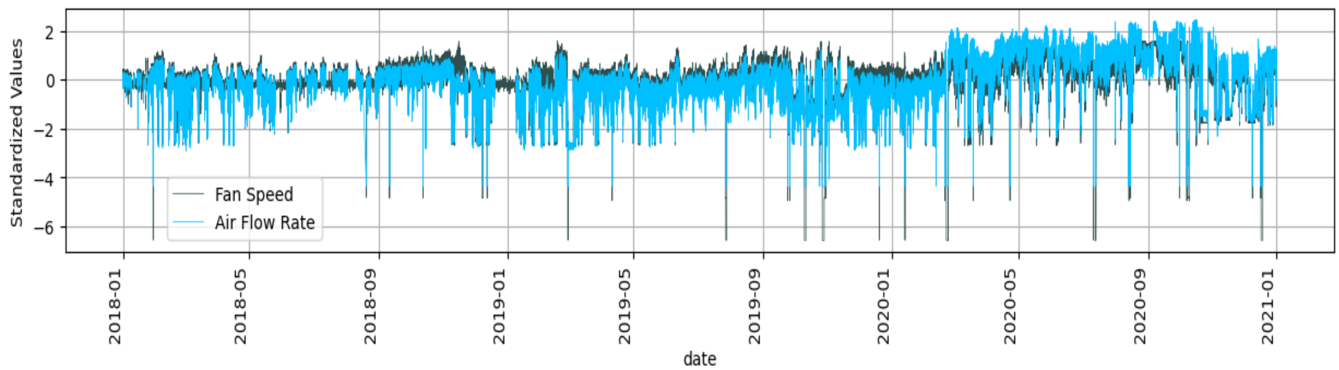

Fig. 11: The supply air flow rate and supply fan speed over the entire date range. The data shows that both the supply air flow rate and the supply fan speed were higher starting in Feb/March 2020 when COVID-19 started.

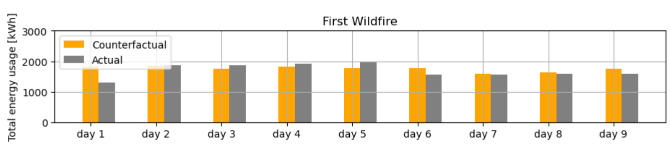

Fig. 12: Comparison of Actual energy use and Counterfactual predicted energy use during the first Wildfire by LR.

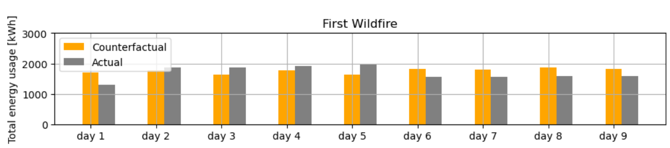

Fig. 13: Comparison of Actual energy use and Counterfactual predicted energy use during the first Wildfire by kNN.

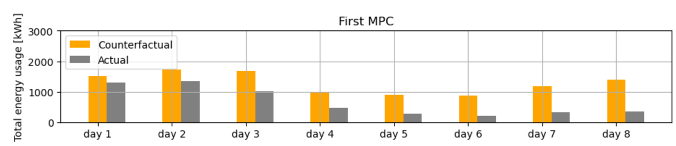

Fig. 14: Comparison of Actual energy use and Counterfactual predicted energy use during the first MPC Testing by LR.

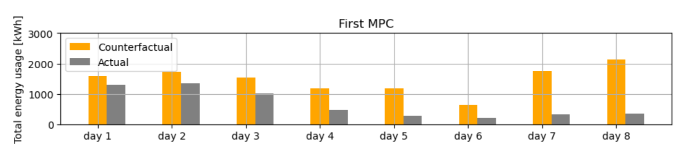

Fig. 15: Comparison of Actual energy use and Counterfactual predicted energy use during the first MPC Testing by kNN.

due to the use of the MPC controller. The Average Counterfactual Difference shows that there were savings in energy in all MPC Testings. In comparison to the Average Naive Difference, there were less energy savings in counterfactually predicted energy use for all the MPC Testings. The goal of the MPC controller was to optimize the efficiency of the HVAC system. The study in [10] showed that the MPC saved approximately 40% of HVAC energy over the existing control. Overall our results corroborate this prior study.

VI. Related Works

Analyses of energy use in different types of buildings have been widely studied. These include comparison and applicability of different machine learning (ML) models on building energy prediction. There are a number of papers that use LR or kNN to make predictions of energy use [12]–[19]. The study in [12] built a LR model for electricity load prediction for commercial and industrial buildings, including an office, a bakery, and a furniture store. They studied the time of the week and the outdoor temperature as explanatory features and concluded that the time of the week was a more accurate predictor. Prediction of daily energy use in an educational building is reported in [13]. The study concluded that using enthalpy and Cooling Degree Days as features improved model accuracy.

A comparative analysis of different ML models including kNN, Support Vector Machine (SVM), and other models applied to two educational buildings is reported in [18]. The study proposed a stacking model and compared it with other models. A similar study comparing ML models including kNN, ANN, and SVM to model energy use for a large biomedical facility is reported in [19]. They conclude that the error is reduced by more than 50% when ML models are used to model energy use. Prediction based on kNN has been shown to increase the accuracy on two commercial buildings in [14]. A noteworthy study in [15] makes hourly predictions of energy use in a community with different types of buildings including an office, an exhibition center, and an auditorium among others. The study in [16] applied a variant of kNN to a group of small commercial and residential buildings to make an accurate forecast of energy use and showed that it gives highly accurate predictions. A kNN based approach to optimize the operation of a chiller system to reduce carbon emissions in medium and large-size buildings is reported in [17]. The study found that the returned chilled water and the outdoor temperature are the best explanatory features for the model.

A number of recent works have studied the impact on building energy use due to COVID-19 [20]–[23]. An approach to measure the impacts of different stages of the COVID-19 lockdown policy is proposed in [20]. They analyzed 451 residential building and found that the energy consumption increased as the lockdown took place. Comparisons of energy use during pre- mid- and post- lockdown stages are reported in [21]. The five-floor commercial building included a basement floor that includes an exhibition, offices, and a performance center which normally welcomed many tourists every day. Unsurprisingly, the study found that the lockdown caused a significant decrease in energy use. A clustering based approach to evaluate energy use with open source data is reported in [22]. The study concluded that impact on building energy use during COVID-19 was different by types of buildings, such as educational buildings, research buildings, and residential buildings. A detailed analysis of the impact of COVID-19 on office buildings across the United States is explored in [23]. While the energy use increased due to COVID-19 because of the new measures, the energy use increased for zones that are above mid-humid climate and others decreased.

Models based on LR have been used to analyze the impact of COVID-19 on building energy use [3], [24], [25]. A study of 225 residential homes [24] shows that the energy use increase correlates well with the time when people used to be out, which is between 10 a.m. and 5 p.m. Also, they compared the differences between the income categories. Interestingly, they found that households with higher and the lowest incomes showed higher change in energy use. The study in [3] reports energy use of commercial buildings to measure the impact of COVID-19. City-level analysis with both hourly and daily predictions of different events such as COVID-19 lockdowns, drought, and blackouts is reported in [25]. They used an LR based counterfactual analysis to perform this study.

There are studies that estimate and/or analyze the effect of building retrofit (update) [26]–[29]. Han et al. [26] use a data-driven method to measure the effect of building retrofit on office buildings and laboratories. They use LR for predicting only outdoor environmental dependent features of energy and used other models to predict the energy use which decreased after the retrofit. Another study [27] estimates the energy saved by building retrofits using data-driven models including LR with simulated and actual data. An estimate of the energy saving by a retrofit to the cooling system of an office building with 25 floors is reported in [28]. An LR model using occupancy, number of working days, Cooling Degree Days, and their combinations as features is developed to make a prediction of energy use. They found out that the effectiveness of the feature was the largest with occupancy, then the number of working days, and lastly the Cooling Degree Days. kNN is also used in [29]. This study analyzed the difference in the energy use before and after a retrofit to a chiller system using the kNN algorithm. They found that the annual electricity is reduced by about 17%. The results of our case study are consistent with this prior work.

Overall, there are studies using one or more ML models to analyze the energy use of a building, a community, or a city by comparing before and after an event such as COVID-19 lockdowns, natural disasters, or building retrofits. However, to the best of our knowledge, none of these studies have used a counterfactual approach to accurately estimate the change in energy use at a granularity of a day from actual data observed from an office building.

VII. Conclusion

Accurately estimating the change in the energy use of a building due to updates and/or external events is essential to optimize energy use and reduce financial and environmental costs. Besides the updates and/or external events, other factors such as the weather affect energy consumption. In this study, we developed a method to perform a counterfactual analysis that considers the ambient weather conditions into account. We developed two models one based on Linear

Regression and the other based on kNN and identified the statistical features of the weather that were most predictive of the energy use. For a given event, we employed the models to make a counterfactual prediction of the energy use, i.e., the prediction of the energy use had the event not occurred. This provides an accurate baseline to estimate the change in energy use as a result of the external events. We considered the difference between the actual energy use and the counterfactually predicted energy use to obtain the true energy saving or loss. We considered four events, namely HVAC System Update, Remote Work due to COVID-19, Wildfires, and tests of MPC-based controller for the HVAC system. The results show that the HVAC System Update and MPC Tests saved energy, but Remote Work consumed more energy due to the increased airflow of the HVAC system. The impact of Wildfires was inconclusive. This study has the potential to optimize building energy use taking different measures into consideration. Our future work is to include the occupancy data in the model. It could not be done in this study because there were many missing values in the occupancy data and the available data was not consistent. As such the counterfactual analysis approach is applicable not only for building energy use but also for other scenarios that include data of measurement before and after a change.

Acknowledgement

This work was supported by the Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and also used resources of the National Energy Research Scientific Computing Center.

References

[1] Energy information administration (EIA)- commercial buildings energy consumption survey (CBECS). https://www.eia.gov/consumption/commercial/. Accessed: 2023-8-25.

[2] Singapore BCA. Singapore: leading the way for green buildings in the tropics, 2016.

[3] Senhong Cai and Zhonghua Gou. Impact of covid-19 on the energy consumption of commercial buildings: A case study in singapore. *Energy and Built Environment*, 2022.

[4] Occupational Safety and Health Administration. Osha worker rights and protections. https://https://www.osha.gov/workers. Accessed: 2023-6-20.

[5] Peder Bacher and Henrik Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511–1522, 2011.

[6] Shyh wang hall achieves LEED gold certification. https://www.nersc.gov/news-publications/nersc-news/nersc-center-news/2017/shyh-wang-hall-achieves-leed-gold-certification/, March 2017. Accessed: 2023-8-22.

[7] Na Luo, Zhe Wang, David Blum, Christopher Weyandt, Norman Bourassa, Mary Ann Piette, and Tianzhen Hong. A three-year dataset supporting research on building energy management and occupancy analytics. *Sci Data*, 9(1):156, April 2022.

[8] Irfan Qaisar and Qianchuan Zhao. Energy baseline prediction for buildings: A review. *Results in Control and Optimization*, 7:100129, 2022.

[9] Taehoon Kim, Jaesik Choi, Dongeun Lee, Alex Sim, C. Anna Spurlock, Annika Todd, and Kesheng Wu. Predicting baseline for analysis of electricity pricing. *International Journal of Big Data Intelligence*, 5(1/2):3–20, 2018.

[10] David Blum, Zhe Wang, Chris Weyandt, Donghun Kim, Michael Wetter, Tianzhen Hong, and Mary Ann Piette. Field demonstration and implementation analysis of model predictive control in an office hvac system. *Applied Energy*, 318:119104, 2022.

[11] Sebastian Raschka. *Python machine learning*. Packt publishing ltd, 2015.

[12] Johanna L Mathieu, Phillip N Price, Sila Kiliccote, and Mary Ann Piette. Quantifying changes in building electricity use, with application to demand response. *IEEE Trans. Smart Grid*, 2(3):507–518, September 2011.

[13] Minjae Shin and Sung Lok Do. Prediction of cooling energy use in buildings using an enthalpy-based cooling degree days method in a hot and humid climate. *Energy Build.*, 110:57–70, January 2016.

[14] Mel Keytingan M Shapi, Nor Azuana Ramli, and Lilik J Awalin. Energy consumption prediction by using machine learning for smart building: Case study in malaysia. *Developments in the Built Environment*, 5:100037, March 2021.

[15] Goopyo Hong, Gyeong-Seok Choi, Ji-Young Eum, Han Sol Lee, and Daeung Danny Kim. The hourly energy consumption prediction by KNN for buildings in community buildings. *Buildings*, 12(10):1636, October 2022.

[16] Meritxell Gómez-Omella, Iker Esnaola-Gonzalez, Susana Ferreiro, and Basilio Sierra. k-nearest patterns for electrical demand forecasting in residential and small commercial buildings. *Energy Build.*, 253:111396, December 2021.

[17] W T Ho and F W Yu. Chiller system optimization using k nearest neighbour regression. *J. Clean. Prod.*, 303:127050, June 2021.

[18] Ran Wang, Shilei Lu, and Wei Feng. A novel improved model for building energy consumption prediction based on model integration. *Appl. Energy*, 262:114561, March 2020.

[19] Colm V Gallagher, Ken Bruton, Kevin Leahy, and Dominic TJ O'Sullivan. The suitability of machine learning to minimise uncertainty in the measurement and verification of energy savings. *Energy and Buildings*, 158:647–655, 2018.

[20] V Todeschi, K Javanroodi, R Castello, and others. Impact of the COVID-19 pandemic on the energy performance of residential neighborhoods and their occupancy behavior. *Sustainable Cities and*, 2022.

[21] Yuan Su, Haoyuan Cheng, Zhe Wang, and Linwei Wang. Impacts of the COVID-19 lockdown on building energy consumption and indoor environment: A case study in dalian, china. *Energy Build.*, 263:112055, May 2022.

[22] Sebin Choi and Sungmin Yoon. Energy signature-based clustering using open data for urban building energy analysis toward carbon neutrality: A case study on electricity change under COVID-19. *Sustainable Cities and Society*, 92:104471, May 2023.

[23] Nuno D Cortiços and Carlos C Duarte. COVID-19: The impact in US high-rise office buildings energy efficiency. *Energy Build.*, 249:111180, October 2021.

[24] Emily Kawka and Kristen Cetin. Impacts of COVID-19 on residential building energy use and performance. *Build. Environ.*, 205:108200, November 2021.

[25] Gianlucca Zuin, Rob Buechler, Tao Sun, Chad Zanocco, Francisco Galuppo, Adriano Veloso, and Ram Rajagopal. Extreme event counterfactual analysis of electricity consumption in brazil: Historical impacts and future outlook under climate change. *Energy*, 281:128101, October 2023.

[26] Jung Min Han, Sunghwan Lim, Ali Malkawi, Xu Han, Elence Xinzhu Chen, Shide Salimi, Tor Helge Dokka, Tine Hegli, and Kristian Edwards. Data-informed building energy management (DiBEM) towards ultra-low energy buildings. *Energy Build.*, 281:112761, February 2023.

[27] Abdurahman S Alrobaie and Moncef Krarti. Measurement and verification building energy prediction (MVBEP): An interpretable data-driven model development and analysis framework. *Energy Build.*, 295:113321, September 2023.

[28] Suhaidi Mohd Aris, Nofri Yenita Dahlan, Mohd Nasrun Mohd Nawi, Tengku Ahmad Nizam, and Mohamad Zamhari Tahir. Quantifying energy savings for retrofit centralized hvac systems at selangor state secretary complex. *Jurnal Teknologi*, 77(5):93–100, 2015.

[29] WT Ho and FW Yu. Measurement and verification of energy performance for chiller system retrofit with k nearest neighbour regression. *Journal of Building Engineering*, 46:103845, 2022.