2022 IEEE 4th International Conference on Cognitive Machine Intelligence (CogMI) | 978-1-6654-7406-1/22/$31.00 ©2022 IEEE | DOI: 10.1109/CogMI56440.2022.00023

A Methodology for Energy Usage Prediction in Long-Lasting Abnormal Events

Gabriele Maurina[1], Hajar Homayouni[2], Sudipto Ghosh[1], Indrakshi Ray[1], Gerald P. Duggan[3]
[1]*Department of Computer Science, Colorado State University*
[2]*Department of Computer Science, San Diego State University*
[3]*Energy Institute, Colorado State University*
[1,3]{gmaurina,ghosh,iray,jerry.duggan}@colostate.edu, [2]hhomayouni@sdsu.edu

***Abstract*—Accurate energy consumption prediction is critical for proper resource allocation, meeting energy demand, and energy supply security. This work aims at developing a methodology for accurately modeling and predicting electricity consumption during abnormal long-lasting events, such as COVID-19 pandemic, which considerably affect consumption patterns in different types of premises. The proposed methodology involves three steps: (A) selects among multiple models the most accurate one in energy consumption prediction under normal conditions, (B) uses the selected model to analyze the impact of a specific abnormal event on energy consumption for various classes of premises, and (C) investigates which features contribute most to energy consumption prediction for abnormal conditions and which features can be added to improve such predictions.**

We use COVID-19 as a case study with datasets obtained from Fort Collins Utilities, which contain energy consumption data for residential and different sizes of commercial and industrial premises in the city of Fort Collins, Colorado, USA. We also use temperature records from NOAA and COVID-19 public orders from Larimer County.

We validate the methodology by demonstrating that the methodology can help design a model suited for the pandemic situation using representative features, and as a result, accurately predict the energy consumption. Our results show that the MLP model selected by our methodology performs better than the other models even when they all use the COVID-related features. We also demonstrate that the methodology can help measure the impacts of the pandemic on the energy consumption.

I. Introduction

Accurate energy consumption prediction is becoming increasingly important for efficient energy management. Rapidly changing consumption patterns can be damaging to the energy provider. Accurate prediction is possible by taking into account the features that affect consumption while modeling the consumption behavior. COVID-19 has affected our normal energy consumption patterns as several changes took place when the rate of infections grew. Students attended remote classes from home, several small businesses closed down, and working from home became the norm in several companies. According to a study by Bartik et al. [1], out of 5,800 samples of small businesses surveyed in 7 days (March 28, 2020, to April 4, 2020), 43% were temporarily closed during the period because of COVID-19.

Machine learning-based models have gained widespread usage for predicting energy demand because of their effectiveness and efficiency but these studies focused on specific types of premises [2], [3]. Moreover, they have been used to predict the energy consumption for a given time period assuming a stable environment. Such approaches may be inadequate during uncertain situations, such as the COVID-19 pandemic because of changes in human behavior and consumption patterns.

The objective of this paper is to present a methodology with three steps to analyze, understand, and predict energy consumption in abnormal situations. Step A is to compare machine learning-based models on the task of energy consumption prediction in a stable situation. Step B is to use the best model to analyze the effect of a long duration abnormal event on the energy consumption patterns. Finally, step C is to create a model that is able to make accurate predictions of energy consumption during that event.

We use the COVID-19 pandemic in the city of Fort Collins, Colorado, as a case study. The research questions that this paper aims to answer for this case study in each class of premises are listed below.

RQ1. What is the most effective machine learning algorithm for predicting energy consumption before the COVID-19 pandemic?

RQ2. What is the impact of COVID-19 on energy consumption?

RQ3. What pandemic-related features can be added to improve the energy consumption prediction during the COVID-19 period? Which features affect energy consumption the most?

To answer RQ1, we use temporal and non-temporal machine learning models that have been demonstrated to be effective in energy consumption prediction in a stable environment [2]–[6]. The temporal models include Long Short Term Memory (LSTM) [7] and the non-temporal ones include Support Vector Regression (SVR) [8] and Multi-Layer Perceptron (MLP) [9]. We use Jan 2017—Feb 2020 data (before the COVID-19 pandemic) for training and evaluating the models.

To answer RQ2, we analyze how the expected energy consumption differs from the actual consumption during the COVID-19 era (Mar 2020—Dec 2021). We used the best model from RQ1 to make the predictions for RQ2. We show how the energy consumption by different types of premises was affected, and identify those that were the most and least affected by the pandemic.

978-1-6654-7406-1/22/$31.00 ©2022 IEEE
DOI 10.1109/CogMI56440.2022.00023

To answer RQ3, we introduce COVID-19-related features, which describe various characteristics of the pandemic. These characteristics include the number of cases, number of hospitalization, and public orders on closures. We measure how the prediction of the best model of RQ1 improves with this addition. We also run a feature importance analysis on the whole dataset to determine which features had the highest impact on the prediction.

The methodology helped select the best model to predict energy consumption, use that model to effectively measure the impact of an abnormal event on the energy consumption, and adapt it to accurately model the energy consumption during the event. Our findings show that the best model to predict energy consumption is the MLP. Our data shows that across all classes of premises, air temperature and time of the day are the strongest predictors of energy consumption. Among COVID-19 related features we added to improve the model, the strongest predictors of energy consumption are number of cases and number of hospitalizations.

The paper is organized as follows. Section II describes the datasets used in this paper. Section III summarizes the temporal and non-temporal machine learning models used in the work. Section IV presents the proposed methodology. Section IV-A compares these approaches on the task of predicting energy consumption before the COVID-19 pandemic. Section IV-B uses the best model of section IV-A and analyzes the impact of COVID-19 on the energy consumption. Section IV-C studies feature importance and improves the best model of section IV-A to predict energy consumption during the COVID-19 pandemic. Section V discusses the related work. Finally, Section VI concludes the paper and outlines directions for future work.

II. Datasets

The work in this paper uses several features that have an impact on energy consumption. Some of these features are known to have an impact on energy consumption (e.g. air temperature) and some are novel features that we have decided to include to help the models and improve the prediction (e.g. number of COVID-19 cases).

We use the readings of energy consumption from the city of Fort Collins. We analyze the consumption of energy for five classes of premises:

- *Residential*: The main use is a place of residence.
- *General Service (GS)*: The yearly average peak-power demand is below 25 kW.
- *General Service-25 (GS-25)*: The yearly average peak-power demand is between 25 and 50 kW.
- *General Service-50 (GS-50)*: The yearly average peak-power consumption is between 50 and 750 kW.
- *General Service-750 (GS-750)*: The yearly average peak-power demand is above 750 kW.

Colorado State University Energy Institute and Fort Collins Utilities use these five classes of premises to categorize the data. Particularly, Fort Collins Utilities uses these classes to set different rates to different types of premises. We have access to the data collected from January 1, 2017, to December 31, 2021, in 15 minute intervals. We aggregate the data into 12-hour intervals: from midnight to noon and from noon to midnight. The datasets have a total number of 76,927 unique premises with a distribution shown in Figure 1. There are

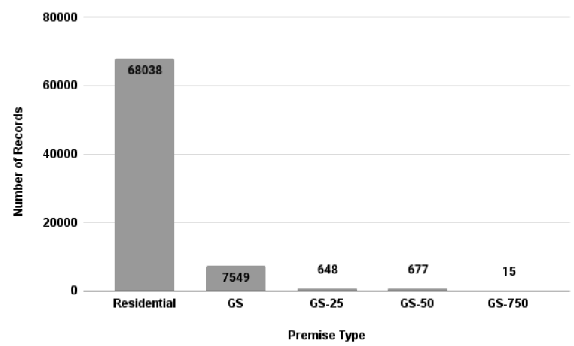

Fig. 1: Distribution of records by premise type: Residential, GS (small-size commercial), GS-25 (medium-size commercial), GS-50 (large-size commercial), and GS-750 (Industrial).

different consumption patterns in these five classes because they are different types of premises (e.g. single family homes, retail shops, restaurants, offices, and industrial premises). We expect different changes in consumption patterns as a result of the COVID-19 pandemic, because it impacted residential, commercial and industrial premises differently.

We use the air temperature data from the National Oceanic and Atmospheric Administration (NOAA) [10]. Air temperature data is necessary because we expect that energy consumption is correlated with temperature. People use more air-conditioning in the summer and more heating during winter. NOAA records the average daily temperature, the minimum daily temperature and the maximum daily temperature for multiple stations in the United States. We choose one station in the geographical center of the city of Fort Collins. We associate the minimum recorded temperature in a particular day with the 12 hour interval from midnight to noon, because that is when the minimum temperature generally occurs during a day. Similarly, we associate the maximum recorded temperature in a particular day with the 12 hour interval from noon to midnight, because that is when the maximum temperature generally occurs during a day.

To answer RQ3 we want to improve the model by adding COVID-19 related features that can influence energy consumption patterns. We then use these features to make predictions of the energy consumption during the COVID-19 pandemic. We use official COVID-19 data from the Larimer County [11] of Colorado. Specifically, we use the datasets: 7-day average number of cases, 7-day average number of hospitalizations, and daily confirmed deaths. We also track school closures and restaurant and bar closures due to the COVID-19 lockdown according to the public health orders and executive orders of the State of Colorado [12]. For schools we also track summer and winter closures following the Poudre School District calendar [13]. Finally we have added a feature called *post_pandemic* that is equal to one on March 10, 2020 and thereafter, and is equal to zero before that date. This feature is meant to represent any long term change in our society that started with the COVID-19 pandemic and influences energy consumption patterns. Machine learning models can use the *post_pandemic* feature to better model changed patterns in energy consumption caused by the COVID-19 pandemic. An example is the fact that many companies today allow workers to work from home one or more days per week indefinitely (i.e., regardless of the status of the COVID-19 pandemic), thus changing energy consumption patterns. It is the task of the models to understand the usefulness of the *post_pandemic* feature in the prediction of energy consumption during the unsupervised learning.

TABLE I: Features of energy data

Feature name	Data type	Description
delivered_kwh	Float	Energy consumed in 12-hour period
temperature	Float	Air temperature
time_of_year	Float	Time of the year, where January 1 equals one, July 1 equals zero, and other days of the year are interpolated between zero and one
weekend	Boolean	1 during the weekend, 0 otherwise
am	Boolean	1 from midnight to noon, 0 otherwise
holiday	Boolean	1 during federal holidays, 0 otherwise
cases	Float	7-day average of COVID-19 cases in Larimer County
hospitalizations	Float	7-day average of COVID-19 hospitalization in Larimer County
deaths	Float	confirmed COVID-19 related deaths in Larimer County
school_closure	Boolean	1 during school closure, 0 otherwise
bar_restaurant_closure	Boolean	1 during bar and restaurants closure, 0 otherwise
post_pandemic	Boolean	1 after March 10, 2020, 0 otherwise

Table I summarizes all the features used in this study. The first six features are not COVID-19 related and are used to answer all three research questions. The last six features are COVID-19 related and are used only to answer RQ3.

The *time_of_year* feature represents the day of the year of a record. It repeats every year (i.e., a particular day of the year will have the same value every year) and it is used to understand the seasonality of energy consumption. It tells the models which period of the year it is (e.g., Summer or Winter). This is important because temperature is correlated with energy consumption. For this reason, days that on average have similar temperatures, should have similar values. Not only May 1 and May 2 should have similar values, but also November 1 and March 1 should have similar values. It is particularly necessary to help non-temporal models understand the cyclic nature of energy consumption. The values of the *time_of_year* feature span from one to zero, with one on January 1, zero on July 2, and all other days with an interpolated value between one and zero. The formula is $|(day\_of\_year - half\_year)/half\_year|$, where day_of_year goes from 0 to 364 (365 on a leap year) and half_year is equal to 182 (183 on a leap year). This feature is created in this fashion because neural networks work best with normalized input values (generally between -1 and 1 or between 0 and 1). The date as a string could definitely not be given to a network as input, nor could the day of the year as an integer between 1 and 365 (366 on leap years). In the first case because a neural network cannot handle a string as input and in the second case because the value is not normalized and the first and the last day of the year would have values at the opposite ends of the spectrum despite being part of the same season and have similar weather. Figure 2 shows the value of the feature *time_of_year* over the course of three years.

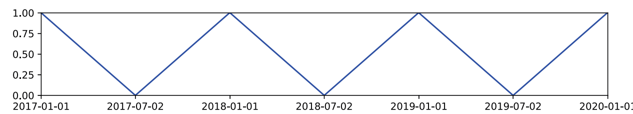

Fig. 2: Value of the *time_of_year* feature over the course of five years

The *holiday* feature indicates whether or not a specific date is a federal holiday or not. Its value is equal to one during a federal holiday and it is equal to zero all other days.

Before feeding the data to the machine learning models, all features are normalized using min-max normalization. This improves the optimization of neural networks during training.

III. Background on ML Models

We use three machine learning models based on temporal and non-temporal features that are known to be effective in energy prediction tasks [2]–[6], [14]. Energy consumption patterns appear to be temporally dependent [15]. For example, the energy consumed by a residential user depends upon the time of the day, the day of the week, and the season. As a result, we use temporal models based on Long Short Term Memory (LSTM) networks, which preserve long-term temporal dependencies among data records in their predictions [7]. On the other hand, the energy consumed by any premise is highly dependent on the daily temperature [2]. Thus, we also use non-temporal techniques (Support Vector Regression (SVR) and Multi-Layer Perceptron (MLP)), which base the prediction of energy consumed on a specific date primarily on the temperature on the same day.

Our objective is to determine whether a non-temporal model can effectively predict the energy consumption without involving the past record values in the prediction. We use these features to help the non-temporal models: *time_of_year*, *weekend*, *am*. These features allow the non-temporal models to understand the different patterns of energy consumption based on time of the year, day of the week and time of the day. However, we do use these feature with the temporal models as well, this is because adding new features could improve, but not worsen, the model prediction.

We describe the temporal and non-temporal models in Sections III-A and III-B respectively.

A. Temporal Machine Learning Model

Our LSTM model is based on a Recurrent Neural Network (RNN) [16] that contains loops in its structure to allow information to persist and make networks learn sequential dependencies among data records [7]. The original RNNs can only learn short-term dependencies among data records by using the recurrent feedback connections [17]. LSTMs extend RNNs by using specialized gates and memory cells in their neuron structure to learn long-term dependencies. The computational units (neurons) of an LSTM are called *memory cells*. An LSTM has the ability to remove or add information to the memory cell state by using *gates*. The gates are defined as weighted functions that govern information flow in the memory cells. The gates are composed of a *sigmoid layer* and a *point-wise operation* to optionally let information through. The sigmoid layer outputs a number between zero (to let nothing through) and one (to let everything through). There are three types of gates, namely, *forget*, *input*, and *output*.

- *Forget gate*: Decides what information to discard from the memory cell.
- *Input gate*: Decides which values to use from the network input to update the memory state.
- *Output gate*: Decides what to output based on the input and the memory state.

In our experiment, we use two LSTM networks. The first network is an original LSTM model that takes a sequence of past records ($R_{t-i},\ (i = 1, ..., p)$) as input to predict a single record (R_{t-i}) as output. The second network is an LSTM-autoencoder model as a sequence-to-sequence predictor, which takes a sub-sequence of records ($R_{t-i},\ (i = 1, ..., p)$) as input to predict a one-step-ahead sub-sequence of records ($R_{t-(i+1)},\ (i = 1, ..., p)$) as output. The LSTM-autoencoder model is capable of extracting and encoding significant features of the input sub-sequence into a new representation, resulting in accuracy improvements and reducing overfitting risk in comparison with the original LSTMs [18].

B. Non-Temporal Machine Learning Models

We use two models, which are Support Vector Regression (SVR) and Multilayer Perceptron (MLP). The SVR is a support-vector machine that is generally used to compute a linear function of its inputs (however, we are going to use non-linear kernels), whereas the MLP is a feedforward artificial neural network that computes a non-linear function of its inputs.

Support Vector Regression (SVR). Support Vector Machine (SVM) is a machine-learning method based on structural risk minimization technique, which balances fitting the training data against model complexity [19]. An SVM sets one or more hyperplanes in a high-dimensional space to perform classification by finding the hyperplane that maximizes the margin between the two data classes [20], [21]. The version of SVM for regression estimation is known as Support Vector Regression (SVR) [8]. SVR fits errors within a certain threshold while minimizing the error rate. Research has shown that SVR performs better than traditional regression methods [22]–[24].

The performance of SVR relies on a kernel, which is a function that helps find a hyperplane in a high-dimensional space with low computational cost [6], [25]. There are four types of kernel functions, namely, linear, polynomial, Radial Basis Function (RBF), and sigmoid. Polynomial kernel shapes a curved line, and RBF kernel creates complex regions. In this work, we test all three non-linear kernels, to make a comparison with the non-linear MLP. It has been shown that polynomial, RBF and sigmoid kernels are effective in modeling non-linear patterns in data [2].

Multi Layer Perceptron (MLP). An MLP is an Artificial Neural Network [26] with an extremely popular usage in energy consumption prediction because of its capability in estimating continuous non-linear functions [27]. Tosun et al. [28] studied the difference in prediction using regression models and neural networks, which shows that the usage of complex non-linear networks can result in great effectiveness for the prediction tasks. An MLP consists of an interconnection of a number of neurons. The input layer receives the input data. The intermediate layers, called hidden layers, perform computations on the input data. An MLP uses a supervised technique called backpropagation [29] for training.

IV. Proposed Methodology

Figure 3 shows an overview of the proposed methodology with three steps. Step A identifies the best model to predict energy consumption in a normal period (i.e. a period in which there are no abnormal events that impact energy consumption, such as pandemics, flooding, and hurricanes). Step B measures the impact of the phenomenon that causes the abnormal situation on energy consumption. The impact is measured as the difference between the expected energy consumed and the actual energy consumed. Step C creates a new model that is able to predict energy consumption during the event. This step is conducted by using the best performing model (previously selected in section IV-A) and adding new features that describe the evolution of the situation. These features provide the model with more information that can be used to learn changing energy consumption patterns and to improve energy consumption prediction. Sections IV-A to IV-C describe these steps using the COVID-19 pandemic case study.

A. Model Comparison

We compare the six (MLP, SVR-poly, SVR-rbf, SVR-sig, LSTM, and LSTM-autoencoder) machine learning models presented in Section III on the task of predicting energy consumption for the period Jan 2017—Feb 2020. The goal is to find the most effective model (i.e., the model with the lowest prediction error). The models take the *temperature*,

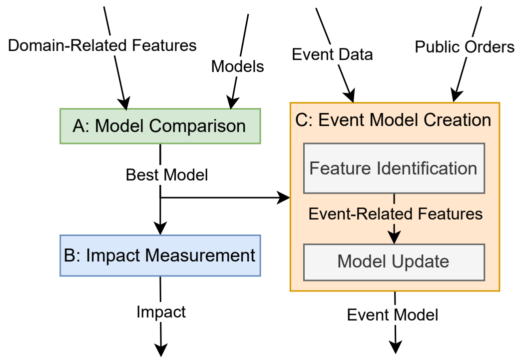

Fig. 3: Methodology Overview

time_of_year, *weekend*, *am* and *holiday* features of an individual record i as input and predict the value of *delivered_kwh*$_i$ for the same record.

Many hyperparameters need to be tuned in order for the models to make the best predictions. We use grid search to find the best combination for each model. The loss function [26] for the neural networks used in this work is the Mean Squared Error (MSE), which is the most commonly used regression loss function and is defined as the sum of squared distances between the target and predicted values. For the MLP, we use two hidden layers of size 20 and the Adam optimizer with a learning rate of 0.001. For the LSTM models we use the LSTM layers followed by one dense layer of size 20 and the Adam optimizer with a learning rate of 0.001.

The following hyperparameters of the SVR models need to be tuned: regularization, gamma, and epsilon. Values of 100, 0.1, and 0.1 for the hyperparameters respectively give the best results. Finally, for the polynomial kernel of the SVR we use a degree of 3.

To compare the models we use the MSE computed on the validation set of each model for each premise class on 12-hour intervals. We use MSE during training because it can amplify larger error and because it is widely used in regression problems. Table II contains the MSE for all models and all classes of premises. All models achieve a low (less than 1%) MSE on the validation set of residential and commercial premises (GS, GS-25, and GS-50). MLP has a 0.1% MSE on residential and commercial premises, which is the lowest of all the models. SVR-poly and SVR-rbf average around 0.2% MSE for residential and commercial premises. Temporal networks perform worst on residential and commercial premises, with an average MSE of around 0.4% and 0.6% for LSTM-autoencoder and LSTM respectively. SVR-sigmoid also has an average of 0.6%.

Industrial premises (GS-750) result in higher MSEs of no more than 3.7%. On industrial premises temporal networks performed best with a MSE of 2.1% and 2.4% for LSTM-autoencoder and LSTM respectively. Of the non-temporal networks, MLP has the lowest MSE on industrial premises (2.8%). SVR models have a slightly higher MSE of 3.0%, 3.2%, and 3.7% for SVR-poly, SVR-rbf, and SVR-sigmoid respectively.

Temporal models perform better on industrial premises and non-temporal models perform better on residential and commercial premises. A possible explanation could be that the class GS-750 only contains 15 premises, making the data more erratic and unpredictable, and in turn making the next value easier to predict by knowing the previous. This is supported by the fact that the MSE on the class GS-750 is higher than the MSE on all other classes of premises for all models. The fact that the class GS-750 only contains 15 premises does not entail that the dataset used to train the models is smaller. The dataset contains the same number of datapoints (i.e. the same number of 12-hour interval energy cunsumption readings). The difference with the other classes is that those readings are from a smaller (only 15) set of premises. Figures 4 shows the different energy consumption pattern for GS and GS-750 premises in the year 2019. Indeed the measured consumption for GS-750 premises appears more random and less predictable.

Overall, all models achieve a low MSE on all classes of premises on the task of predicting energy consumption before the COVID-19 pandemic. MLP is overall the best performing model, closely followed SVR-poly and SVR-rbf.

TABLE II: MSE (%) of models with pre-COVID-19 data

Model	Res	GS	GS-25	GS-50	GS-750	Avg
MLP	0.1	0.1	0.1	0.1	2.8	0.6
SVR-poly	0.3	0.2	0.1	0.2	3.0	0.8
SVR-rbf	0.2	0.2	0.2	0.1	3.2	0.8
SVR-sigmoid	1.0	0.8	0.5	0.3	3.7	1.3
LSTM	1.0	0.6	0.6	0.4	2.4	1.0
LSTM-autoencoder	0.6	0.3	0.3	0.3	2.1	0.7

The answer to RQ1 is MLP, because it has a prediction MSE % of 0.1 in all but one premise class, and has the lowest average MSE among all models (0.6%). MLPs have also been confirmed by researchers to be well suited for various prediction tasks [30], [31]. MLPs outperform SVMs in predicting features of big data as SVMs might become overwhelmed by the curse of dimensionality, and as a result, have poor generalization properties when dealing with high-dimensional data [32], [33]. Moreover, MLPs are capable of learning the relevant features from the data as a result of using a multi-layer architecture with several layers of non-linearity. Although LSTMs are known to be superior to MLPs in case of sequential data [34], the complex architecture of the LSTM model requires more training data points than the MLP model to achieve a higher prediction accuracy. Moreover, the LSTM model requires more cycles of annual data for its training to capture changing factors, such as effects of global warming and technology developments on the energy consumption.

Figure 5 shows the measured and predicted energy consumption for residential premises for the year 2019 for the models MLP, SVR-poly, SVR-rbf, SVR-sig, LSTM, and LSTM-autoencoder, respectively. The blue line in the plots

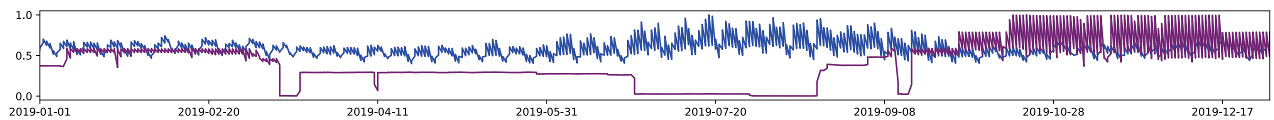

Fig. 4: Normalized measured energy consumption for GS premises (blue) and GS-750 premises (purple)

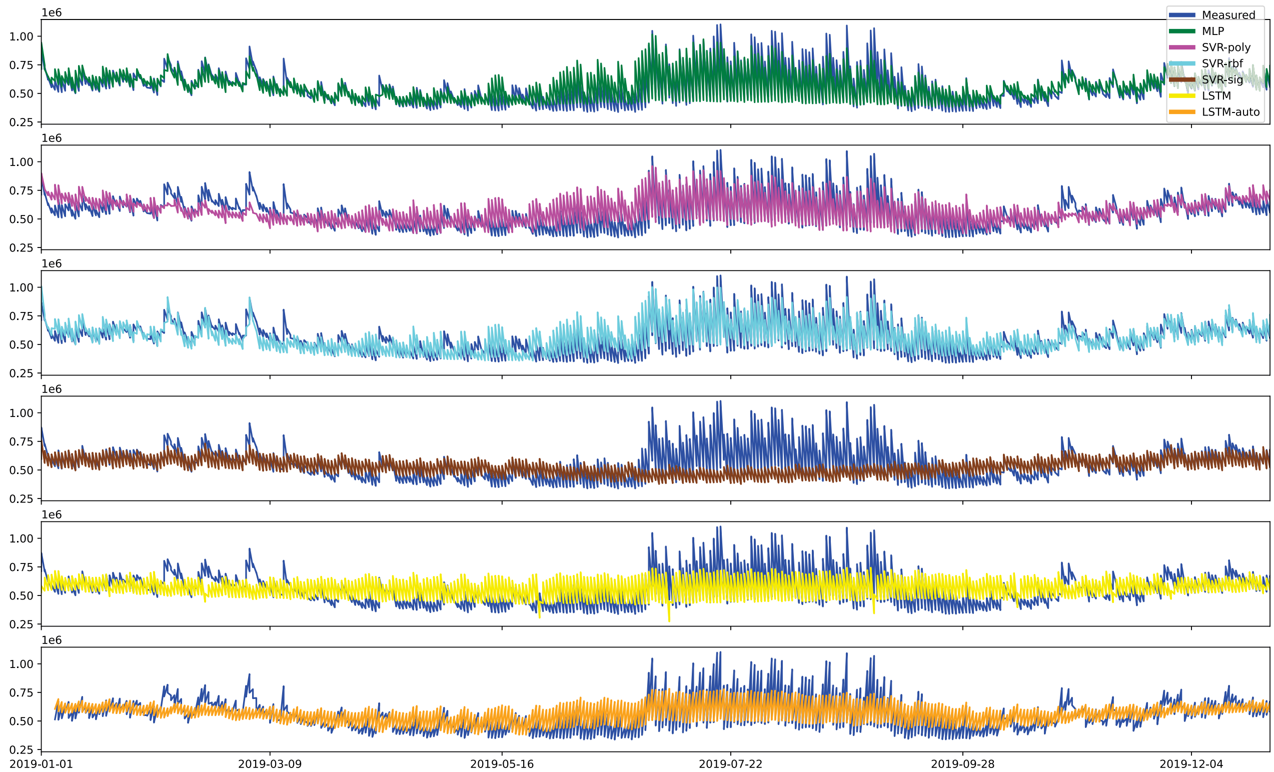

Fig. 5: Predicted energy consumption (KWh) for residential premises in the year 2019

represents the measured energy consumed by residential premises. There is a clear section in the summer period with higher spikes in measured consumption probably caused by the use of air conditioning. This section repeats with similar intensity every summer in the data that we analyzed (2017-2021). The measured consumption also appear to be slightly higher in Winter than in Spring or Fall, with some spikes around February. This finding is consistent in every year that we analyzed (2017-2021). This could be caused by the use of electric heating in some premises. Overall there is great variation in the measured energy consumption from one day to another. This is clearly visible in the plots, as the gray line is not smooth at all. Furthermore, a sawtooth pattern appears to be present both in the measured and in the predicted energy consumption. A sawtooth pattern is sequence that alternates high and low points regularly (i.e. like a sawtooth wave). This is explained by the fact that we analyze and predict energy consumption on 12-hour intervals. So, if the consumption is generally higher in the afternoons compared to the mornings, then the plot will appear to go up and down within a period of 24 hours.

The first plot in figure 5 shows the prediction made by the MLP model (in green) compared to the actual measured consumption (in blue). It is clearly visible in the plots that the MLP prediction closely follows the actual measured energy consumption, as expected by the results shown in Table II.

The second and third plots in figures 5 show the prediction of SVR-poly (in magenta) and SVR-rbf (in cyan) respectively. The predictions closely follow the measured energy consumption and the models are overall slightly less accurate then the MLP. This could be caused by the lower capacity of the two models to understand high dimensional data when compared to a MLP.

The fourth plot in figure 5 shows the prediction of the SVR-sig model (in brown). The plot clearly indicates that the prediction is not accurate and most of what the model is doing is following the sawtooth pattern.

The fifth and sixth plots in figures 5 show the prediction of the temporal models LSTM (in yellow) and LSTM-autoencoder (in orange) respectively. It is visible in the picture that these two models hardly predict the energy consumption. They mostly follow the sawtooth pattern that repeats every 24 hours. It is possible that by analyzing 12-hour intervals, the 24 hour cycle of the energy consumption overpowers any other possible temporal pattern (e.g. seasons and weekends), that could be learnt by the LSTM model, thus making it harder to make accurate predictions.

B. Impact Measurement

The best performing model (MLP) of section IV-A is used to measure the impact of the COVID-19 pandemic on energy consumption for the city of Fort Collins. We use the MLP model trained on pre COVID-19 data to make a prediction for the COVID-19 period. We compare the predicted energy consumption to the actual energy consumption. The difference is the result of different energy consumption patterns during COVID-19. This is because the model takes into account all major factors (e.g., *temperature*, *time_of_year*, and *weekend*) except for the COVID-19 pandemic.

Table III shows the impact of COVID-19 on energy consumption for the period Mar 2020—Dec 2021. The table shows the Mean Absolute Error (MAE) which is the mean absolute difference between the measured and the predicted energy consumption of each 12-hour interval. We use the MAE because it shows the absolute error, (which is more understandable by humans in the context of COVID-19 impact on energy consumption), rather than the squared error in the case of MSE, or the root squared error in the case of RMSE. The table also shows the total measured, total predicted and total difference for the period Mar 2020–Dec 2021.

Energy consumption for residential premises over the period Mar 2020—Dec 2021 increased by 2.66%. This is explained by the fact that people stayed at home more (some working remotely and some attending classes remotely) and thus consumed more energy in their homes than in previous years. The MAE for residential premises is 4.85%, which is higher than the overall difference. This entails that in some 12-hour intervals the consumption was higher than expected and in some it was lower than expected. This means that the consumption patterns, not only increased on average, but also changed w.r.t. when the energy is consumed.

Energy consumption for commercial premises over the period Mar 2020 - Dec 2021 decreased by on average 11%. This decrease is much larger in magnitude than the increase in consumption for residential premises (only 2.66%). This is explained by the fact that lots of bars, restaurants and other commercial businesses shutdown completely during the COVID-19 lockdown, and after that many continued working for months with a decreased number of customers, thus consuming less energy overall. The average MAE is 7.25%, which is higher than the percentage total difference. This, again, proves that the consumption patterns changed w.r.t. when the energy is consumed.

Energy consumption for industrial premises over the period Mar 2020—Dec 2021 decreases by an almost negligible 0.46%. This is explained by the fact that large industrial premises largely did not shutdown during the COVID-19 lockdown because they were deemed essential. Furthermore, among all classes of premises, industrial has the fewest number of premises and is the one with the largest MSE in Section IV-A. This entails that the prediction for industrial premises is the least reliable. This is also shown by the higher MAE of 17.22%, which is unjustified (there is no apparent reason why the energy consumption patterns in industrial premises would change so much).

TABLE III: Impact of COVID-19 on energy consumption

Metric	Res	GS	GS-25	GS-50	GS-750
MAE (%)	4.85	6.07	7.7	7.99	17.22
MAE (MWh)	66.7	14.5	10.1	25.9	3.8
TotM (GWh)	773.4	169.2	93.3	265.2	9.4
TotP (GWh)	753.4	186.5	106.3	299.7	9.4
TotD (%)	2.66	-9.25	-12.2	-11.5	-0.46
TotD (GWh)	20.0	-17.2	-13.0	-34.5	-0.04

MAE: Mean Absolute Error
TotM: Total energy measured
TotP: Total energy predicted
TotD: Total energy difference (TotM - TotP)

The answer to RQ2 is: overall the consumption was reduced by 44.7 GWh or 3.4%. This reduction is plausible because the difference between the total energy consumed during the period Mar 2018 - Dec 2019 and the COVID-19 period (i.e., Mar 2020 - Dec 2021) is -33.3 GWh or 2.5%, which is very similar to the difference between the actual and the predicted energy consumed for the COVID-19 period.

C. Event Model Creation

We add features that describe the COVID-19 pandemic in order to improve the prediction of energy consumption during the COVID-19 pandemic. Here, the goal is no longer to use the best model to measure the impact of the COVID-19 pandemic, but rather to create a model that understands the COVID-19 pandemic and can predict the energy usage with a lower error.

We add the six features discussed in section I related to COVID-19: *cases*, *hospitalizations*, *deaths*, *school_closure*, *bar_restaurant_closure*, and *post_pandemic*. We trained a MLP model with the same hyperparameters described in section IV-A, but with the additional six features.

Table IV contains the MSE, computed on the COVID-19 period, of the MLP model for all classes of premises both with and without COVID-19 features. The table shows a consistent reduction in the MSE for all classes of premises for the model that includes the six COVID-19 features. On average the improvement is 0.9%. Particularly on the classes GS-25 and GS-50 the improvement is 8 fold and 9 fold respectively. These results show that we can improve the energy consumption prediction during the COVID-19 pandemic by adding COVID-19 related features. This model can be used to predict energy consumption during future pandemics that are not related to COVID-19. This is because we can use the same COVID-19-related features: *cases*, *hospitalizations*, *deaths*, *school_closure*, *bar_restaurant_closure*, and *post_pandemic*, but with the data of a different pandemic. This model could also be adapted to a different pandemic by adding or removing features, depending on the nature of the pandemic considered.

TABLE IV: MSE (%) of models with and without COVID-19 features

	Res	GS	GS-25	GS-50	GS-750
without covid features	0.3	0.6	0.8	0.9	5.0
with covid features	0.1	0.2	0.1	0.1	2.6

The proposed methodology indicates that the best model selected in step A, should be used in step C to create the event model. The rational is that the best model in step A is also the best model in step C. We have verified this by training and evaluating all models from section IV-A using the pandemic-related features. Our results showed that MLP is still the best model.

The second part of RQ3 asks about the importance of different features when making a prediction. To answer that question, we compute two metrics: feature importance and feature correlation. We define feature importance for a feature *f* as the MSE computed on the validation set when shuffling the values of feature *f*. The higher the MSE, the more important feature f is. This is because a higher MSE means a greater prediction error, which means that by shuffling the feature (effectively nullifying it, such that the model cannot rely on it in the prediction) the error is increased. Feature correlation for a feature *f* is the Pearson correlation between feature *f* and energy consumption. The higher the correlation, the more important feature *f* is. This is because a high correlation (positive or negative) entails that the two variables increase or decrease at a similar rate. Thus, when one is high, it is easy to predict that the other one will be high too, and vice versa.

Table V contains the feature importance metric (MSE when shuffling a feature) of the MLP model for all classes of premises. *Temperature* and *am* are by far the overall most important features in the prediction, followed by *time_of_year*. The most important COVID-19 feature is *post_pandemic*, followed by *hospitalizations*, *cases*, and *deaths*. Overall, the whole GS-750 class has much higher MSE. This is because of the fact that the MLP model as a higher MSE on the GS-750 class to begin with. For this reason, comparisons cannot be made across classes, but only within.

For residential premises the most important feature is *am*, followed by *temperature*. The feature *am* is very important as shown by the plots in section IV-A, where the sawtooth pattern is very visible and it is directly related to the *am* feature.

For commercial premises the most important feature is *temperature*. This is explained by the fact that air temperature determines heating and cooling of buildings, which in turn determines energy consumption. Air temperature is particularly impactful for commercial premises because they generally have a much larger size and more open internal layout, which require a lot more energy to heat or cool.

For industrial premises the most important feature is *time_of_year*. This seems inexplicable because large industrial premises are not expected to have seasonality in energy consumption. A factor contributing to this result might again be that the prediction for industrial premises is overall less reliable (because of a higher MSE).

TABLE V: Feature importance: MSE (%) of models when shuffling feature

Feature	Res	GS	GS-25	GS-50	GS-750
temperature	1.1	1.0	1.3	1.2	5.5
cases	0.3	0.3	0.2	0.2	5.5
hospitalizations	0.5	0.1	0.2	0.1	3.7
deaths	0.3	0.1	0.1	0.1	3.4
weekend	0.2	0.4	0.3	0.3	3.3
am	1.4	0.5	0.5	0.2	4.5
post_pandemic	0.6	0.4	0.4	0.4	5.5
bar_restaurants_closure	0.2	0.2	0.2	0.2	4.6
school_closure	0.2	0.2	0.2	0.2	5.9
holiday	0.2	0.1	0.1	0.1	3.3
time_of_year	0.7	0.4	0.4	0.3	6.9

Table VI contains the Pearson correlation between features and energy consumption for all classes of premises. High correlation (positive or negative) is a proxy for feature importance in the prediction. The two strongest correlations to energy consumption are *temperature* and *am*. In contrast with table V, in table VI the correlation can be compared between different classes of premises.

As with the feature importance analysis, the feature *am* has the highest correlation for residential premises at -0.6, followed by the feature *temperature* at 0.38. Again, as with the feature importance analysis, the feature *temperature* has the highest correlation for commercial premises.

For industrial premises, the feature *time_of_year* has the highest correlation to energy consumption.

These results show a very high degree of consistency between feature importance and Pearson correlation.

TABLE VI: Feature correlation to energy consumption

Feature	Res	GS	GS-25	GS-50	GS-750
temperature	0.38	0.54	0.64	0.59	-0.21
cases	-0.01	-0.19	-0.26	-0.32	0.04
hospitalizations	0.0	-0.2	-0.28	-0.34	0.05
deaths	0.01	-0.13	-0.19	-0.22	0.0
weekend	0.05	-0.37	-0.27	-0.36	-0.01
am	-0.6	-0.54	-0.61	-0.49	-0.2
post_pandemic	0.04	-0.21	-0.29	-0.38	-0.06
bar_restaurants_closure	-0.02	0.13	0.13	0.18	0.21
school_closure	-0.16	-0.14	-0.14	-0.1	0.4
holiday	0.05	-0.08	-0.09	-0.08	0.04
time_of_year	0.01	-0.22	-0.28	-0.3	0.44

The answer to RQ3 is: The following features can be added to improve the energy consumption prediction during the COVID-19 period: *cases*, *hospitalizations*, *deaths*, *school_closure*, *bar_restaurant_closure*, and *post_pandemic*. Not all COVID-19 related features contribute the same way to the improvement of the prediction. The most important COVID-19 related features are, in order of importance, *post_pandemic*, *hospitalizations*, *cases*, and *deaths*. *school_closure* and *bar_restaurant_closure* have a low importance. A possible reason is that these two features are very coarse grained (i.e. they are a boolean, not a float, and they change very seldom during the course of the pandemic) and that makes them less predictive of energy consumption. For the feature *bar_restaurant_closure*, another reason is that bars and restaurants are a small subset of all commercial premises and their closure impacts mildly the aggregated consumption of all commercial premises. Overall the average MSE improvement caused by the COVID-19 related features is 0.9%. The features that overall affect energy consumption prediction the most are: *temperature*, *am*, *time_of_year* and *school_closure*.

V. Related Work

Machine learning models have been widely used in time series analysis. Empirical research has demonstrated that machine learning algorithms outperform statistical models in complicated time series prediction problems [35].

Recently, the prediction of energy consumption using machine learning models has been highly valued [2], [3], [36]. Seyedzadeh et al. [2] show that Artificial Neural Network (ANN) and Support Vector Machine (SVM) were the most adopted machine learning techniques for energy and electricity prediction. Weather-based features, such as temperature (65% of studies used this feature), humidity (40%), and solar radiation (25%) were the most common features used in those studies. In these studies, machine learning-based approaches have been shown to be effective in energy consumption prediction for specific premise types [37], [38] with relatively stable environmental conditions.

Researchers have also modeled energy consumption under unstable conditions, such as COVID-19 pandemic [39]. Dalcali et al. [40] proposed a hybrid multiple linear regression-feedforward artificial neural network (ANFIS) algorithm to analyze the impact of COVID-19 on the total energy consumption in Bursa, Turkey in 2020. This approach uses environmental conditions (i.e., daily average temperature, wind speed, pressure, and humidity), days of the week, and COVID-19 pandemic precautions (i.e., restrictions governments have applied to reduce and control the impact of the pandemic) as inputs to the algorithm. Their results show that environmental factors had a more pronounced effect on the estimation of electrical energy consumption than days of the week and COVID-19 precautions. Lu et al. [14] used a support vector machine model to predict daily electricity demand of US from January to May 2020. This approach uses the number of daily infections, the number of daily deaths, and a GRSI factor (i.e., an indicator of the degree of lockdown proposed by the Oxford University) as the model's input and demonstrated that using the daily infections results in the highest prediction accuracy and stability. As these approaches were trained and evaluated using only the during-COVID data, they do not measure the effect of pandemic-related features in the model improvement.

Huang et al. [41] used a rolling mechanism called IMSGM that identifies the gap between predicted and the actual energy consumption by industries to analyze the impact of the pandemic on electricity in different time periods relating to the local lockdown policies in China. This approach uses a univariate time-series modeling technique that takes pre-COVID consumption sequence as input to predict the consumption during the pandemic and does not involve any COVID-related features into its modeling. Moreover, none of these approaches provide distinct prediction analysis for different premises.

VI. Conclusions

We proposed a three-step methodology to study, understand and predict energy consumption in abnormal situations. We used the COVID-19 pandemic in the City of Fort Collins as case study to demonstrate the effectiveness of the methodology. The outputs of the proposed methodology are the measured impact of the abnormal event on the energy consumption and a model that can predict energy consumption during the abnormal event. We compared the effectiveness of six machine learning models in predicting energy consumption for five classes of premises before the COVID-19 pandemic. We demonstrated that an MLP model is more accurate in predicting energy consumption compared to more structurally complex models, such as LSTM and SVR. The MLP model achieved an MSE of 0.1%.

We improved by 0.9% the MSE of an MLP model in making energy consumption predictions during the COVID-19 pandemic by including features related to the COVID-19 pandemic (e.g., cases, hospitalizations, and deaths). This model can be used to make predictions about energy consumption during future pandemics, as it contains features that can be adapted to different pandemics (e.g. *cases*, *hospitalizations*, and *deaths*). We showed that the most important features when making an energy consumption prediction are: air temperature, time of the day and day of the year.

Our results show that the energy consumption in residential premises was increased. Out results also show that the energy consumption in small-size (GS), medium-size (GS-25), and large-size commercial (GS-50) premises was reduced. This can be explained by the fact that at the start of the pandemic people mostly stayed at home and bars and restaurants were closed. Industrial premises show a vary minor reduction in energy consumption during the COVID-19 pandemic.

The methodology presented in this paper can be applied in future in case of other pandemics or large scale phenomenons (e.g. hurricanes, fires, and tornandos) that effect energy consumption pattern. Our future work involves the applicability of our models in other abrupt changes of the environment, such as unanticipated weather.

Acknowledgment

This work was supported in part by funding from NSF under Award Numbers CNS 1822118, IIS 2027750, OAC 1931363, Statnett, Cyber Risk Research, ARL, AMI, New Push LLC, and NIST. We would like to thank Pablo Bauleo, Jonathan Care, Michael Authier, and the Fort Collins Utilities for giving us access to their data and feedback on the analysis.

References

[1] A. W. Bartik, M. Bertrand, Z. Cullen, E. L. Glaeser, M. Luca, and C. Stanton, "The Impact of COVID-19 on Small Business Outcomes and Expectations," *Proceedings of the National Academy of Sciences*, vol. 117, no. 30, pp. 17 656–17 666, 2020.

[2] S. Seyedzadeh, F. P. Rahimian, I. Glesk, and M. Roper, "Machine Learning for Estimation of Building Energy Consumption and Performance: A Review," *Visualization in Engineering*, vol. 6, no. 1, pp. 1–20, 2018.

[3] Z. Liu, D. Wu, Y. Liu, Z. Han, L. Lun, J. Gao, G. Jin, and G. Cao, "Accuracy Analyses and Model Comparison of Machine Learning Adopted in Building Energy Consumption Prediction," *Energy Exploration & Exploitation*, vol. 37, no. 4, pp. 1426–1451, 2019.

[4] H.-x. Zhao and F. Magoulès, "A Review on the Prediction of Building Energy Consumption," *Renewable and Sustainable Energy Reviews*, vol. 16, no. 6, pp. 3586–3592, 2012.

[5] S. Paudel, M. Elmitri, S. Couturier, P. H. Nguyen, R. Kamphuis, B. Lacarrière, and O. Le Corre, "A Relevant Data Selection Method for Energy Consumption Prediction of Low Energy Building based on Support Vector Machine," *Energy and Buildings*, vol. 138, pp. 240–256, 2017.

[6] H. Zhong, J. Wang, H. Jia, Y. Mu, and S. Lv, "Vector Field-based Support Vector Regression for Building Energy Consumption Prediction," *Applied Energy*, vol. 242, pp. 403–414, 2019.

[7] Y. Yu, X. Si, C. Hu, and J. Zhang, "A Review of Recurrent Neural Networks: LSTM Cells and Network Architectures," *Neural Computation*, vol. 31, no. 7, pp. 1235–1270, 2019.

[8] I. B. Aydilek and A. Arslan, "A Hybrid Method for Imputation of Missing Values Using Optimized Fuzzy C-Means with Support Vector Regression and a Genetic Algorithm," *Information Sciences*, vol. 233, pp. 25–35, 2013.

[9] H. Taud and J. Mas, *Multilayer Perceptron (MLP)*. Cham: Springer International Publishing, 2018, pp. 451–455.

[10] NOAA, "National oceanic and atmospheric administration," accessed: 2022-08-01. [Online]. Available: https://www.ncdc.noaa.gov/cdo-web/datasets

[11] L. County, "National oceanic and atmospheric administration," accessed: 2022-08-01. [Online]. Available: https://www.larimer.gov/health/communicable-disease/coronavirus-covid-19/larimer-county-positive-covid-19-numbers

[12] C. D. of Public Health and Environment, "Public health orders and executive orders," accessed: 2022-08-01. [Online]. Available: https://covid19.colorado.gov/public-health-orders-and-executive-orders

[13] P. S. D. N. Colorado, "Calendars and schedules," accessed: 2022-08-01. [Online]. Available: https://www.psdschools.org/your-district/calendars

[14] H. Lu, X. Ma, and M. Ma, "A Hybrid Multi-objective Optimizer-based Model for Daily Electricity Demand Prediction Considering COVID-19," *Energy*, vol. 219, p. 119568, 2021.

[15] R. K. Jain, K. M. Smith, P. J. Culligan, and J. E. Taylor, "Forecasting Energy Consumption of Multi-Family Residential Buildings Using Support Vector Regression: Investigating the Impact of Temporal and Spatial Monitoring Granularity on Performance Accuracy," *Applied Energy*, vol. 123, pp. 168–178, 2014.

[16] "Understanding LSTM Networks, Recurrent Neural Networks," https://colah.github.io/posts/2015-08-Understanding-LSTMs/ Accessed (21-7-2021).

[17] T. Guo, Z. Xu, X. Yao, H. Chen, K. Aberer, and K. Funaya, "Robust Online Time Series Prediction with Recurrent Neural Networks," in *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, 2016, pp. 816–825.

[18] A. H. Mirza and S. Cosan, "Computer Network Intrusion Detection using Sequential LSTM Neural Networks Autoencoders," in *2018 26th Signal Processing and Communications Applications Conference (SIU)*, 2018, pp. 1–4.

[19] B. Karaçalı, R. Ramanath, and W. E. Snyder, "A Comparative Analysis of Structural Risk Minimization by Support Vector Machines and Nearest Neighbor Rule," *Pattern Recognition Letters*, vol. 25, no. 1, pp. 63–71, 2004.

[20] R. Collobert and S. Bengio, "SVMTorch: Support Vector Machines for Large-Scale Regression Problems," *Journal of machine learning research*, vol. 1, no. Feb, pp. 143–160, 2001.

[21] K.-R. Müller, A. J. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, and V. Vapnik, "Predicting Time Series with Support Vector Machines," in *International Conference on Artificial Neural Networks*. Springer, 1997, pp. 999–1004.

[22] J.-f. Yang, Y.-j. Zhai, D.-f. Wang, and D.-p. Xu, "Time Series Prediction Based on Support Vector Regression," *PROCEEDINGS-CHINESE SOCIETY OF ELECTRICAL ENGINEERING*, vol. 25, no. 17, p. 110, 2005.

[23] S. Li, Y. Ouyang, W. Wang, and B. Sun, "Multi-Document Summarization Using Support Vector Regression," in *Proceedings of DUC*. Citeseer, 2007.

[24] R. Maity, P. P. Bhagwat, and A. Bhatnagar, "Potential of Support Vector Regression for Prediction of Monthly Streamflow Using Endogenous Property," *Hydrological Processes: An International Journal*, vol. 24, no. 7, pp. 917–923, 2010.

[25] W. Gani, H. Taleb, and M. Limam, "Support Vector Regression based Residual Control Charts," *Journal of Applied Statistics*, vol. 37, no. 2, pp. 309–324, 2010.

[26] C. M. Bishop and P. o. N. C. C. M. Bishop, *Neural Networks for Pattern Recognition*. Clarendon Press, 1995.

[27] M. Ilbeigi, M. Ghomeishi, and A. Dehghanbanadaki, "Prediction and Optimization of Energy Consumption in an Office Building Using Artificial Neural Network and a Genetic Algorithm," *Sustainable Cities and Society*, vol. 61, p. 102325, 2020.

[28] E. Tosun, K. Aydin, and M. Bilgili, "Comparison of Linear Regression and Artificial Neural Network Model of a Diesel Engine Fueled with Biodiesel-Alcohol Mixtures," *Alexandria Engineering Journal*, vol. 55, no. 4, pp. 3081–3089, 2016.

[29] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton, "Backpropagation and the Brain," *Nature Reviews Neuroscience*, vol. 21, no. 6, pp. 335–346, 2020.

[30] E.-H. A. Rady and A. S. Anwar, "Prediction of Kidney Disease Stages Using Data Mining Algorithms," *Informatics in Medicine Unlocked*, vol. 15, p. 100178, 2019.

[31] X. Zhang, S. N. Mohanty, A. K. Parida, S. K. Pani, B. Dong, and X. Cheng, "Annual and non-Monsoon Rainfall Prediction Modelling using SVR-MLP: an Empirical Study from Odisha," *IEEE Access*, vol. 8, pp. 30 223–30 233, 2020.

[32] M. A. Ghorbani, H. A. Zadeh, M. Isazadeh, and O. Terzi, "A Comparative Study of Artificial Neural Network (MLP, RBF) and Support Vector Machine Models for River Flow Prediction," *Environmental Earth Sciences*, vol. 75, no. 6, pp. 1–14, 2016.

[33] S. Osowski, K. Siwek, and T. Markiewicz, "MLP and SVM Networks-A Comparative Study," in *Proceedings of the 6th Nordic Signal Processing Symposium, 2004. NORSIG 2004.* IEEE, 2004, pp. 37–40.

[34] M. Djerioui, Y. Brik, M. Ladjal, and B. Attallah, "Heart Disease Prediction using MLP and LSTM Models," in *2020 International Conference on Electrical Engineering (ICEE)*. IEEE, 2020, pp. 1–5.

[35] A. R. S. Parmezan, V. M. Souza, and G. E. Batista, "Evaluation of Statistical and Machine Learning Models for Time Series Prediction: Identifying the State-of-the-Art and the Best Conditions for the Use of each Model," *Information sciences*, vol. 484, pp. 302–337, 2019.

[36] D. B. Crawley, J. W. Hand, M. Kummert, and B. T. Griffith, "Contrasting the Capabilities of Building Energy Performance Simulation Programs," *Building and environment*, vol. 43, no. 4, pp. 661–673, 2008.

[37] M. K. M. Shapi, N. A. Ramli, and L. J. Awalin, "Energy Consumption Prediction by Using Machine Learning for Smart Building: Case Study in Malaysia," *Developments in the Built Environment*, vol. 5, p. 100037, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S266616592030034X

[38] S. Bourhnane, M. R. Abid, R. Lghoul, K. Zine-Dine, N. Elkamoun, and D. Benhaddou, "Machine Learning for Energy Consumption Prediction and Scheduling in Smart Buildings," *SN Applied Sciences*, vol. 2, no. 2, pp. 1–10, 2020.

[39] L. Zhang, H. Li, W.-J. Lee, and H. Liao, "COVID-19 and energy: Influence mechanisms and research methodologies," *Sustainable Production and Consumption*, vol. 27, pp. 2134–2152, 2021.

[40] A. Dalcali, H. Özbay, and S. Duman, "Prediction of Electricity Energy Consumption Including COVID-19 Precautions Using the Hybrid MLR-FFANN Optimized with the Stochastic Fractal Search with Fitness Distance Balance Algorithm," *Concurrency and Computation: Practice and Experience*, vol. 34, no. 15, p. e6947, 2022, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.6947.

[41] L. Huang, Q. Liao, R. Qiu, Y. Liang, and Y. Long, "Prediction-based Analysis on Power Consumption Gap under Long-term Emergency: A Case in China under COVID-19," *Applied Energy*, vol. 283, p. 116339, Feb. 2021.