



Machine-Learning-Driven, Site-Specific Weather Forecasting for Grid-Interactive Efficient Buildings

Preprint

Rui Yang, Jun Hao, Huaiguang Jiang, and Xin Jin

National Renewable Energy Laboratory

*Presented at the 2020 ACEEE Summer Study on Energy Efficiency in Buildings
August 17–21, 2020*

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5D00-77768
September 2020



Machine-Learning-Driven, Site-Specific Weather Forecasting for Grid-Interactive Efficient Buildings

Preprint

Rui Yang, Jun Hao, Huaiguang Jiang, and Xin Jin

National Renewable Energy Laboratory

Suggested Citation

Yang, Rui, Jun Hao, Huaiguang Jiang, and Xin Jin. 2020. *Machine-Learning-Driven, Site-Specific Weather Forecasting for Grid-Interactive Efficient Buildings: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-5D00-77768. <https://www.nrel.gov/docs/fy20osti/77768.pdf>.

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-5D00-77768
September 2020

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Building Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Machine-Learning-Driven, Site-Specific Weather Forecasting for Grid-Interactive Efficient Buildings

Rui Yang, Jun Hao, Huaiguang Jiang, and Xin Jin, National Renewable Energy Laboratory

ABSTRACT

Emerging grid-interactive efficient buildings (GEBs) have great potential to provide much-needed demand flexibility to electric grids while fulfilling their own control targets by co-optimizing smart appliances, solar photovoltaics, electric vehicles, and energy storage at buildings. To enable the optimal operation of GEBs, site-specific weather information—such as temperature, solar irradiance, relative humidity, and wind speed—is crucial; however, this information is generally unavailable or expensive to obtain. This paper develops advanced machine learning methods to provide precise weather forecasts for individual building sites using readily available weather station data. Support vector regression and artificial neural networks have been employed to learn the spatiotemporal correlations between the weather conditions at nearby weather stations and the individual building site. The proposed site-specific weather forecasting methods have been validated using 1-year actual weather measurement data collected in the Denver metro area. Results show that the developed machine-learning-driven methods can accurately forecast the temperature at the target building site 1 hour ahead with mean absolute error less than 0.72°C and a 48% improvement over the persistence method. Site-specific weather forecasts will improve the understanding of the microclimate effect and its impact on building energy consumption. This information will drive efficiency upgrades and adjustments of building control strategies to improve energy savings and increase flexibility in building loads.

Introduction

Residential and commercial buildings accounted for nearly 40% of total energy use and more than 74% of electricity use in the United States in 2019 (EIA 2020). Efficiency upgrades and operation automation could significantly reduce building energy consumption across the nation. Moreover, by integrating and coordinating smart appliances, solar photovoltaics, electric vehicles, and energy storage, the emerging grid-interactive efficient buildings (GEBs) can provide much-needed demand flexibility to power systems. They can participate in demand response programs to shed or shift their energy use during peak-demand hours (Jin et al. 2017; Yoon, Baldick, and Novoselac 2014; Zhou, Zhao, and Wang 2011) and provide grid services such as frequency regulation (Lin et al. 2015; Kim, Fuentes, and Norford 2015; Zhao et al. 2013) and voltage support (Zhu et al. 2018; Utkarsh et al. 2020).

To unlock the building energy savings and grid benefits, it is crucial to accurately forecast building energy consumption and optimally coordinate controllable resources in buildings for mutual benefits to the grid, building owners, and occupants. Because weather-driven loads such as heating and space cooling contribute to a large portion of loads in both commercial and residential buildings, accurate, site-specific weather information—such as temperature, solar irradiance, relative humidity, and wind speed—is key to understanding building behavior patterns and enabling optimal building controls. Weather information is widely incorporated as input to improve the prediction accuracy of building energy consumption using

various types of data-driven methods (Deb et al. 2017; Xu, Wang, and Tang 2019; Kwak et al. 2013). On the other hand, advanced building control technologies, such as model predictive control, also use weather forecasts to inform the optimal control strategies for appliances and distributed energy resources in buildings; therefore, buildings can proactively respond to heating and cooling demands as well as grid events to reduce their energy cost and provide grid services (Jin et al. 2017; Li and Jin 2014; Florita and Henze 2009; Lazos, Sproul, and Kay 2014; Thieblemont et al. 2017). Without accurate, site-specific weather information, the flexibility provided by GEBs might not be accurately captured or fully harvested to enhance the reliability and resilience of the electric grid.

Site-specific weather information, however, is generally unavailable or expensive to obtain for individual buildings. Most existing building analyses and studies use either representative weather data, such as typical meteorological year (TMY) data (Wilcox and Marion 2008), or historical weather measurements collected at designated weather stations, mostly at airports (NOAA 2001). These available weather data might not accurately reflect local weather conditions for individual buildings because TMYs represent median weather conditions averaged during multiple years and buildings might be located far from and in very different environments than the weather stations. Moreover, deviations in weather conditions could lead to significant uncertainty for building energy consumption. For example, one study showed that monthly building energy consumption could vary by $\pm 40\%$ using different weather data sets for the same location (Bhandari, Shrestha, and New 2012). Therefore, accurate, site-specific weather forecasts for individual buildings need to be developed for better building energy prediction and controls.

Although a variety of weather forecasting methods and products exist—including both physics-based (Al-Yahyai, Charabi, and Gastli 2010; Mathiesen and Kleissl 2011; Xie, Sengupta, and Dudhia 2016) and data-driven approaches (Ren, Suganthan, and Srikanth 2015; Voyant et al. 2017; Dobbs et al. 2017)—none have been adopted to forecast weather conditions at specific building locations. Physics-based approaches, such as numerical weather prediction models, have been employed to predict weather conditions, such as wind speed and solar irradiance, down to a geographic scale of 2 km by 2 km (Freedman et al. 2014; Cheung et al. 2015); however, they lack the capability to further scale down to individual building sites because of the significantly increased complexity of modeling atmospheric dynamics at finer geographic resolutions. On the other hand, data-driven approaches, such as machine learning algorithms, have shown promising accuracy for short-term weather forecasts. These data-driven methods directly use time-series weather data to predict future weather conditions; however, they rely on the available weather data and have not been investigated or employed to forecast site-specific weather conditions for individual buildings.

In this paper, machine-learning-driven forecasting methods are developed to provide accurate, site-specific weather forecasts for individual building sites in the short-term future using readily available weather station data. Support vector regression (SVR) and artificial neural networks (ANNs) are employed to learn the spatiotemporal correlations between the weather conditions at nearby weather stations and the building site. Once the SVR and ANN models are trained, only the measurement data from nearby weather stations are needed to predict the local weather conditions at the building site. The developed site-specific weather forecasting methods are validated to forecast the outdoor dry bulb temperature, one of the most important weather variables for building energy consumption (Fikru and Gautier 2015), 1 hour ahead using 1-year actual temperature measurements collected from 13 weather stations in the

Denver metro area. The forecasting accuracy of the developed site-specific weather forecasting methods are evaluated and improved considering different input data as well as structures and hyperparameters associated with SVR and ANN models.

Site-specific weather forecasts provided by the machine learning algorithms will improve the understanding of the microclimate effect and its impact on building energy consumption. This information will drive efficiency upgrades and adjustments of building control strategies to improve energy savings and increase flexibility in building loads.

Site-Specific Weather Forecasting Methods

To forecast the weather conditions at specific building locations using historical measurements from nearby weather stations, the goal is to learn a function that maps historical weather station data to site-specific weather conditions; therefore, the inputs to the machine learning models are the time-series weather measurements from a certain time period collected at available weather stations around the target building site, and the output are the predicted weather conditions at the target building site in the short-term future. A function with certain structures will be learned in the training stage and then used to forecast site-specific weather conditions in the online operation. SVR and ANNs both have the capability to model a complex relationship between the input and output data; therefore, in this paper, these two methods are employed to forecast site-specific weather conditions.

The developed machine-learning-driven methods have superior flexibility by considering different weather prediction variables, forecasting horizons, and input data horizons. In this paper, the outdoor dry bulb temperature is used as an example to demonstrate the developed site-specific weather forecasting methods because outdoor temperature is one of the most important weather variables that impacts building energy consumption and controls. The developed site-specific weather forecasting methods can be employed to forecast other weather variables, such as relative humidity and wind speed, for the specific building site by changing the input and output of the machine learning models to the considered weather variables. Alternatively, the developed machine-learning-based methods can jointly forecast the temperature, relative humidity, and wind speed at the individual building site by taking the historical temperature, relative humidity, and wind speed at nearby weather stations as inputs. Also, the forecasting horizon can be adjusted according to the specific applications for buildings, with appropriate input data selected that yield the best forecasting performance.

In the following sections, the SVR and ANN models are presented in detail.

Support Vector Regression

SVR is a variant of the classic support vector machine (SVM) that is widely adopted for regression problems (Smola and Schölkopf 2004; Burges 1998). The goal of the SVR is to find a function, f , between the input, x , and output, y , in the form of:

$$f(x) = \langle w, \phi(x) \rangle + b \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, and $\phi(x)$ is a function that maps the input, x , to some feature space. To determine the parameters w and b in the function f , an optimization problem is formulated (Smola and Schölkopf 2004):

$$\min_d \left[\frac{1}{2} w^T w + C \sum_{i=1}^{N_s} (\xi_i + \xi_i^*) \right] \quad (2)$$

$$\text{s. t. } y_i - w^T \phi(x_i) - b \leq \varepsilon + \xi_i, \quad (3)$$

$$w^T \phi(x_i) + b - y_i \leq \varepsilon + \xi_i^*, \quad (4)$$

$$\xi_i, \xi_i^* \geq 0, i = 1, \dots, N_s \quad (5)$$

where N_s is the number of training samples. The deviation of the estimated values of the output from the actual values is bound by the tolerance, ε , and the slack variables ξ_i and ξ_i^* represent the potential violations of the given tolerance. Constant C is a predetermined hyperparameter that represents the trade-off between minimizing the violations of the deviation bounds and the flatness of the function f .

The dual problem of the formulated optimization is:

$$\min \left[\frac{1}{2} \sum_{i=1}^{N_s} \sum_{j=1}^{N_s} (\alpha_i - \alpha_i^*)(\alpha_j - \alpha_j^*) k(x_i, x_j) \right. \\ \left. + \varepsilon \sum_{i=1}^{N_s} (\alpha_i + \alpha_i^*) - \sum_{i=1}^{N_s} y_i (\alpha_i - \alpha_i^*) \right] \quad (6)$$

$$\text{s. t. } \sum_{i=1}^{N_s} (\alpha_i - \alpha_i^*) = 0 \quad (7)$$

$$\forall i: 0 \leq \alpha_i \leq C \quad (8)$$

$$\forall i: 0 \leq \alpha_i^* \leq C \quad (9)$$

where α_i and α_i^* are the Lagrange multipliers associated with Equations (3) and (4) for the i -th training sample, and α_j and α_j^* are the Lagrange multipliers associated with Equations (3) and (4) for the j -th training sample. The function $k(x_i, x_j)$ is the kernel function:

$$k(x_i, x_j) = \langle \phi(x_i), \phi(x_j) \rangle \quad (10)$$

where $\phi(x)$ is a transformation that maps x to a high-dimensional space.

By using the kernel, the function f can be written as:

$$f(x) = \sum_{i=1}^{N_s} (\alpha_i - \alpha_i^*) k(x_i, x) + b \quad (11)$$

Nonzero α_i and α_i^* correspond to the support vectors; therefore, the function f depends only on the support vectors. By solving the dual optimization problem, the function f can be estimated as a linear combination of kernels. Because training an SVR model requires solving a quadratic optimization problem, it is usually computationally intensive with large data sets.

Commonly used kernel functions include the linear kernel, polynomial kernel, and (Gaussian) radial basis function (RBF) kernel. The formulations of these kernels are:

$$\text{Linear: } k(x, x') = \langle x, x' \rangle \quad (12)$$

$$\text{Polynomial: } k(x, x') = (\langle x, x' \rangle + c)^p \quad (13)$$

$$\text{RBF: } k(x, x') = \exp(-\gamma \|x - x'\|^2) \quad (14)$$

In this paper, all three types of kernels are employed with different hyperparameters considered, such as the tolerance bound ε and regulation parameter C .

Artificial Neural Networks

To model the complex relationship between the input and output data, an ANN is also widely adopted (Bishop 1995). ANNs comprise several neurons, and each neuron has an activation function mapping its input to output. The neurons in two consecutive layers are interconnected to form a multilayer network, which is used to define the relationship between the input variables and output variables. In this paper, ANN models are trained to find the mapping between the weather conditions at specific building locations and weather measurements from nearby weather stations using historical data. The mapping function is then used to forecast site-specific weather conditions in the short-term future.

Figure 1 depicts a general structure for ANNs, which usually contain the input layer, the hidden layer, and the output layer. Similar to an SVR, the goal of ANNs is to learn a function, f , that maps the input, x , to the output, y , though how the function f is constructed is different in these two methods. To illustrate how the mapping function, f , is constructed in ANNs, we employ a simple ANN structure with one hidden layer.

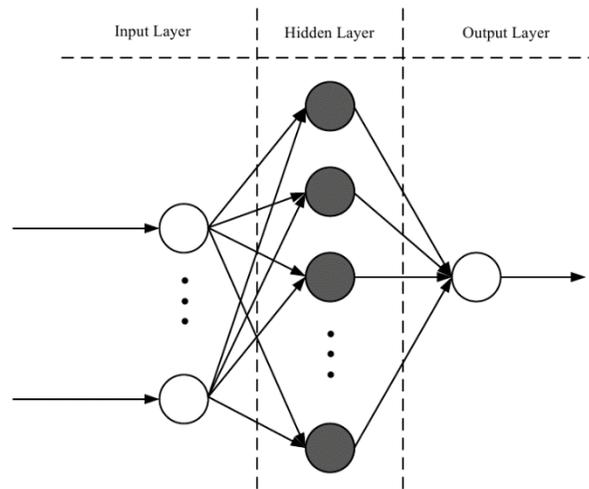


Figure 1. General structure of ANNs (Jiang and Zhang 2016).

Given the input, x , each neuron in the hidden layer calculates a nonlinear input-output mapping by imposing an activation function on the weighted summation of the inputs as:

$$a = g(W_1^T x + b_1) \quad (15)$$

where a represents the output of the hidden layer, W_1 represents the weights of the hidden layer, b_1 is the bias added to the hidden layer, and the function $g(\cdot)$ is the activation function. Commonly used activation functions include the rectified linear unit, hyperbolic tangent function, and sigmoid function.

From the hidden layer to the output layer, the estimated output \hat{y} can be written as:

$$\hat{y} = W_2^T a + b_2 \quad (16)$$

where W_2 represents the weights of the output layer, and b_2 is the bias added to the output layer; therefore, the function f can be written as:

$$f(x) = W_2^T g(W_1^T x + b_1) + b_2 \quad (17)$$

To determine the optimal weights and biases, the objective is to minimize the deviations of the estimated output values and actual output values. Hence, a loss function \mathcal{J} is formulated and minimized given the training samples $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$:

$$\min \quad \mathcal{J} = \frac{1}{N_s} \sum_{i=1}^{N_s} \|y_i - \hat{y}_i\|_2^2 + \frac{\lambda}{2N_s} \sum_{l=1}^L \|W_l\|_F^2 \quad (18)$$

where N_s is the number of training samples, and L is the total number of layers (except the input layer) in the considered ANN. $\|\cdot\|_2$ and $\|\cdot\|_F$ are the ℓ_2 norm of a vector and the Frobenius norm of a matrix, respectively. λ is a regularization parameter that balances fitting the training set and generalization over unseen data. The optimal parameters in an ANN are usually determined using gradient-based methods, such as the backpropagation algorithm (Bishop 1995).

To decide the number of neurons in each hidden layer, we reference the following equation:

$$\frac{N_s}{\beta(N_{in} + N_{out})} \quad (19)$$

where N_s is the number of training samples, N_{in} is the number of inputs, and N_{out} is the number of outputs. β is a constant, which is usually within the range from 2 to 10.

Validation Results

To validate the performance of the developed SVR- and ANN-based site-specific weather forecasting methods, 1-year actual weather measurement data collected at 13 weather stations in the Denver metro area are used. Because 1-hour-ahead outdoor dry bulb temperature forecasts can be used to inform building controls for preheating or precooling the buildings to save their energy consumption, in this paper, the developed machine-learning-based site-specific weather forecasting methods are demonstrated to accurately predict the outdoor temperature at the specific building site 1 hour ahead. As mentioned, the developed forecasting methods can be extended to forecast site-specific weather conditions—including temperature, relative humidity, and wind speed—in different forecasting horizons. The SVR and ANN models have been fine-tuned considering different input data horizons, model structures, and hyperparameters to improve the forecasting accuracy.

In the following sections, the weather data used in this paper are presented first. Then the SVR and ANN models with the best forecasting accuracy are presented. The impact of the input data horizon on site-specific weather forecasting accuracy is also discussed. Finally, the effectiveness of the developed SVR- and ANN-based methods is demonstrated by comparing their performance to a baseline method.

Validation Data Set

The actual weather measurement data used in this paper are obtained from Earth Networks for the entire year of 2017. The weather sensors of Earth Networks are mostly deployed at schools; therefore, these weather sensors measure the weather conditions at school buildings. Based on the availability and quality of the data, 13 weather stations are selected that cover the Denver metro area. Figure 2 shows the locations of these 13 stations. The weather station denoted by “DNVCC” in downtown Denver (red circle in Figure 2) is the target building

site, and the historical measurements from the remaining 12 stations are used to forecast the weather conditions at the target site.

The actual temperature measurement data are preprocessed by taking the hourly average to create the data set for training and testing the developed machine-learning-based forecasting methods. This 1-year data set is randomly divided into a training set with the data from 274 days and a testing set with the data from the remaining 91 days. The goal is to forecast the temperature at the target building site 1 hour ahead using the temperature measurements from up to 24 hours beforehand at the remaining 12 weather stations in the Denver metro area.

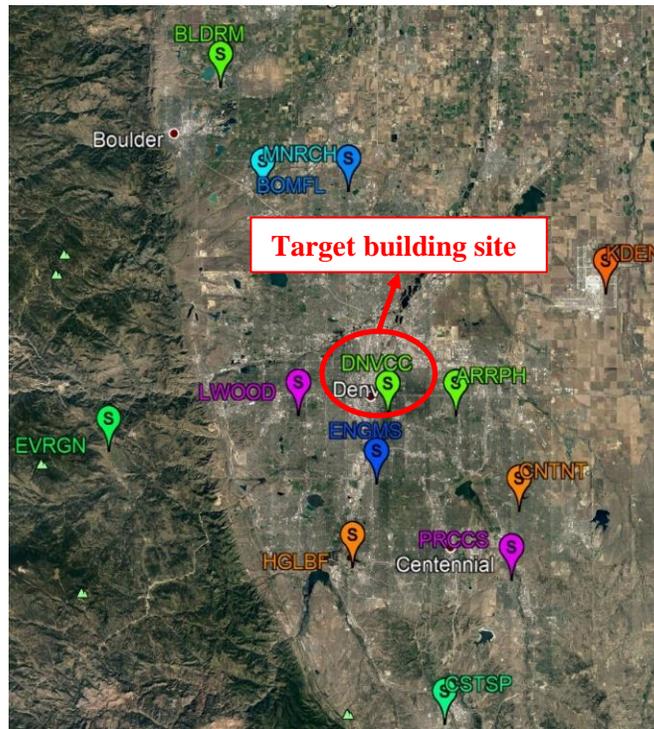


Figure 2. Selected weather stations in the Denver metro area for algorithm validation.

SVR Results

The SVR models using linear, polynomial, and RBF kernels with different hyperparameters have been extensively tested. The best-performing SVR model with the smallest mean absolute error (MAE) for the 1-hour-ahead temperature forecasts at the target building site employs the linear kernel with the tolerance, ϵ , as 0.1 and regularization parameter, C , as 1 and uses the temperature measurements from the past 16 hours at nearby weather stations as inputs. Figure 3 shows the 1-hour-ahead temperature forecasting errors (actual temperature minus predicted temperature) versus the predicted temperature values for the target building site. Figure 4 shows the distribution of the forecasting errors in both the training and testing sets.

As shown in these two figures, the forecasting errors for the temperature at the target building site mostly fall around 0 in both the training and testing sets and the error distribution follows a normal distribution. The MAE is 0.73°C in the training set, with 95% of the absolute forecasting errors less than 2.17°C . The forecasting accuracy in the testing set is slightly better than the training set. The MAE is 0.72°C , with 95% of the absolute forecasting errors less than

2.10°C. As shown in Figure 3, there are three training samples with forecasting errors greater than 8°C. These outliers correspond to times when the temperature at the target building site changes significantly from one hour to the next, rendering it more challenging to accurately predict these changes. Although outliers are observed in the training set, this SVR model generalizes well in the unseen testing data set with a slightly smaller average error and significantly reduced maximum error. Overall, the SVR-based method can accurately forecast the temperature at the target building site using the temperature measured at weather stations in the area.

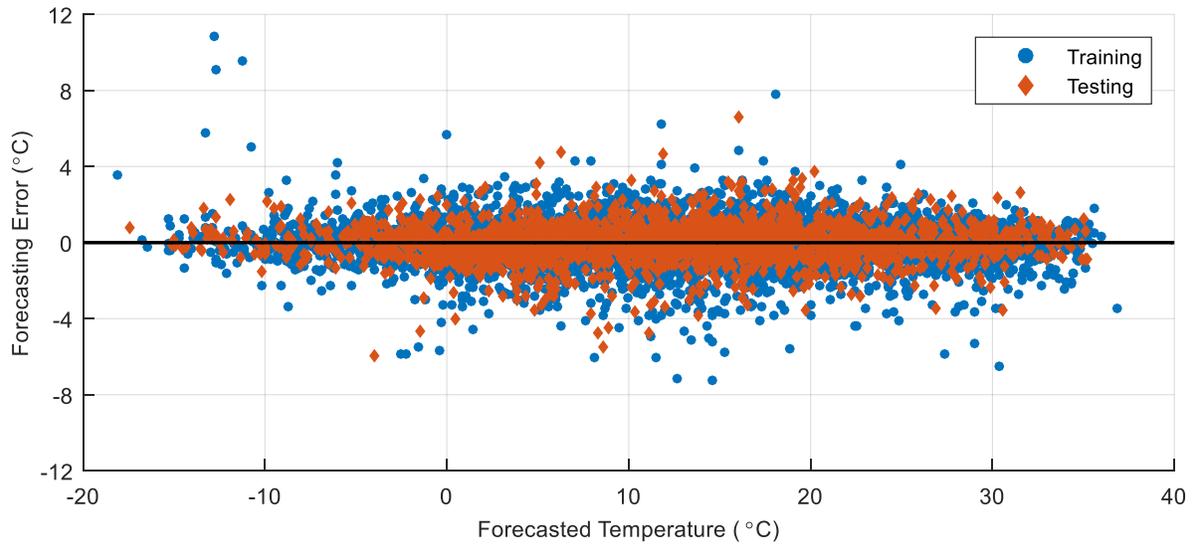


Figure 3. One-hour-ahead temperature forecasting errors versus forecasted temperature in the testing set using SVR.

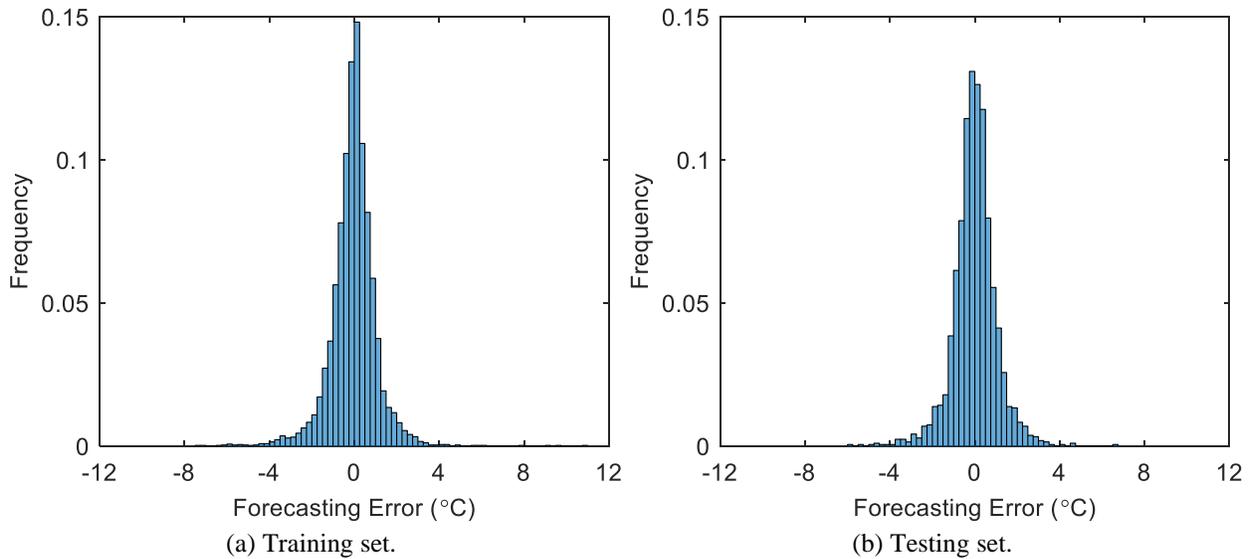


Figure 4. Error distribution of the forecasted temperature at the target building site using SVR.

ANN Results

For the ANN models, different numbers of hidden layers, numbers of neurons, and activation functions are considered. We observed that ANN models with a simple structure, i.e., a single hidden layer, already provide accurate forecasting results for the temperature at the

target building site. The ANN model with the smallest MAE for the 1-hour-ahead temperature forecasts contains one hidden layer with 10 neurons and uses a rectified linear unit as the activation function. The inputs to this best-performing ANN model are the temperature measurements from the past 2 hours at the weather stations near the target building site. Figure 5 shows the 1-hour-ahead temperature forecasting errors (actual temperature minus predicted temperature) versus the predicted temperature values for the target building site. Figure 4 shows the histograms of the forecasting errors in both the training and testing sets.

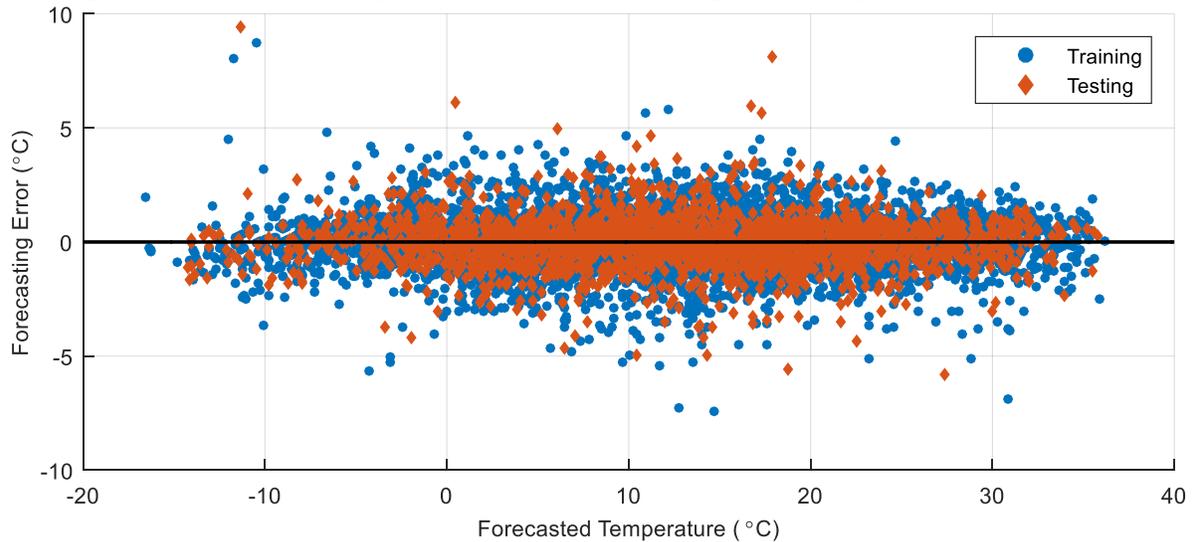


Figure 5. One-hour-ahead temperature forecasting errors versus forecasted temperature in the testing set using ANN.

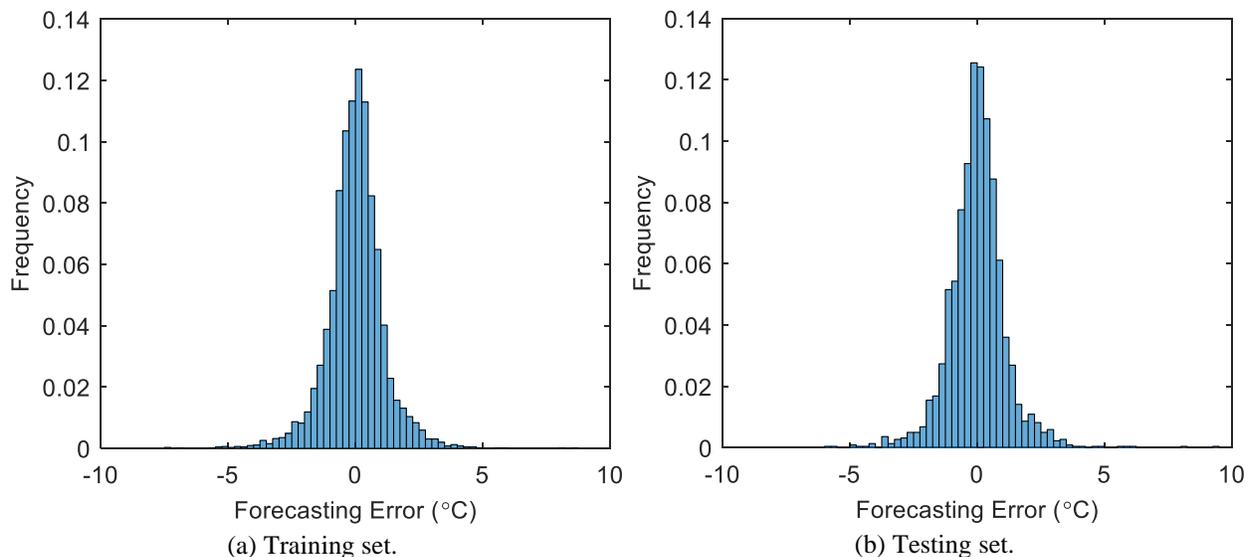


Figure 6. Error distribution of the forecasted temperature at the target building site using ANN.

The forecasting performance of employing the ANN method is comparable to that of the SVR method. In the training set, the MAE is 0.78°C , with 95% of the absolute forecasting errors less than 2.35°C . In the testing set, the MAE is 0.79°C , with 95% of the absolute forecasting errors less than 2.34°C . Although there are two training samples and one testing sample with forecasting errors close to 10°C , which are caused by the big changes in the temperature at the

target building site, the ANN-based forecasting method generally performs well. Compared to the results using SVR in Figure 4, both the MAE and the 95th percentile value for 1-hour-ahead temperature forecasts using ANN are slightly higher; however, SVR models require a longer training time because of the large training set used in this work. The performance comparison of the developed SVR and ANN models will be further discussed in a later section.

Impact of Input Data Horizons

Because the developed machine learning methods aim to learn a relationship that maps historical weather station data to site-specific weather conditions, selecting the input data plays an important role for the forecasting performance. In this paper, the impact of different input data horizons on the forecasting accuracy is studied. The input data horizons considered here are from 1 hour before, which means using only the temperature measurements at time t and $t - 1$ from the nearby weather stations to forecast the temperature at the target building site at time t , up to 24 hours before, which means using the measurements from time t to time $t - 24$, with the increment of 1 hour. Because similar trends are observed for the SVR and ANN models, the forecasting results using the ANN models are shown to demonstrate the impact of input data horizons. Training an ANN model requires solving a nonconvex optimization problem; hence, each run of training the ANN model might converge to a different local optimum. In this paper, 300 runs of training the ANN model are performed, and the parameters of the ANN model that yield the smallest MAE in the training set are selected. Figure 7 depicts the MAE and root mean square error (RMSE) for 1-hour-ahead temperature forecasts in the testing set using different horizons of input data with the best parameters selected from the 300 runs.

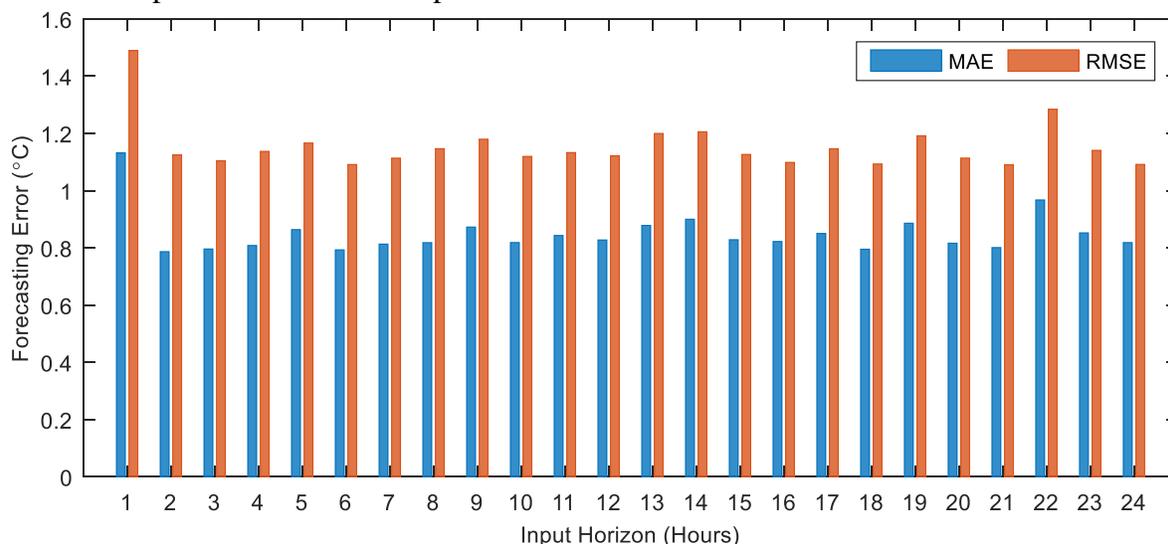


Figure 7. One-hour-ahead temperature forecasting accuracy with different horizons of input data using ANN.

When using only the temperature measurements at the nearby 12 weather stations at the current time step and the past 1 hour to predict the temperature at the target building site for the next 1 hour, the developed machine-learning-based methods can provide accurate forecasts with MAE less than 1.14°C. By using historical weather measurements from the past 2 hours to 24 hours, the forecasting accuracy can be significantly improved compared to the accuracy by using only the past 1 hour of data as inputs. For example, by using the weather measurements from up to 2 hours before, the MAE for 1-hour-ahead forecasts at the target building site can be reduced

by 30% and RMSE by 36%. This means that by using 2 hours of data as inputs, the developed ANN-based forecasting can reduce the average forecasting error and concentrate the forecasting errors in a narrower band than the forecasting performance by using only the past 1 hour of data. When the input data horizon keeps increasing from 2 hours to 24 hours, the differences in the forecasting accuracy are insignificant. Using the past 18 hours of weather station data as inputs yields the smallest RMSE in all the tests, with a 3% decrease compared to the RMSE using the past 2 hours of data as inputs. The main reason for this is that we aim to forecast the temperature at the target building site in the next 1 hour. Given the short forecasting horizon and the fact that the temperature might not dramatically change during a short time, using the past 2 hours of weather station data already provides very accurate forecasts for the target building site. Using longer input data horizons might not significantly reduce the forecasting errors, but it might increase the computation time. Another observation from Figure 7 is that using more historical data could lead to a slightly worse forecasting accuracy, such as the results shown with the input data horizon at 22 hours. This is because training an ANN model requires solving a nonconvex optimization problem, and each training session could end up with a different local minimum. Better initialization and fine-tuning the hyperparameters could further reduce the forecasting errors.

Performance Comparison

To further demonstrate the effectiveness of the developed machine-learning-based methods, the forecasting performance of the developed methods is compared to a persistence method as the baseline, where the 1-hour-ahead temperature at the target building site equals the current temperature measured at the closest weather station. The closest station is the station that has the smallest average absolute temperature deviation during the entire year compared to the temperature at the target building site. In this test, the closest station is denoted by “ARRPH” in Figure 2. Both the SVR and ANN models take the weather station measurements from the past 2 hours as inputs to provide 1-hour-ahead forecasts at the target building site. The MAE, RMSE, and 95th percentile value for the absolute temperature forecasting errors are used as the evaluation metrics. Figure 8 shows the comparison of the forecasting performance using the SVR, ANN, and persistence methods.

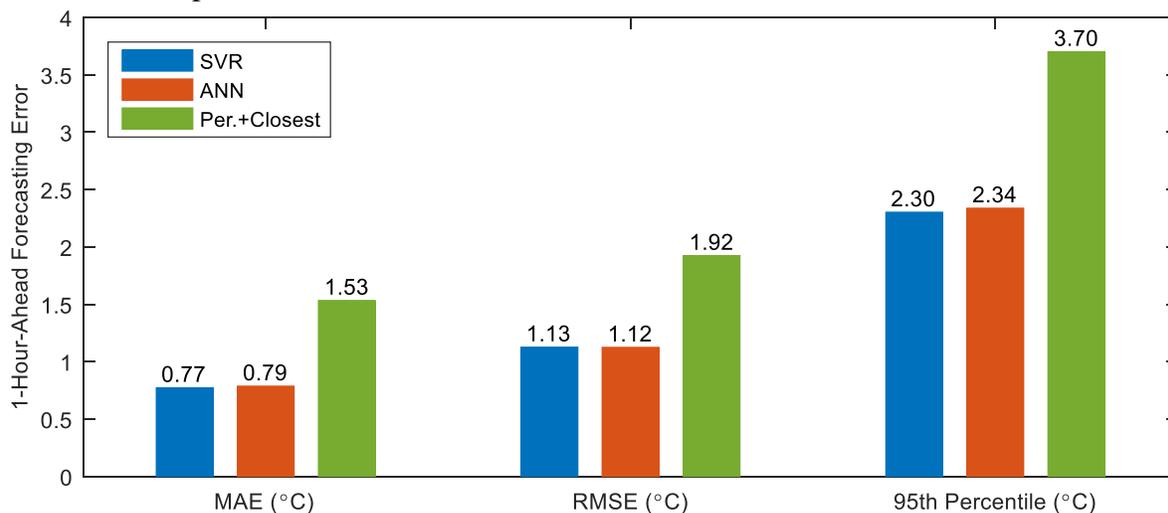


Figure 8. Comparison of 1-hour-ahead temperature forecasting performance using SVR, ANN, and persistence methods.

The two models based on machine learning methods have similar forecasting performance. The SVR model with a linear kernel has a slightly smaller MAE, whereas the ANN model with one hidden layer yields a slightly better RMSE. The persistence method does not perform well because the temperatures measured at the closest weather station might not represent the temperatures at the target location well. By employing the machine learning methods, such as SVR and ANN, the MAE can be reduced by 48%, the RMSE by 42%, and the 95th percentile value by 37%.

In summary, the developed SVR- and ANN-based methods can accurately forecast the temperature at the specific location using historical temperature measurements from nearby weather stations.

Scalability

To evaluate the scalability of the developed site-specific weather forecasting methods, a different weather station, denoted by “ENGMS” in Figure 2, is chosen as the target building site. Historical weather measurements from the other 12 stations in only 91 days are used to train the developed SVR and ANN models, and the performance is evaluated using the remaining 274 days. The same model structure and hyperparameters are employed as those used for forecasting the temperature at “DNVCC”. Here, results using the ANN model are shown to demonstrate the scalability of the developed site-specific weather forecasting methods. Figure 9 shows the distribution of the 1-hour-ahead temperature forecasting errors for “ENGMS” in both the training and testing sets. Figure 10 shows the comparison of MAE, RMSE, and 95th percentile value by using the developed ANN-based method and a persistence method using the data of the closest station “PRCCS” shown in Figure 2.

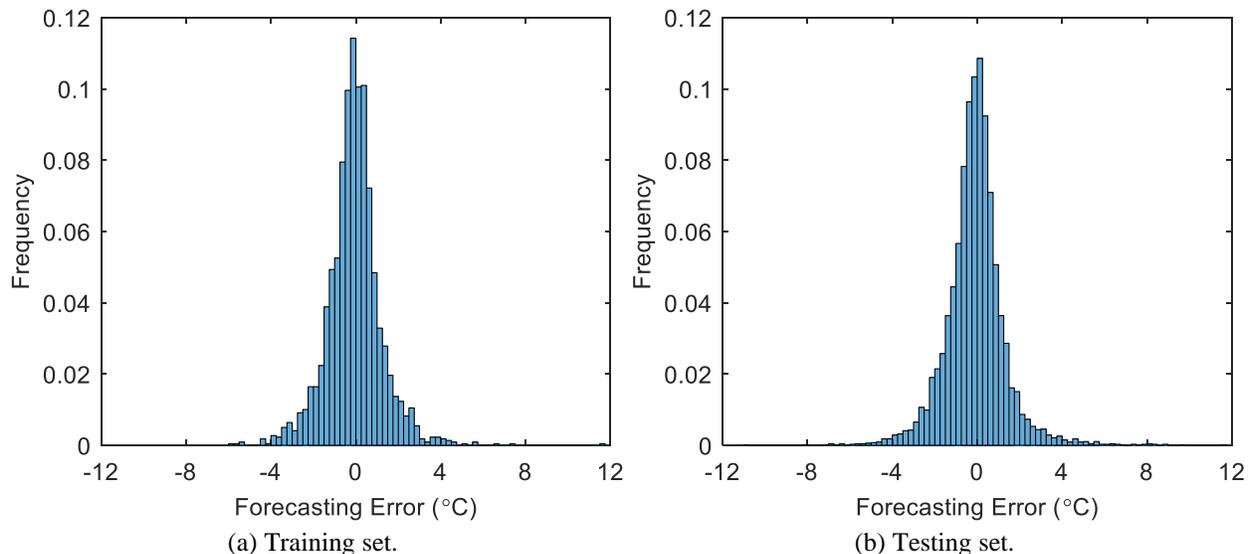


Figure 9. Error distribution of the forecasted temperature at the target building site using ANN.

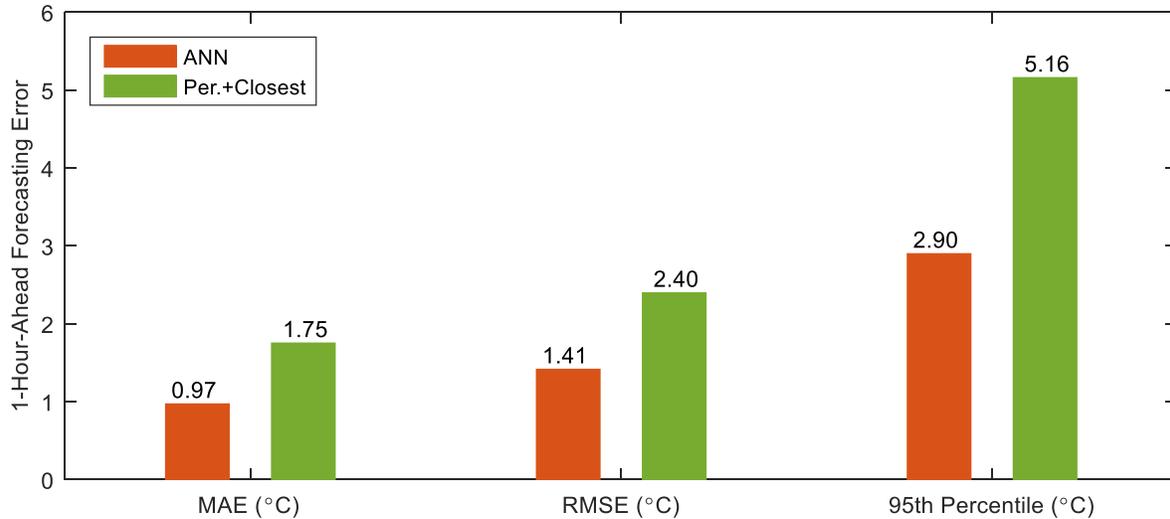


Figure 10. Comparison of 1-hour-ahead temperature forecasting performance using ANN and persistence methods.

In the training set that contains only 91 days, the MAE for 1-hour-ahead temperature forecasts is 0.92°C, with 95% of the absolute forecasting errors less than 2.63°C. In the testing set that contains 274 days, the MAE is 0.97°C, with 95% of the absolute forecasting errors less than 2.90°C. Compared to the forecasting performance for “DNVCC”, the MAE, RMSE, and 95th percentile value only slightly increase when using the same model to forecast the temperature at a different site with significantly fewer training data. The ANN-based method again significantly outperforms the persistence method using the closest station’s data. The MAE is reduced by 45%, the RMSE by 41%, and the 95th percentile value by 44%; therefore, the developed SVR- and ANN-based methods can be generalized to accurately forecast the temperature at different locations using a limited amount of training data. In the future, the scalability of the developed machine-learning-based site-specific weather forecasting methods to different climate zones will be further explored.

Conclusion

This paper presents machine-learning-driven methods to forecast site-specific weather conditions for individual building sites using readily available weather station data. SVR and ANN are adopted to learn the spatiotemporal correlations between the historical weather measurements at nearby weather stations and weather conditions at the building site. In this paper, 1-hour-ahead site-specific temperature forecasting has been used as an example to demonstrate the effectiveness of the developed machine-learning-based forecasting methods. Validation results using 1-year actual weather measurement data from the Denver metro area have shown that the developed SVR- and ANN-based methods can accurately predict the temperature at the target building site 1 hour ahead, with MAE less than 0.72°C and 48% improvement over the persistence method. Future work includes: 1) expanding the developed machine-learning-driven forecasting framework to provide site-specific predictions for other weather variables that are important for building controls, such as relative humidity and wind speed; 2) evaluating the scalability of the developed machine-learning-based site-specific weather forecasting methods to different climate zones; and 3) incorporating other machine learning algorithms to further improve the forecasting accuracy.

Acknowledgement

The authors would like to thank Earth Networks for providing historical weather station data used in this study and Mr. Mark Hoekzema and Mr. Randy Smith for their valuable feedback.

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Building Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

- Al-Yahyai, S., Y. Charabi, and A. Gastli. 2010. "Review of the use of numerical weather prediction (NWP) models for wind energy assessment." *Renewable and Sustainable Energy Reviews* 14 (9): 3192-3198.
- Bhandari, M., S. Shrestha, and J. New. 2012. "Evaluation of weather datasets for building energy simulation." *Energy and Buildings* 49: 109-118.
- Bishop, C. M. 1995. *Neural networks for pattern recognition*. Oxford university press.
- Burges, C. J.C. 1998. "A tutorial on support vector machines for pattern recognition." *Data mining and knowledge discovery* 2 (2): 121-167.
- Cheung, W., J. Zhang, A. R. Florita, B.-M. Hodge, S. Lu, H. F. Hamann, Q. Sun, and B. Lehman. 2015. "Ensemble Solar Forecasting Statistical Quantification and Sensitivity Analysis." in the 5th International Workshop on Integration of Solar Power into Power Systems. Brussels, Belgium.
- Deb, C., F. Zhang, J. Yang, S. E. Lee, and K. W. Shah. 2017. "A review on time series forecasting techniques for building energy consumption." *Renewable and Sustainable Energy Reviews* 74: 902-924.
- Dobbs, A., T. Elgindy, B.-M. Hodge, and A. R. Florita. 2017. "Short-Term Solar Forecasting Performance of Popular Machine Learning Algorithms." in the 7th International Workshop on the Integration of Solar Power into Power Systems. Berlin, Germany.
- EIA (U.S. Energy Information Administration). 2020. *Monthly Energy Review*. Washington, DC: EIA. <https://www.eia.gov/totalenergy/data/monthly/>
- Fikru, Mahelet G., and Luis Gautier. "The impact of weather variation on energy consumption in residential houses." *Applied Energy* 144 (2015): 19-30.

- Florita, A. R., and G. P. Henze. 2009. "Comparison of short-term weather forecasting models for model predictive control." *HVAC&R Research*, 15 (5): 835-853.
- Freedman, J. M., J. Manobianco, J. Schroeder, B. Ancell, K. Brewster, S. Basu, V. Banunarayanan, B.-M. Hodge, and I. Flores. 2014. "The Wind Forecast Improvement Project (WFIP): A public/private partnership for improving short term wind energy forecasts and quantifying the benefits of utility operations. The Southern Study Area, Final Report." Albany, NY: AWS Truepower, LLC.
- Jiang, H., and Y. Zhang. 2016. "Short-term distribution system state forecast based on optimal synchrophasor sensor placement and extreme learning machine." in *IEEE Power and Energy Society General Meeting*. Boston, MA.
- Jin, X., K. Baker, D. Christensen, and S. Isley. 2017. "ForeseeTM: A user-centric home energy management system for energy efficiency and demand response." *Applied Energy* 205: 1583–1595.
- Kim, Y.-J., E. Fuentes, and L. K. Norford. 2015. "Experimental study of grid frequency regulation ancillary service of a variable speed heat pump." *IEEE Transactions on Power Systems* 31 (4): 3090-3099.
- Kwak, Y., D. Seo, C. Jang, and J.-H. Huh. 2013. "Feasibility study on a novel methodology for short-term real-time energy demand prediction using weather forecasting data." *Energy and Buildings* 57: 250-260.
- Lazos, D., A. B. Sproul, and M. Kay. 2014. "Optimisation of energy management in commercial buildings with weather forecasting inputs: A review." *Renewable and Sustainable Energy Reviews* 39: 587-603.
- Li, X., and W. Jin. 2014. "Review of building energy modeling for control and operation." *Renewable and Sustainable Energy Reviews* 37: 517-537.
- Lin, Y., P. Barooah, S. Meyn, and T. Middelkoop. 2015. "Experimental evaluation of frequency regulation from commercial building HVAC systems." *IEEE Transactions on Smart Grid* 6 (2): 776-783.
- Mathiesen, P., and J. Kleissl. 2011. "Evaluation of numerical weather prediction for intra-day solar forecasting in the continental United States." *Solar Energy* 85 (5): 967-977.
- NOAA National Centers for Environmental Information. 2001. Integrated Surface Dataset.
- Ren, Y., P. N. Suganthan, and N. Srikanth. 2015. "Ensemble methods for wind and solar power forecasting—A state-of-the-art review." *Renewable and Sustainable Energy Reviews* 50: 82-91.
- Smola, A. J., and B. Schölkopf. 2004. "A tutorial on support vector regression." *Statistics and computing* 14 (3): 199-222.

- Thieblemont, H, F. Haghghat, R. Ooka, and A. Moreau. 2017. "Predictive control strategies based on weather forecast in buildings with energy storage system: A review of the state-of-the art." *Energy and Buildings* 153: 485-500.
- Utkarsh, K., F. Ding, C. Zhao, H. Padullaparti, and X. Jin. 2020. "A model-predictive hierarchical-control framework for aggregating residential DERs to provide grid regulation services." the 2020 IEEE Conference on Innovative Smart Grid Technologies. Washington, DC.
- Voyant, C., G. Notton, S. Kalogirou, M. L. Nivet, C. Paoli, F. Motte, and A. Fouilloy. 2017. "Machine learning methods for solar radiation forecasting: A review." *Renewable Energy* 105: 569-582.
- Wilcox, S., and W. Marion. 2008. "Users manual for TMY3 data sets." NREL/TP-581-43156. Golden, CO: National Renewable Energy Laboratory.
- Xie, Y., M. Sengupta, and J. Dudhia. 2016. "A Fast All-sky Radiation Model for Solar applications (FARMS): Algorithm and performance evaluation." *Solar Energy* 135: 435-445.
- Xu, L., S. Wang, and R. Tang. 2019. "Probabilistic load forecasting for buildings considering weather forecasting uncertainty and uncertain peak load." *Applied Energy* 237: 180-195.
- Yoon, J. H., R. Baldick, and A. Novoselac. 2014. "Dynamic demand response controller based on real-time retail price for residential buildings." *IEEE Transactions on Smart Grid* 5 (1): 121-129.
- Zhao, P., G. P. Henze, S. Plamp, and V. J. Cushing. 2013. "Evaluation of commercial building HVAC systems as frequency regulation providers." *Energy and Buildings* 67: 225-235.
- Zhou, Z., F. Zhao, and J. Wang. 2011. "Agent-based electricity market simulation with demand response from commercial buildings." *IEEE Transactions on Smart Grid* 2 (4): 580-588.
- Zhu, X., J. Wang, N. Lu, N. Samaan, R. Huang, and X. Ke. 2018. "A hierarchical VLSM-based demand response strategy for coordinative voltage control between transmission and distribution systems." *IEEE Transactions on Smart Grid* 10 (5): 4838-4847.