Prediction of HVAC System Parameters Using Deep Learning

Sirine Maalej[1], Zoubeir Lafhaj[2], Jean Yim[3], Pascal Yim[2], Colin Noort[1]
[1] British Columbia Institute of Technology, Burnaby, BC, Canada
[2] Centrale Lille, Lille, France
[3] Brain Analytics Technologies, Lille, France

Abstract

Heating, ventilation, and air conditioning (HVAC) systems consume between 10 – 20% of developed countries' energy annually. Up to 30% of this energy is often wasted due to mismanagement or improper control strategies. In order to overcome this issue and optimize energy consumption, this paper proposes a predictive modeling technique to effectively forecast HVAC system parameters using machine and deep learning models. A case study of an air handling unit (AHU) at the British Columbia Institute of Technology (BCIT, Canada) is used to test and confirm the results of this research. Five models were applied to predict the supply air temperature. Each model was compared with actual supply air temperature and its accuracy was explored. The results reveal that all investigated models were successful in predicting the supply air temperature, and that the combination of a Convolutional Neural Network (CNN) and a Long Short-Term Memory (LSTM) models has obtained the highest accuracy.

Introduction

Global Energy Consumption

Over the last few decades, the world has seen a substantial increase in its energy consumption. The International Energy Agency has estimated that over a period of two decades primary energy use has grown by 49%, and CO_2 emissions have increased by 43%. This translates to a global average annual increase of about 2% and 1.8% respectively (Perez-Lombard, Ortiz, & Pout, 2008). Discussions regarding this increased energy consumption have brought to light connections between global warming, climate change, increased pollution, and the growth in global energy consumption. With worldwide population growth expected to increase in coming years, and worldwide energy demands expected to continue to grow, reducing energy consumption worldwide is a topic deserving of further inquiry and analysis.

The building sector sits among the top worldwide energy consumers, surpassed only by the food, and transportation industries. The operation of heating, ventilation, and air conditioning (HVAC) systems accounts for about 50% of the energy consumed in the building sector worldwide (Moran, et al., 2017). This translates to between 10 – 20% of the total energy consumption in developed countries (Perez-Lombard, Ortiz, & Pout, 2008; Ramesh, Prakash, & Shukla, 2010; Verbert, Babuska, & De Schutter, 2017) Additionally, degradation and mismanagement of these HVAC systems, and improper control strategies can often lead to 30% of this energy being wasted (Verbert, Babuska, & De Schutter, 2017). Improved control of HVAC systems, through data-driven modeling and prediction methods, has the potential to significantly reduce the overall energy consumption of these systems.

Existing Control & Prediction Methods

Most modern building management systems monitor and record HVAC system data (temperature, pressure, humidity, etc.) using software and sensors connected to various system components (dampers, fans, cooling, heating coils, etc.). Over time, this leads to a substantial amount of historical HVAC system data being available. This data can be used to develop data-driven models using predictive modeling techniques. The models would make the simulation and optimization of HVAC systems possible and contribute to the energy reduction in buildings.

Various data-driven approaches have been explored and established as effective at overcoming the shortcomings of analytical methods. For example, Fan et al. (2017) showed that these data-driven methods are effective at cooling load predictions in buildings, while Korolija et al. (2013) showed them to be effective in energy consumption modeling. Many of these data-driven methods rely on artificial intelligence (AI) and machine learning (ML) to do most of the data processing involved. AI has proven to be proficient in modeling complex systems such as HVAC systems, due to its information processing and real-time decision-making capabilities (Manic, Amarasinghe, Rodriguez-Andina, & Rieger, 2016). Additionally, the heterogeneous nature of an HVAC system's datasets makes AI a perfect tool for data-

driven HVAC system modeling applications. AI and ML have already been shown to be capable of predicting future HVAC system states by identifying patterns and learning from past system datasets.

Although many studies in the literature reviewed proved good results in indoor air temperature predictions using classical machine learning models, few works investigated the deep learning prediction performance. In recent years, deep learning (DL) techniques have emerged as subsets of ML and AI. Where ML achieves AI through established algorithms and training, DL is differentiated by its use of artificial neural networks (ANN). DL has been used in a wide variety of analytical tasks, such as image classification, and speech recognition (Fan et al., 2019). Because of DL's more complex architecture, it can describe more accurately heterogeneous and non-linear relationships seen in datasets. As such, DL has the potential of being more suited for applications in data-driven HVAC system predictive modeling.

This paper looks at a specific application case of an Air Handling Unit (AHU) in an HVAC system at the British Columbia Institute of Technology. Artificial intelligence, in combination with various machine and deep learning models, is used to develop predictive models of the preheated air serving 7 rooms in this AHU. This paper will first break down the specific application case, provide a brief background on AI and deep learning, discuss the methodology employed in this specific application case, and finally outline the results.

Methodology

Predictive modeling has been widely accepted as an effective tool by both academic researchers and building professionals. Predictive modeling is a mathematical process that aims to predict future events based on patterns present in past historical datasets. Predictive algorithms present an opportunity for the building sector to reduce failures, maintenance costs and energy consumption while improving the comfort of the occupants. The larger the dataset, the more accurate the prediction will be. Although this data could theoretically be analyzed manually by a human operator, this would be incredibly inefficient. A much more effective and efficient tool which can be used to analyze this data is artificial intelligence, specifically machine learning and deep learning.

Artificial intelligence and machine learning encompass a wide range of algorithms and techniques (Manic, Amarasinghe, Rodriguez-Andina, & Rieger, 2016). Machine learning algorithms such as Linear Regression, Logistic Regression, and Dimensionality Reduction Algorithms are well established and have been effectively used for predicting future events from past datasets (Manic, Amarasinghe, Rodriguez-Andina, & Rieger, 2016). Implementing a predictive approach presents several challenges, such as the availability of valuable data and the accuracy of predictive algorithms. In this paper, learning algorithms belonging to classical ML and DL are compared. The main purpose of the study is to identify an algorithm that accurately predicts the supply air temperature of an AHU unit and that can be used to build a digital twin for energy optimization in future work.

This study focuses on predicting the supply air temperature that allows the adjustment of the primary delivery temperature in an HVAC system. It is a critical parameter in controlling the AHU equipment and the Variable Air Volumes (VAVs) in the served zones.

Predicting the supply air temperature is a time series problem. Therefore, the selected models in this study are chosen based on their proven efficacy in time series scenarios.

Machine Learning

Tree-based models are built by recursively partitioning the considered observations following some criteria. The criteria are created by comparing all possible splits in the data and choosing the one that offers the highest mean squared-error (MSE) reduction in Variance of child nodes.

Ensemble methods (e.g. Bagging and Boosting) combine several decision trees predictors to produce better performance and build stronger predictive models. The bagging technique uses average predictions from different trees constructed by creating several data subsets. While Boosting is an iterative technique that fits consecutive trees constructed on random samples. At every step simple models are fit to the data and data is analyzed for errors. The effectiveness of models based on ensemble methods (such as Random Forest, XGBoost, and Extra Tree) in forecasting time series data sets and HVAC systems, in particular, was examined by a few studies (Alawadi, et al., 2020; Ampomah, Qin, & Nyame, 2020; Manivannan, Behzad, & Rinaldi, 2017).

In this study, we selected one bagging model and one boosting model: Extra Tree and XGBoost. The Extra Tree is very similar to Random Forest except that the decision rule during tree construction is randomly selected. Alawadi et al. (Alawadi, et al., 2020) found that the Extra Tree model is the most accurate and performant model for HVAC temperature prediction. Thus, the Extra Tree regressor was selected as the bagging model in this study. XGBoost is built based on parallel tree boosting and is designed for high efficiency. The model has been widely used for supervised learning problems including HVAC analysis (Chakraborty & Elzarka, 2019; Divina, Torres, Vela, & Noguera, 2019).

Deep Learning

Deep learning is a powerful tool and is quickly gaining popularity, especially in the field of data analytics. In recent

years, Deep Learning methodologies outperformed traditional learning models. It has taken predictive modeling one step further and continues to revolutionize numerous fields through its strength in predictive modeling applications. Common applications of deep learning include image recognition, speech recognition, and computer vision applications (Manic, Amarasinghe, Rodriguez-Andina, & Rieger, 2016). Deep learning refers to a collection of machine learning algorithms that take a deeper approach to discovering knowledge (Fan, Xiao, & Zhao, 2017). It transforms the input data multiple times in either a linear or non-linear manner before deriving an output. Conversely, most conventional machine learning algorithms tend only to translate the input data once or twice before producing an output, resulting in a shallower model incapable of modeling some more complex systems.

In this section, the adopted deep learning models for this study are briefly presented: Convolutional Neural Networks (CNN), and Long Short Term Memory (LSTM). These algorithms are selected in this paper because they are relevant to the considered application and the specific case study discussed later.

Long Short-Term Memory (LSTM)

Long short-term memory (LSTM) networks are a type of Recurrent Neural Network (RNN) capable of learning order dependence in sequence prediction problems. Unlike standard feedforward neural networks, LSTM has feedback connections. LSTMs are typically used in generative models, especially in natural language processing tasks. However, LSTM networks are very well suited to classifying, processing, and making predictions based on time series data (Sülo, Keskin, Dogan, & Brown) & (Xu, Chen, Wang, Guo, & Yuan, 2019). These models can capture features and remember them over time and therefore they were considered in this investigation.

Convolutional Neural Networks (CNN)

The convolutional neural network (CNN) is a type of artificial neural network which has so far been most used for analyzing images. Although image analysis has been the most common application for CNN's, they can be used for any type of classification problem (DeepLizard, 2017). CNN's have been designed to learn spatial hierarchies of features automatically and adaptively, from low-to-high-level patterns (Yamashita, Nishio, Gian Do Kinh, & Togashi, 2018). A CNN is typically composed of three types of layers: convolution, pooling, and fully connected layer. The convolution and pooling layers perform the feature extraction, while the fully connected layer maps the extracted features into a final output, such as a classification (Yamashita, Nishio, Gian Do Kinh, & Togashi, 2018). CNN has been proven effective in time series problems (Krizhevsky, Sutskever, & Hinton, 2012) and is thus considered in this study.

Case Study - BCIT AHU-07

To reduce the ecological footprint of the British Columbia Institute of Technology's Burnaby campus, various initiatives have been put in place. Among these initiatives is a larger goal of reducing the greenhouse gas emissions at BCIT by 33% by the year 2023 (British Columbia Institute of Technology, 2021). The purpose of this case study is to contribute in a meaningful way to meeting this goal.

HVAC System Structure and Components

The British Columbia Institute of Technology, British Columbia, Canada (BCIT) is a post-secondary institute with a total of five campuses and 62 buildings. The HVAC system being looked at is owned and operated by BCIT - The Burnaby Campus. In this study, a single Air Handling Unit (AHU-07) located on the fourth floor of the North-East 1 (NE01) building of BCIT is analyzed. The building mainly consists of offices, classrooms, washrooms, and mechanical rooms, with a gross floor area of 20,076.88 m^2. The building zones temperature is controlled by a forced-air heating system. A total of 10 AHUs serve the four floors. AHU-07 serves 7 zones on the fourth floor: 6 classrooms and 1 office.

Figure 1 shows the schematic diagram of AHU-07 with the monitored variables. Outdoor air is mixed with air returning from the seven served zones. The composition of the mixture is controlled by three dampers regulating the percentages of air exhausted from the system, entering the system, and recirculating in the system. The mixed air is cooled or heated when passing through the cooling and heating coils. The heating coil uses hot water and is connected to the boiler. The amount of water flowing through the heating coil is controlled by a valve. Pressure in the AHU system is maintained by the return and supply fans. A description of the monitored variables and the role they play in the overall system is provided in Table 1. These variables are the input variables for the predictive models as will be detailed later. Detailed technical information of the monitored equipment can be found through ASHRAE & Consulting-Specifying Engineer (2019). The study being conducted was designed to investigate the impacts of AHU parameters on the supply air temperature at temperature

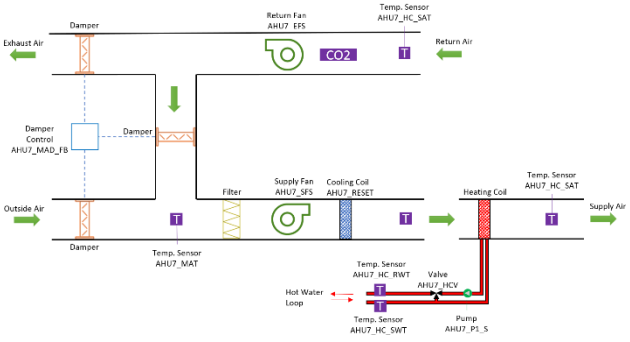

Figure 1 Schematic Diagram of AHU System

sensor AHU7_HC_SAT. The operation of the AHU system is typically impacted by the outside weather condition. The latter was represented by the Outside Air Temperature variable (AHU7_OAT).

Table 1 AHU7 Parameters description

Variable label	Description	Unit
AHU7_RESET	Cooling Coil Valve	% open
AHU7_HCV	Hot Water Valve	% open
AHU7_HC_SWT	Hot Water Supply Temperature	°C
AHU7_HC_RWT	Hot Water Return Temperature	°C
AHU7_MAD_FB	Dampers position	% open
AHU7_HC_SAT	Supply Air Temperature	°C
AHU7_MAT	Mixed Air Temperature	°C
AHU7_RAT	Return Air Temperature	°C
AHU7_SF_SPD	Supply Fan Speed	% max speed
AHU7_EF_SPD	Return Fan Speed	% max speed
AHU7_OAT	Outside Air Temperature	°C
AHU7_P1_S	Hot Water Pump	A

Data Collection

Over the last few years, the facilities division at BCIT has established a modern and sophisticated building management system (BMS). The BMS interface and dedicated archival equipment have enabled the precise monitoring and control of the HVAC systems at BCIT. The data from the various system sensors is recorded, archived and can be extracted.

Minimal campus occupancy is being observed in the years 2020 and 2021 due to the COVID-19 pandemic. Thus, data from these two years was not considered to be useful for this predictive modeling application. Therefore, the data used for this study is from the period of January 2016 – December 2019. This represents a total of four years and about 360,000 data points per system component. Each row of this data corresponds directly to the state of the system, and each row is separated by a 15-minute time interval. Using ML and DL data mining techniques, patterns can be found, and system trends and correlations can be derived from this data. Through closer inspection of these trends, it is possible to eliminate some redundant system components when developing models of the system.

Application

The application framework of the prediction models to the BCIT case study is composed of three steps: Data collection and preparation, model development, and performance evaluation. The data collection step has been discussed above. The performance evaluation is detailed in the next section. The data preparation and algorithm development are discussed and detailed below.

Data Preparation

Data preparation refers to the process of constructing relevant features and data sets as model inputs by cleaning and transforming the raw data. In the presented case study, all equipment and sensor data were extracted from the Building Management System (BMS). In total there were 19 features and this number was reduced to 12 features that are meaningful to the Supply Air Temperature (SAT) prediction (see table 1). Return Air Temperature and Return fan speed were dropped due to high correlations with other features.

The considered case study is a campus where the HVAC system is working according to a schedule. The latter is based on operational rules that vary depending on the day of the week (weekday/weekend) and the hour of the day in addition to Holidays. The SAT was affected by this operation schedule. Figure 2 illustrates the variation of the SAT across the hours of the day for the four considered years. A variance in SAT across the 24h can be observed. The figure shows a constant SAT between 11:00 pm and 4:00 am. That is when the HVAC system is off. The system is programmed for an Optimal Start at 4:00 am during weekdays and that extends over two hours.

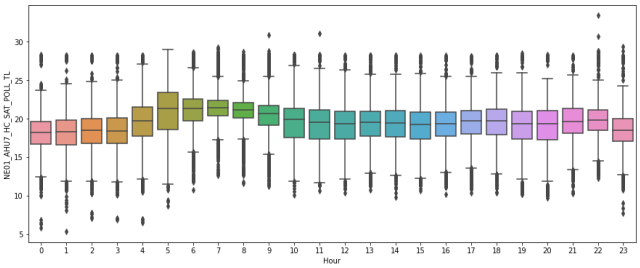

Figure 2 The change in Supply Air Temperature per hour

In an attempt to capture the time dependency in the model, the parameters that represent the seasonality including the month of the year, the hour of the day, and the day of the week (weekday/weekend) are also utilized as inputs. Table 2 summarizes the schedule-related variables defined as input features for the prediction models.

Table 2 Schedule related features

Input variable	Description	N° of Inputs
Month of the year	Month of the recorded data	12
Day of week	Weekday vs Weekend	2
Hour of the day	Day time, evening, night	3

The data was split into three subsets to ensure the accuracy of the models. For all five architectures, the first three years were used as the training data and the last year was used as testing data. A validation set was needed to ensure that the models are unbiased and to select the best iteration of the models based on the mean squared error (MSE).

Historical HVAC features are known to impact the SAT of the AHU system. Hence a sliding window algorithm was

employed for the sampling process to efficiently train the considered models.

The first training sample of the input variable "x" is constructed as a matrix of dimension $W \times F$, where W is the window length (W=n*15 min, if n=4 then W=1h) and F is the number of features (26 features). Then the window slides by 1 step (15 min) and the readings from the time step i+1 make the second sample and so on (see figure 3). The length of this sliding window determines the influence of past events in the prediction (e.g. 15 min, 30 min, 1h, etc.).

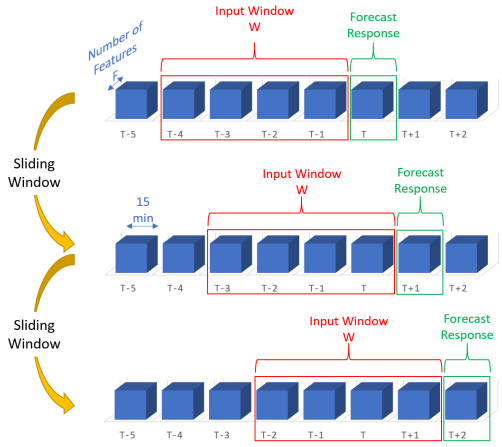

Figure 3 Sliding windows sampling process

Considering 15 min step, there is an overlap between the sliding windows, and a reading from a single time step belongs to multiple windows. If a total of N rows of data was considered, the final training dataset shape is a matrix [N, n, 26]. For each sample (window of length W), the model predicts the supply air temperature for T+i (variable "y") where i is the horizon response in the feature. Each model was developed to effectively predict the temperature of supply air at two different forecast responses: T + 1 and T+2 (15 and 30-min in the future).

A number of 128 consecutive samples generated by the sliding windows technique are placed in a group referred to as the batch. The learning takes place one batch at a time. A model is always trained only on the current batch.

Predictive Models Development

To develop and select the most effective predictive model, various algorithms including ExtraTree, XGboost, LSTM, and CNN are employed. These algorithms were implemented on an open-source Python™-based dedicated machine learning package. Python 3.9 programming language and the Keras deep learning framework were used.

To ensure that the optimal model is selected, we evaluated the following five models:

- Model1 – Extra Tree: Used classical Extra Tree Machine learning model. This model was selected as the bagging method-based tree model.
- Model2 - XGBoost: Used XGBoost model to predict the Supply Air Temperature as the boosting method-based tree model.
- Model3- LSTM: Used an LSTM layer and a dense layer to build the model. This model has shown better performances over other RNN models tested in this application. It allows better control of which features will be useful for the predictions and avoids the vanishing gradient problem when larger sequences will be considered to build a digital twin of the system.
- Model4 - CNN: Used two CONV1D filters and 1 dense layer
- Model5 - CNN+LSTM: Considering that CNN can learn local features and can reduce the number of parameters and that LSTM can capture variation in long/short term dependencies, a hybrid method based on LSTM and CNN is used to benefit from the potentials of both models. In this proposed model the LSTM layer supports sequential dependencies in the input data, while the CNN layers inform the model by extracting the effective features using multiple filters. Two CNN layers and one LSTM layer are used for building the model. The loss function is MSE, the optimizer is Stochastic Gradient Descent (SGD), and the metrics is MAE.

Results and Discussion

In this section, predictions obtained by implementing the models' framework developed above on the BCIT AHU system are discussed. The purpose of this investigation is first, to assess which technique is more suitable in predicting HVAC parameters and second, to select the optimal model settings to predict the Supply Air Temperature.

The selected prediction model will be used to adjust the equipment setting of an AHU in real-time without human involvement and preset rules. The equipment response time is generally very short and a 15-min to 30-min ahead forecast should be sufficient.

Since the size of the time delay of input variables (sliding window) and the forecast horizon (T+1, T+2, etc.) have a significant impact on the prediction model accuracy, different combinations of the two parameters are set to analyze their influence. The lookback time delay was parameterized and sliding windows of time-steps 1, 2, 4, 6, and 8 (8 windows = 2 hours) were considered. The models were used to predict the SAT for 15 and 30 minutes ahead forecast.

A simple regression algorithm was also tested. The model assumes that the SAT at T+1 and T+2 is equal to the temperature at T for the 15 and 30 minutes ahead forecast.

Figures 4 & 5 show the performance of the different architectures on the supply air temperature test data. They show a snapshot of the prediction result of 1-year from January to December 2019 with 15- and 30-minute ahead forecast and 1-h lookback window. It is noticed that the overall prediction performance of all models is comparable.

Overall, all models are able to predict the 15min-ahead and the 30min- ahead Supply Air Temperature within ±0.5°C and ±0.8°C error respectively.

Although the considered ML and DL models are able to predict the SAT with an acceptable error, the authors examined the achieved accuracy differences. The model sensitivity is important because this investigation is part of a bigger study aiming to build a digital twin of a complex HVAC system with all served zones and with extended future time predictions.

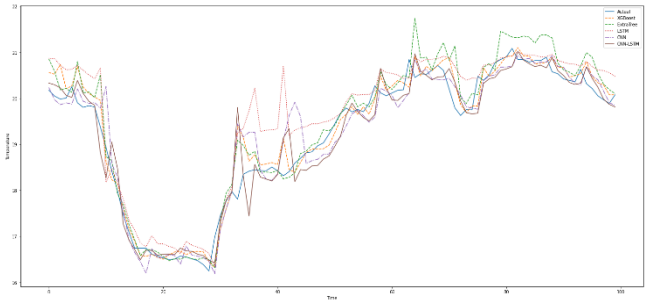

Figure 4 - 15min ahead and 1-h sliding window predicted temperature for the five Models

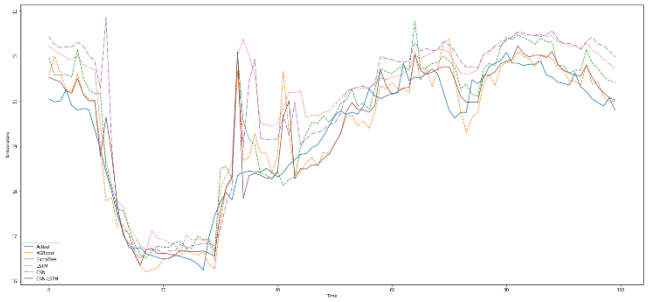

Figure 5 - 30min ahead and 1-h sliding window predicted temperature for the five Models

To investigate what model outperforms the others, two metrics were considered: the Mean Absolute Error (MAE) and the Squared Mean Squared Error (RMSE). The Mean Absolute Percentage Error (MAPE)was exculed because it is scale-sensitive and penalizes positive errors more than negative errors.

$$MAE = \frac{\sum_{i=1}^{n} |\widehat{y_i} - y_i|}{n} \quad (1)$$

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n} (\widehat{y_i} - y_i)^2}{n}} \quad (2)$$

where y_i and $\widehat{y_i}$ are the actual and predicted values, and n is the total number of observations. MAE and RMSE are commonly used on time series predictions (Galicia, Talavera-Llames, Troncoso, Koprinska, & Martínez-Álvarez, 2019; Fan, Sun, Zhao, Song, & Wang, 2019). They provide the user a straightforward way to quantify the prediction error because the forecast error is expressed in the same units of the variables to be predicted. MAE measures the average magnitude of the errors in the prediction set. RMSE gives a relatively high weight to the large errors since errors are squared before they are averaged. The lower value of MAE and RMSE implies higher accuracy of the considered model.

MAE and RMSE obtained on the validation and test sets for different historical sliding windows are detailed in Appendix A1 and A2. For 15-min ahead prediction, the lowest errors were observed for a sliding window of 1h. the lookback window size does not have a significant impact on the prediction except for when the lookback window is set to 15-min. A 15-min lookback window shows less good predictions for all the studied models. For the 30-min ahead prediction, a lookback window of 1.5h resulted in the optimal temperature forecast across models.

The simple regression model, ML and DL models' best performances at 1h and 1.5h lookback windows for 15-min and 30-min ahead predictions, respectively, are shown in Tables 3 & 4.

The simple regression model has the lowest accuracy compared to all other models. This proves the need for more elaborate ML and DL models that learn time dependency and how the SAT evolves over time.

The MAE for the five ML and DL models' predictions is less than 0.6 °C for 15-min ahead prediction and less than 0.8°C for 30-min ahead prediction. Hence, the approach, data, and models considered for this investigation are proven to provide a reliable prediction of the supply air temperature, considering that the sensors' overall uncertainty is around 0.3 °C. The CNN+LSTM relatively outperforms the other models on test data with an improved prediction of 13% to 29% and 15% to 24% on 15-min and 30-min ahead forecast targets respectively.

Looking at the errors' values, it is noticed that RMSE errors are higher than MAE errors for both the validation and test sets for all the studied models. This refers to the existence of large errors values in the data set. A closer look at the large errors showed that the models are less accurate in predicting sudden changes in the predicted temperature. This will be investigated in future work.

Table 3 - 15 min Prediction Error [℃] with a sliding window of 4 (1h)

Model	Validation Set		Test Set	
	RMSE	MAE	RMSE	MAE
Simple Regression	-	-	1.1043	0.5131
ExtraTree	0.6020	0.3084	0.8187	0.4677
Xgboost	0.5838	0.3189	0.8789	0.4895
LSTM	0.6763	0.4113	0.9340	0.6028
CNN	0.6688	0.3686	0.8177	0.4715
CNN + LSTM	0.5685	0.3119	**0.7287**	**0.4058**

Table 4 - 30 min Prediction Error [℃] with a sliding window of 6 (1.5h)

Model	Validation Set		Test Set	
	RMSE	MAE	RMSE	MAE
Simple Regression	-	-	1.7637	0.9071
ExtraTree	0.8283	0.4726	1.2100	0.7686
Xgboost	0.8755	0.5255	1.2725	0.7962
LSTM	0.9669	0.6125	1.1875	0.8007
CNN	0.9574	0.5761	1.1550	0.7390
CNN + LSTM	0.8620	0.5049	**1.0537**	**0.6431**

Conclusion

When designing smart buildings, energy optimization and user comfort are key performance indicators. Hence, it is important to develop accurate and reliable HVAC systems that optimize energy and improve comfort. The primary purpose of this study was to propose a reliable data-driven model that represents AHU systems. Five algorithms were implemented, over a real data set, to predict the supply air temperature at a BCIT building.

The results show that 15-min and 30-min ahead Supply Air Temperature of an AHU system can be accurately predicted using a designed deep learning model. Hence the study provides the framework to build a predictive data-based model – a digital twin that would include multiple AHU systems and their served rooms in a given building. The digital twin model would be used to simulate and generate the optimal settings of the HVAC systems that would imply a reduction in energy consumption and an improvement of the occupant comfort level. This will be addressed in future work.

Acknowledgment

This work was supported by the Seed Fund and the Institute Research Fund granted by the British Columbia Institute of Technology.

References

Alawadi, S., Mera, D., Fernández-Delgado, M., Alkhabbas, F., Olsson, C. M., & Davidsson, P. (2020). A comparison of machine learning algorithms for forecasting indoor temperature in smart buildings. *Energy Systems*. Retrieved from https://doi.org/10.1007/s12667-020-00376-x

Ampomah, E. K., Qin, Z., & Nyame, G. (2020). Evaluation of Tree-Based Ensemble Machine Learning Models in Predicting Stock Price Direction of Movement. *Information, 11*(332).

ASHRAE. (n.d.). Retrieved 09 21, 2021, from https://www.ashrae.org/technical-resources/standards-and-guidelines

British Columbia Institue of Technology. (2021). *Ecological Footprint*. (British Columbia Institue of Technology) Retrieved August 15, 2021, from https://www.bcit.ca/sustainability/community/ecological-footprint/

British Columbia Institute of Technology. (2021). *Energy*. (British Columbia Institute of Technology) Retrieved August 15, 2021, from https://www.bcit.ca/sustainability/operations/energy/

Chakraborty, D., & Elzarka, H. (2019). Early detection of faults in HVAC systems using an XGBoost model with a dynamic threshold. *Energy & Buildings*(185), 326-344.

Consulting-Specifying Engineer. (2019, January 21). *AHU Types, Components, and Configurations*. (Consulting-Specifying Engineer) Retrieved July 19, 2021, from https://www.csemag.com/articles/ahu-types-components-and-configurations/

DeepLizard. (2017, December 9). *Convolutional Neural Networks (CNNs) Explained*. (YouTube) Retrieved July 19, 2021, from https://www.youtube.com/watch?v=YRhxdVk_sIs

Divina, F., Torres, M. G., Vela, F. G., & Noguera, J. V. (2019). A Comparative Study of Time Series Forecasting Methods for Short Term Electric Energy Consumption Prediction in Smart Buildings. *Energies, 12*(1934).

Fan, C., Sun, Y., Zhao, Y., Song, M., & Wang, J. (2019). Deep learning-based feature engineering methods for improved building. *Applied Energy, 240*, 35-45.

Fan, C., Xiao, F., & Zhao, Y. (2017). A Short-Term Building Cooling Load Prediction Method Using Deep Learning Algorithms. *Applied Energy, I*(195), 222-233.

Galicia, A., Talavera-Llames, R., Troncoso, A., Koprinska, I., & Martínez-Álvarez, F. (2019). Multi-step forecasting for big data time series based on ensemble learning. *Knowl.-Based Syst, 163*, 830-841.

Korolija, I., Zhang, Y., Halburd-Marjanovic, L., & Hanby, V. I. (2013). Regression Models for Predicting UK Office Building Energy Consumption from Heating and Cooling Demands. *Energy and Buildings, I*(59), 214-227.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems 25 (NIPS 2012).*

Manic, M., Amarasinghe, K., Rodriguez-Andina, J., & Rieger, C. G. (2016). *Intelligent Buildings of the Future: Cyberaware, Deep Learning Powered, and Human*

Interacting. Idaho: U.S. Department of Energy National Laboratory.

Manivannan, M., Behzad, N., & Rinaldi, F. (2017). Machine Learning-Based Short-Term Prediction of Air-Conditioning Load through Smart Meter Analytics. *Energies, 10*.

Moran, A., Alonso, S., Prada A., M., Fuertes J., J., Diaz, I., & Dominguez, M. (2017). Analysis of Parallel Process In HVAC Systems Using Deep Autoencoders. In *Communications in Computer and Information Science* (pp. 15-26). Engineering Applications of Neural Networks.

Perez-Lombard, L., Ortiz, J., & Pout, C. (2008). A Review on Building Energy Consumption Information. *Energy and Buildings, I*(40), 394 - 398.

Ramesh, T., Prakash, R., & Shukla, K. (2010). Life Cycle Energy Analysis of Buildings: An Overview. *Energy and Buildings, I*(42), 1592 - 1600.

Sülo, I., Keskin, S. R., Dogan, G., & Brown, T. (n.d.). Energy efficient smart buildings: LSTM neural networks for time series prediction. *2019 International Conference on Deep Learning and Machine Learning in Emerging Applications (Deep-ML)* (pp. 18-22). IEEE (2019).

Verbert, K., Babuska, R., & De Schutter, B. (2017). Combining Knowledge and Historical Data for System-Level Fault Diagnosis of HVAC Systems. *Engineering Applications of Artifical Intelligence, I*(59), 260-273.

Xu, C., Chen, H., Wang, J., Guo, Y., & Yuan, Y. (2019). Improving prediction performance for indoor temperature in public buildings based on a novel deep learning method. *Build. Environ., 148*, 128-135.

Yamashita, R., Nishio, M., Gian Do Kinh, R., & Togashi, K. (2018). Convolutional Neural Networks: An Overview and Application in Radiology. *Insights Into Imagine, I*(9), 611-629.

Appendix

Table A 1 - 15 min Prediction errors with each model

Model	Validation Set		Test Set		Sliding Window Size
	RMSE	MAE	RMSE	MAE	
ExtraTree	0.6094	0.3199	0.8170	0.5078	1
	0.5939	**0.3034**	**0.8000**	**0.4617**	2
	0.6020	0.3084	0.8187	0.4677	4
	0.5993	0.3061	0.8324	0.4716	6
	0.6065	0.3082	0.8229	0.4736	8
Xgboost	0.6263	0.3560	0.9272	0.5470	1
	0.6027	0.3214	**0.8568**	0.4946	2
	0.5838	**0.3189**	0.8789	**0.4895**	4
	0.5890	**0.3177**	0.8626	0.4921	6
	0.5984	0.3212	0.8913	0.5222	8
LSTM	0.7013	0.43871	1.0142	0.6910	1
	0.6949	0.41303	**0.9090**	0.5941	2
	0.6763	0.41132	0.9340	0.6028	4
	0.6775	**0.40472**	0.9447	0.6228	6
	0.6818	**0.40291**	0.9131	**0.5877**	8
CNN	0.66542	0.38899	0.8681	0.5499	1
	0.65754	0.35488	0.7966	0.4629	2
	0.6688	0.3686	0.8177	0.4715	4
	0.6637	0.3604	0.8223	0.4738	6
	0.6740	0.3704	0.8455	0.4900	8
CNN + LSTM	0.6177	0.3554	0.8148	0.4919	1
	0.5903	0.3149	0.7572	**0.4165**	2
	0.5685	0.3119	**0.7287**	**0.4058**	4
	0.5688	**0.3059**	0.7480	0.4208	6
	0.5927	0.3284	0.7383	0.4274	8

Table A 2 - 30 min Prediction Error with each model

Model	Validation Set		Test Set		Sliding Window Size
	RMSE	MAE	RMSE	MAE	
ExtraTree	0.8752	0.5018	1.2943	0.8572	1
	0.8644	0.4881	1.2242	0.7742	2
	0.8517	0.4829	**1.2044**	**0.7674**	4
	0.8283	**0.4726**	1.2100	**0.7686**	6
	0.8506	0.4823	1.2229	0.7779	8
Xgboost	0.9170	0.5660	1.4472	0.9424	1
	0.9009	0.5361	1.3373	0.8172	2
	0.9023	0.5346	**1.1942**	**0.7514**	4
	0.8755	**0.5255**	1.2725	0.7962	
	0.8958	0.5352	1.2367	0.7834	8
LSTM	1.0144	0.6542	1.2368	0.8644	1
	1.0053	0.6359	1.2030	0.8279	2
	0.9931	0.6301	1.1948	0.8106	4
	0.9669	**0.6125**	**1.1875**	**0.8007**	6
	0.9873	0.6186	1.1983	0.8087	8
CNN	0.9904	0.6229	1.2147	0.8183	1
	1.0301	0.6033	**1.1391**	0.7433	2
	0.9871	0.5862	1.1624	0.7501	4
	0.9574	**0.5761**	1.1550	**0.7390**	6
	0.9797	0.5892	1.1592	0.7427	8
CNN + LSTM	0.9326	0.5753	1.1562	0.7550	1
	0.9219	0.5425	1.1537	0.7231	2
	0.8913	0.5253	1.0990	0.6772	4
	0.8620	**0.5049**	**1.0537**	**0.6431**	6
	0.8907	0.5169	1.0888	0.6741	8