A data-driven approach to automatically label BAS points

Karim El mokhtari, elmkarim@ryerson.ca
Ryerson University, Toronto, Canada

J.J. McArthur, jjmcarthur@ryerson.ca
Ryerson University, Toronto, Canada

Abstract
Legacy Building Automation System (BAS) pose a challenge to modernization due to the lack of a common point naming convention across the built environment. Identifying these points is crucial to analyze the behavior of the system, detect faults, and permit the migration to more integrated BAS systems enabled by Digital Twins. The objective of this research is to automatically identify BAS points from their time series data. To do so, time series patterns characterizing each BAS point type are identified and deep machine learning methods are used to create feature maps to discriminate between them. To reconstruct the BAS topology, small clusters of signals (e.g., emanating from the same room) are next identified using machine learning and this process can be extended to ever-larger system clusters until the full system is mapped into Digital Twin. The findings of this initial study have demonstrated significant potential of this approach, with point type classification accuracies between 87 and 100% and clustering accuracies of 88 - 94%.

Keywords: deep learning, CNN, time series, identification, HVAC

1 Introduction
Smart and Ongoing Commissioning (SOCx) is an increasing focus within the building sector, providing facility managers with a wealth of algorithms and tools to monitor and continuously optimize building system performance. There are many applications found in the literature related to fault detection and diagnosis (Li & O'Neill, 2018) or preventive maintenance (Gunay, et al., 2019); both show the impact of using analytics to extend the system lifetime and maintain the occupants' comfort. Energy management and Smart and Continuous Commissioning are also fields where data significantly leverages optimal energy and equipment use.

In order to successfully implement SOCx, a significant volume of data is required to support a wide range of machine learning and analytics. Building Automation Systems (BAS) provide access to a wealth of such data; however, their traditional design of prescriptive controls and alarms offers limited SOCx capabilities. To add this "smart" component, two solutions can be considered, the first and least expensive is to connect the BAS to external analytical platforms and the second is to update or migrate it to a newer system with extended analytical possibilities. While a new BAS can be designed to support such integration, legacy systems provide a unique challenge. While strategies have been developed to extract such data from legacy systems (Misic, et al., 2020), the BAS points must be known, either through detailed system documentation or a consistent nomenclature. Without a clear identification of these points, any further exploitation is rendered impossible. The main objective of this paper is to develop models to identify the key properties and patterns of HVAC signals to allow their classification. Therefore, signals associated to the same cluster (e.g., room or equipment) will be easily identified and their relationships learned. This is particularly useful in the re-creation of the BAS architecture and topology, enabling the application of SOCx analytics on legacy systems and/or the migration to (or expansion with) a new BAS or integrated Smart Building platform. This research directly supports the development of digital twins (DTs) for both legacy and new buildings. For the former, poor documentation of BAS structure is a significant barrier to adoption, while in all building types, the relationships learned from signal clusters will permit early anomaly and fault detection as well as a automated signal recognition and labeling for new and existing equipment. To attain these objectives, we propose a generalizable, automated approach to identify and label BAS points using a convolutional neural network with a grey-box model. First, the convolutional model identifies the BAS points based on their time-series characteristics; once these have been classified, the grey-box models transform control signals into expected time-series data for related points, thus permitting system clustering.

This paper is organized as follows: Section 2 provides an overview of the methods used in time series classification. Section 3 presents the methodology to classify both sensors and control time series and then uses these outputs to construct the BAS topology. Section 4 discusses the results of the experiments that were performed on data from an academic building in Toronto. Section 5 summarizes the findings, presents the DT integration, and discusses avenues of promising future research.

2 Literature review

There has been significant research interest in Time Series Classification (TSC). Generally, TSC algorithms identify characteristic patterns, which are then used as features for classification. Three dominant approaches for TSC are distance-based, statistical, and deep learning.

Distance-based approaches measure the distance between time-concurrent points between two time series. The KNN-DTW algorithm (Bagnall, et al., 2017) uses an elastic distance measure before applying the nearest neighbor method for classification. Statistical algorithms create features from time series such as the mean and the standard deviation, then train a classifier on those features (Nanopoulos, et al., 2001). Interval feature methods, such as Time Series Forest (TSF) (Deng, et al., 2013), are an extension of this approach, using random temporal intervals to develop a feature vector for each time series. These calculated features (mean, standard deviation, and slope) are then used to predict the time series class. Such methods can provide insight on the time series regions that contain characteristic patterns.

Deep learning methods have attracted significant focus in recent years (Fawaz, et al., 2019) and consist of several types. Recurrent neural networks and particularly the Long Short-Term Memory (LSTM) are effective in locating patterns across long sequences of data. The LSTM was used in many TSC applications such as predicting diagnoses from medical observations (Lipton, et al., 2015) or classifying discrepancies in financial data (Higdon, et al., 2019). Other augmented variants of the LSTM were proposed, namely the one associated with a Fully Connected Network (FCN) (Karim, et al., 2019) and the attention-based LSTM (Karim, et al., 2018) and benchmarked in classifying 85 time series datasets (Dau, et al., 2018). Convolutional Neural Networks (CNNs) were initially created in computer vision and achieved near-human accuracy in image classification (Szegedy, et al., 2015), but they were also efficient in TSC as they can automatically learn new features to classify 1D-data. (Fawaz, et al., 2019) and (Wang, et al., 2017) presented uses of CNNs and FCNs as efficient time series classifiers. Application range from music genre classification (Costa, et al., 2017) to speaker identification (Ravanelli & Bengio, 2018) and environmental sound classification (Pons & Serra, 2019). However, to the best of our knowledge, there are no applications of TSC methods in HVAC signals classification in the existing literature. Some non-documented attempts were found to classify BAS time series based on their associated labels, however, in most of the cases, these labels are only numerical or alphanumerical indices that do not provide any useful information.

3 Methodology

Data was acquired from the BAS of the Daphne Cockwell Complex (DCC) at Ryerson University, Toronto; the data retrieval process is fully described in (Misic, et al., 2020). All BAS sensor, controller, and actuator point change of values are logged before being streamed to a time series database for analytics and event monitoring. Data is stored as a series of timestamped tuples recording the asset name, the meter name (a descriptor of the control or measurement point) and the reading value.

For data querying, a client sends a query to a relational database that contains the building ontology which is a formal model that represents the knowledge of the building operation (Gilani, et al., 2020). This database is referred to as ontology database. To obtain a time series related to a given room or a set of equipment, we send a query to the ontology database to get the corresponding HVAC equipment and their subsequent measurement points, we then query the time series database to obtain data for the desired period. Data from the DCC time series database was retrieved for a period of 16 months ranging from March 2019 to December 2020. Data was sampled at a rate of one sample per 10 minutes and was collected from 5 different rooms located in the first and second levels of the building. A two-day width sliding window was used to create shorter time series for training and testing. Each time series was labelled according to its source. As slow changing or constant instances did not provide any useful patterns that could be learned, any time series with a range of variation below 0.1 was discarded from the training and test sets. The overall methodology using this data is illustrated in Figure 1.

3.1 Point Type Classification

First, the time series are split according to their type. Binary signals are removed and continuous signals are further divided according to range and change frequency. Control signals serve to modulate mechanical equipment operation such as valves and pumps and typically range from 0% to 100%. Set point time series are subject to a few changes during long periods of time (many weeks or months) and keep their value all along this period (e.g., room temperature set point). Sensor's data may have any value in a predefined range and have often rapid changing patterns with noisy observations. In this paper, we present only the sensors and control time series classification and clustering.

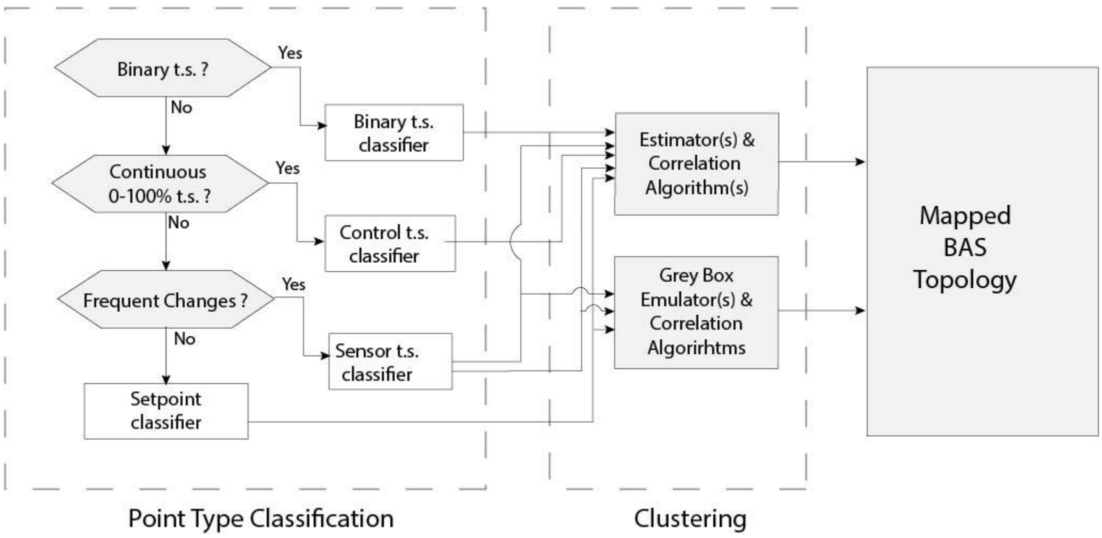

Figure 1 Classification flow chart

3.1.1 Sensor time series classification

In this experiment, our aim is to classify four types of sensor data: 1) T: room temperature, 2) INSLAB_T: in-slab temperature, 2) DA_T: discharge air temperature, and 3) Q: Air quality (CO_2 level). If we exclude Q, the first three time series are all temperature measurements and have the same range of values, thus a statistical method has little differentiation power. However, a thorough examination of the time series and their derivatives in Figure 2 demonstrates a few discriminative patterns.

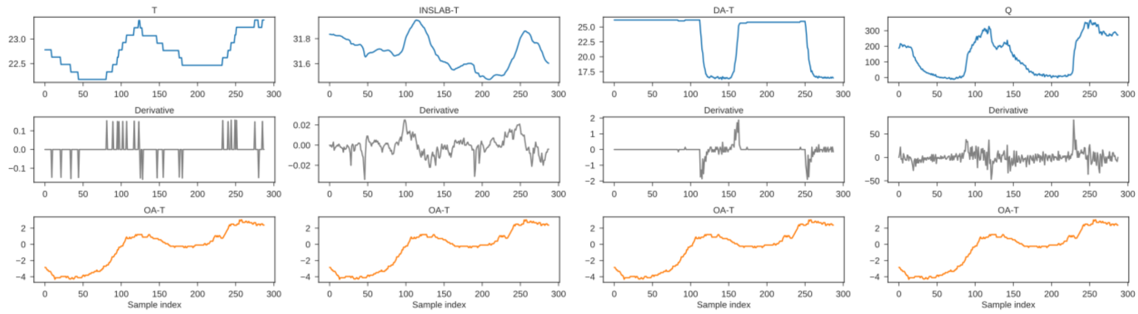

Figure 2 Continuous time series samples for 48 hours of data. Sensor measurement and their derivatives are shown on the first and second row. Outside air temperature is shown last.

The room temperatures exhibit small variations due to heating/cooling control around its setpoint, while in-slab temperatures have a wider range of temperature variation with smooth transitions; both temperatures have some correlation with the outside air temperature OA_T. Conversely, discharge air temperature (DA_T) exhibits sudden changes in temperature due to cool or warm air supply from the HVAC equipment. These air supply variations are readily noticeable as positive and negative spikes in its derivative. Carbon dioxide levels Q increase considerably during human presence in working hours and have a certain correlation with the outside temperature as it is warmer during the day than during the night. In addition, Q observations have a higher level of noise that is noticeable on the derivative chart. More patterns may exist in other samples and may go unnoticed to the human eye, rendering the use of hand-crafted features or rule-based models non-ideal.

Convolutional neural networks, which allow for automatic feature generation, provide a valuable solution. Features are associated with the presence of specific patterns that characterize each class of time series. The 1D-CNN can also find discriminative patterns on multiple time series at once. This property is handled effectively in our case by a three-layered input constructed by the time series, its derivative, and the outside air temperature (Figure 3). By calculating the convolution of a set of learned kernels with each input layer, the 1D-CNN generates a feature map that represents the discriminative patterns. These features are utilized by a fully connected neural network for classification. Additional models tested for benchmarking include 1NN-DTW (Bagnall, et al., 2017), TSF (Deng, et al., 2013), FCN (Wang, et al., 2017) and state-of-the-art LSTM-FCN (Karim, et al., 2019). The FCN only includes convolutional layers and

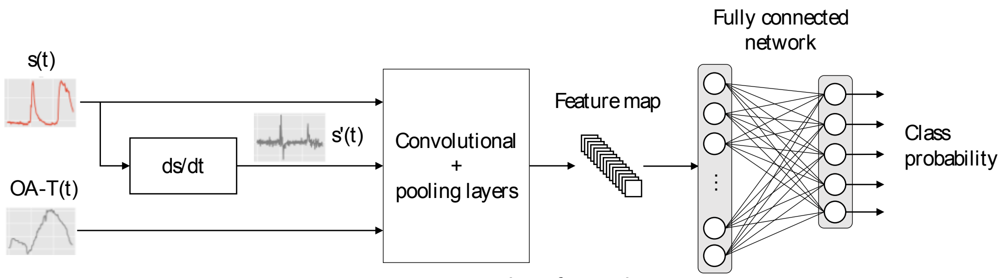

Figure 3 1D-CNN classifier architecture

replaces the cumbersome final dense layers with a global average pooling layer making the FCN a lightweight version of the typical CNN. On the other hand, the LSTM-FCN combines the FCN effectiveness with the LSTM proven capabilities in learning patterns in long time series.

3.1.2 Control and dependent time series classification

HVAC controls are typically prescriptive responses based on sensor measurement, which in turn affect subsequent sensor responses. Our approach, therefore, relies first on classifying sensor time series, then estimating control time series behavior to classify control time series and finally classifying subsequent response time series. The BAS schematic and resultant grey box approach are depicted in Figure 4 for a local HVAC system with cooling and heating coils.

The control valves for these coils respond to the variation in room temperature, T, from its set point, T_SP, so once these inputs have been classified, it becomes possible to generate an estimate of the coil valve actuator positions, CLG_O and HTG_O, which in turn can be used to estimate the discharge air temperature (DA_T) from the equipment.

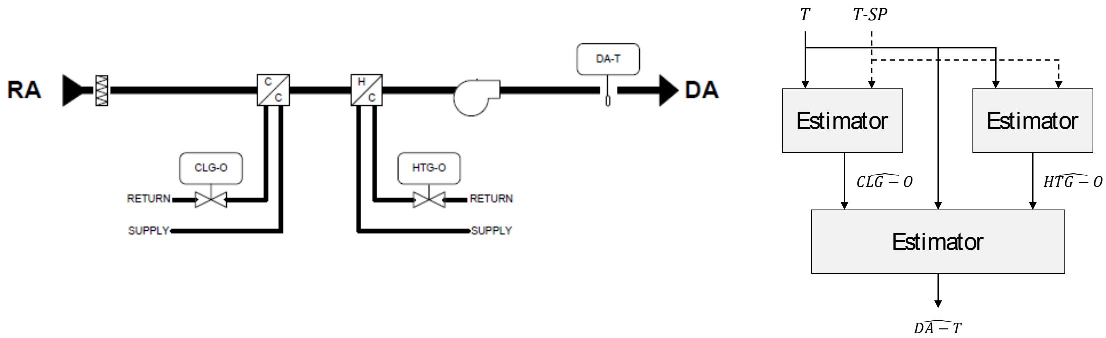

Figure 4 Sample control schematic (left) and associated grey-box-based classification process (right)

Linear control algorithms such as Proportional-Integral (PI) are the most common in HVAC systems (ASHRAE, 2016), whose generalized output is a weighted sum of the error and its integral (Equation 1), where CO is the controller output, K_C the controller gain, and T_i the integral reset time tuning parameter.

$$CO = K_C\, e(t) + K_C/T_i \int e(t)dt \qquad (1)$$

Although K_C and T_i are unknown, the observation of existing time series shows that using the same parameters for PI controllers on various rooms is a valid approximation. Figure 5a shows the full structure of the grey box emulator for PI controllers.

To begin, we assume that the set point is known and has a fixed value. For the DCC building, set points for heating and cooling are $T\_SP + 1$ and $T\_SP - 1$, where T_SP is the room temperature set point. We use the previously identified room temperature time series T to calculate the estimates of heating and cooling controls for the same room $\widehat{CLG\_O}$ and $\widehat{HTG\_O}$. We then use the correlation between these estimates and the unknown time series to determine the ones that are associated with the same room. We chose correlation because it is higher when the compared time series exhibits the same variations, that is why we can apply approximations to PI controller parameters.

Next, we define a dependent time series as one that can be predicted from previously identified ones. For this, we use the same approach as for control time series. The time series that depends directly on heating and cooling controls is Discharge air temperature DA_T which is low when CLG_O is activated and high when HTG_O is activated, otherwise it remains close to room temperature. Consequently, we can estimate DA_T by Equation 2.

$$\widehat{DA-T} = \mathrm{T} + \left(\frac{\mathrm{HTG-O}}{100}.(\mathrm{T_{HA}} - \mathrm{T}) + \frac{\mathrm{CLG-O}}{100}.(\mathrm{T_{CA}} - \mathrm{T})\right) \quad (2)$$

T is the room temperature, HTG_O and CLG_O are mutually exclusive heating and cooling controls (between 0% and 100%), T_{HA} and T_{CA} are approximate values of hot and cool air supply, respectively. This estimate is smoothed by a rolling mean of a window size of two samples. To find the DA_T associated with a given room, we search for the time series that has the highest correlation with $\widehat{\mathrm{DA}-\mathrm{T}}$ as shown in Figure 5b.

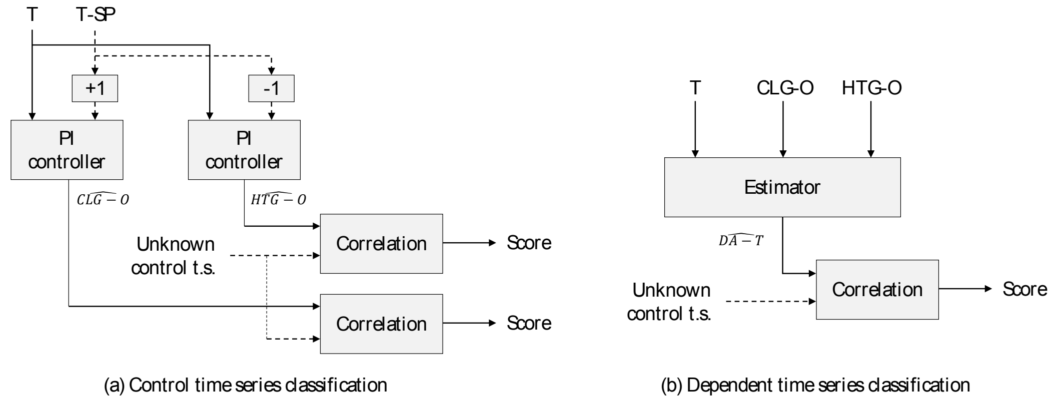

Figure 5 Control and dependent time series classification

4 Findings

While many time series were available, this study was limited to four types of sensor data from the case study building (DCC) – T, DA_T, INSLAB_T, and Q – and two control signals, CLG_O, HTG_O. These were selected because of their frequency (occurring in nearly every room in the building) and broad applicability for other buildings and HVAC system types. To permit a broad range of algorithms to be tested without computational cost becoming restrictive, five sample rooms were selected from the case study building.

4.1 Continuous time series classification

This experiment included two non-deep TSC algorithms: 1NN-DTW and TSF, as well as three deep algorithms: 1D-CNN, FCN and LSTM-FCN. We kept the same parameters of the FCN and LSTM-FCN suggested in (Wang, et al., 2017) and (Karim, et al., 2019) respectively as they provided the best results. The architecture and parameters of the 1D-CNN are detailed in Table 1.

Table 1. 1D-CNN architecture and parameters

Layer type	Parameters
Convolutional	Filters = 32, Kernel size = 5, activation = ReLU
Max pooling	Pool size = 2
Convolutional	Filters = 64, Kernel size = 3, activation = ReLU
Dropout	Rate = 0.5
Max pooling	Pool size = 2
Flatten	-
Dense	Units = 100, activation = ReLU
Dense	Units = 5, activation = Softmax

For all models we applied 8-fold cross validation by training the model on 14 months and testing it on 2 months. We assessed the performance of the model with both accuracy and F1-score macro that averages the F1 measure for each class. The F1-score is defined as the harmonic mean of the precision and recall obtained for each class. The training was performed on a Linux machine with a dual-core CPU@2.30GHz - 13GB of RAM size, equipped with a Testla T4 GPU - 16GB of VRAM. All deep models were trained for 150 epochs with a batch size of 32 and an Adam optimizer.

The confusion matrices in Figure 6 show that while the baseline 1NN-DTW performed poorly with 0.71 accuracy, TSF performed best reaching 0.95 of accuracy. The deep models showed a slightly lower performance with an accuracy of 0.92. Q and INSLAB_T were perfectly classified and that was expected as Q has a different range of measurements and noisy observations, while INSLAB_T change patterns were smooth and slow changing. The only visible confusion on the results was between T and DA_T. This occurs primarily in period when the local unit is not operating: when no cool or warm air is supplied, the duct temperature DA_T comes into equilibrium with the room temperature T.

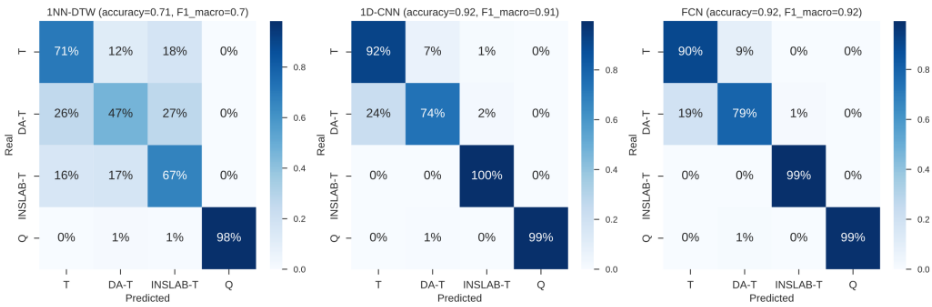

Figure 6. Selected Confusion matrices for continuous time series classification

When looking at the running time in Table 2, 1NN-DTW is by far the most time-consuming, while the lightweight structure of the 1D-CNN allows the fastest training and testing. Although TSF achieved the best performance in terms of classification, it took roughly 3 hours to run. Unfortunately, due its current implementation, it could not benefit from the hardware acceleration provided by GPUs as is the currently case with non-deep learning algorithms. Therefore, the 1D-CNN can be viewed as the model that achieves the best trade-off between performance and running time.

Table 2. Cross-validation total running time

1D-CNN	TSF	FCN	FCN-LSTM	1NN-DTW
1h 40min	2h 57min	5h 20min	8h 40min	11h 15min

4.2 Control and dependent times series classification

Heating and cooling control time series estimates are created from the room temperature T identified in the previous section. We apply the method in Figure 5a to generate $\widehat{CLG\_O}$ and $\widehat{HTG\_O}$, we then calculate the correlation between these estimates and the unknown time series. The temperature set point is supposed to be known. If the correlation score is higher than 0.5, the result is considered positive.

To assess this method, we retrieve three control time series from the time series database for five rooms: CLG_O, HTG_O and SW_O (a control signal for in-slab heating and cooling). We then divide data into frames of one week of width. Thus, for each frame we have up to 15 different time series (3 times series from 5 rooms). The classifier should generate only one positive result for the correct time series among the fifteen. Figure 7 shows an example of the real and estimated control time series for heating and cooling for one week. As small variations of control time series do not help identifying them correctly, we defined a minimum threshold for variation range that we fixed to 30%. Any estimate that does not fall into this range is discarded from the test set.

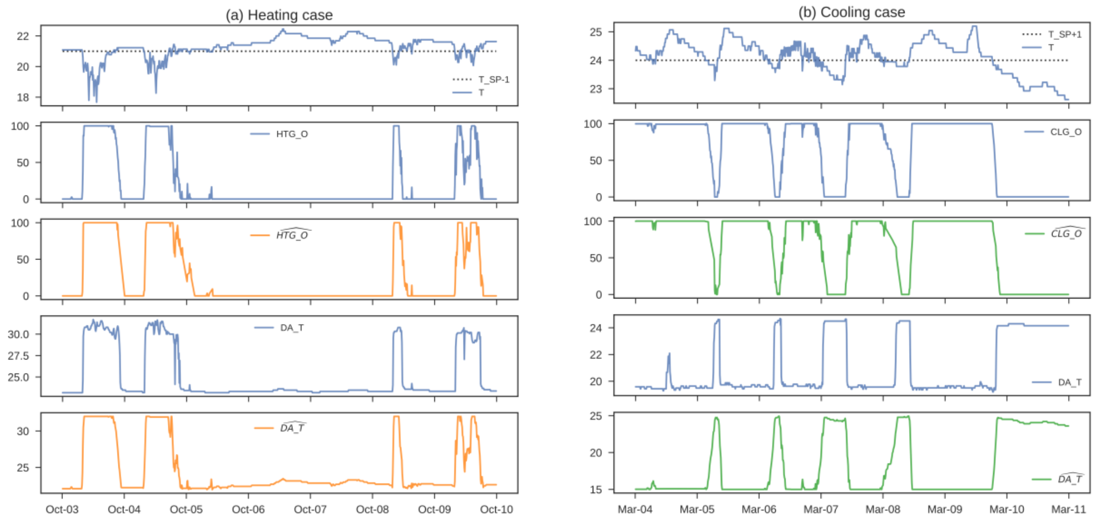

Figure 7 CLG_O, HTG_O and DA_T real and estimated time series for a) heating and b) cooling seasons

The confusion matrices in Figure 8(a) and (b) show that the method allowed to correctly classify instances of HTG_O and CLG_O with an accuracy of 90% and 94% respectively. Figure 8(c) shows the confusion matrix found with that the subsequent time-series classifier achieved a good accuracy of 88% in classifying correct DA_T instances.

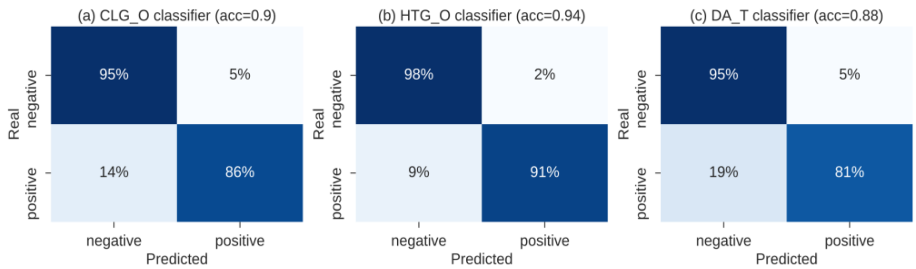

Figure 8 Confusion matrix for control and dependent time series

We follow the same approach we used in classifying control times series. In this case, we consider the following points from five rooms: DA_T, T and INSLAB_T for one-week time frames for assessment. As such, for any given room and time frame, the classifier should only label the correct DA_T time series as positive and the remaining fourteen as negative. To allow enough variations for pattern recognition, we discard any DA_T estimate for which the range of variations does not exceed 5°C as we know the difference between the room temperature and cool or warm supply air must exceed this difference in absolute value. We note that this grey-box approach is quite fast. It takes around 1 ms in average to generate a one-week size control/dependent time series estimate and classify one instance. The whole grey-box classification process of 6615 time series took less than 8s.

5 Discussion and Conclusions

This paper has presented an approach to classify building automation system time series. This method was tested on data from five different rooms in a large academic building and demonstrated a good accuracy. The method consisted in using a convolutional neural network to identify sensor measurement such as room temperatures, then applies the prior knowledge on controller's operation to classify related control time series and any other ones that can be identified by a combination of the room temperature and cooling and heating controls (grey-box model). Both methods resulted in a mean accuracy reaching 95% with TSF and 90% with the grey-box. This classification method enables identifying unlabeled points in building automation systems which allows applying advanced analytics and machine learning algorithms for many applications such as fault or anomaly detection, preventive maintenance, and energy saving. The same approach is useful in upgrading or migrating legacy BAS to newer systems. The use of these classifiers depends heavily on the availability of BAS information and consistent nomenclature. When a complete BAS documentation is available providing clear information about BAS topology, we can map most points without the need for classifiers. However, when the BAS is poorly configured or documentation is lacking, the pretrained continuous time series classifier can be used first to identify sensor data. This same architecture can be leveraged to identify various classes of sensorial data from other buildings logged trends. Opposite to non-deep methods such as TSF, CNNs can learn new patterns through transfer learning while keeping the previously learned primitives. This results in considerably reducing the training time. Once sensor points are identified, the control time series and subsequent points can be estimated and identified with the grey-box method. This is computationally efficient but requires a general understanding of control systems. From a scalability perspective, the linear complexity of this algorithm lends itself to larger-scale applications Further, because it is only the initial identification process that is computationally-intensive, this approach lends itself to event detection as deviations from expected data trends can immediately be identified.

The end use of these data analytics is to facilitate the creation of a university campus DT; a key technical challenge in this project has been the lack of documentation of legacy BAS systems, many of which do not have human-readable point names. Starting with a series of models developed using Autodesk Revit, a Forge interface was developed to query the SQL database in order to determine key point(s) of interest and then map the associated time-series data into the DT. Figure 9 shows a sample of the resultant visualization of point behaviour, where recorded measurements from selected equipment from the mechanical floor is displayed.

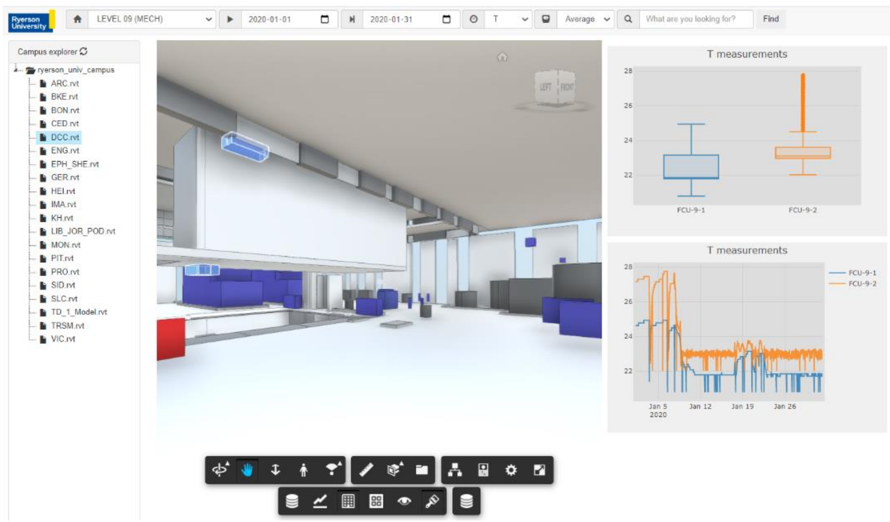

Figure 9 Sample Visualization from DT

This DT is composed of a web server retrieving buildings 3D models from the Autodesk cloud, equipment hierarchy data from a mySQL database, and measurements from an Elasticsearch database securely connected to BAS equipment networks. Future research will enhance this dashboard with event detection, identifying points demonstrating behaviour inconsistent with either their type or system relationships.

This study considered only data from five rooms from the academic building and classifying only the main time series involved in heating and cooling. As more data is streamed from other sources, we will extend this work to classify data related to other equipment such as boilers, chillers, and Energy Recovery (ERV) units. With more complex and rich patterns learned from signal clusters (e.g. room, level, mechanical room), the current research can be extended to expand the DT capabilities to detecting and diagnosing faulty equipment as well as recognizing anomalies. This will leverage the overall HVAC system performance and improve service quality and occupants' comfort.

References

ASHRAE, 2016. *ASHRAE Handbook 2016: HVAC Systems and Equipment.* Atlanta: SI Edition.

Bagnall, A. et al., 2017. The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data mining and knowledge discovery,* 31(3), pp. 606-660.

Costa, Y. M., Oliveira, L. S. & Silla Jr, C. N., 2017. spectrograms, An evaluation of convolutional neural networks for music classification using. *Applied soft computing,* Volume 52, pp. 28-38.

Dau, H. A. et al., 2018. *The UCR Time Series Classification Archive.* [Online] Available at: https://www.cs.ucr.edu/~eamonn/time_series_data_2018/ [Accessed 22 March 2021].

Deng, H., Runger, G., Tuv, E. & Vladimir, M., 2013. A time series forest for classification and feature extraction. *Information Sciences,* Volume 239, pp. 142-153.

Fawaz, H. I. et al., 2019. Deep learning for time series classification: a review. *Data Mining and Knowledge Discovery,* 33(4), pp. 917-963.

Gilani, S., Quinn, C. & McArthur, J., 2020. A review of ontologies within the domain of smart and ongoing commissioning. *Building and Environment,* p. 107099.

Gunay, H. B., Shen, W. & Yang, C., 2019. Text-mining building maintenance work orders for component fault frequency. *Building Research & Information,* 47(5), pp. 518-533.

Higdon, B. P., El Mokhtari, K. & Başar, A., 2019. *Time-series-based classification of financial forecasting discrepancies.* UK, International Conference on Innovative Techniques and Applications of Artificial Intelligence, pp. 474-479.

Karim, F., Majumdar, S., Darabi, H. & Chen, S., 2018. LSTM Fully Convolutional Networks for Time Series Classification. *IEEE Access,* Volume 6, pp. 1662-1669.

Karim, F., Majumdar, S., Darabi, H. & Harford, S., 2019. Multivariate LSTM-FCNs for time series classification. *Neural Networks,* Volume 116, pp. 237-245.

Lipton, Z. C., Kale, D. C., Elkan, C. & Wetzel, R., 2015. Learning to diagnose with LSTM recurrent neural networks. *arXiv preprint arXiv:1511.03677.*

Li, Y. & O'Neill, Z., 2018. *A critical review of fault modeling of HVAC systems in buildings.* Berlin Heidelberg, Springer, pp. 953-975.

Misic, T., Gilani, S. & McArthur, J., 2020. *BAS Data Streaming for Smart Building Analytics.* Sao Paulo, CIB W78, pp. 430-439.

Nanopoulos, A., Alcock, R. & Manolopoulos, Y., 2001. Feature-based classification of time-series data. *International Journal of Computer Research,* 10(3), pp. 49-61.

Pons, J. & Serra, X., 2019. *Randomly weighted cnns for (music) audio classification.* s.l., IEEE, pp. 336-340.

Ravanelli, M. & Bengio, Y., 2018. *peaker recognition from raw waveform with sincnet..* s.l., IEEE, pp. 1021-1028.

Szegedy, C. et al., 2015. *Going deeper with convolutions.* s.l., Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1-9.

Wang, Z., Yan, W. & Oates, T., 2017. *Time series classification from scratch with deep neural networks: A strong baseline.* s.l., International joint conference on neural networks (IJCNN), pp. 1578-1585.