BS 2021 1-3 SEPT BRUGES

INTERNATIONAL BUILDING PERFORMANCE SIMULATION ASSOCIATION

A gap-filling method for room temperature data based on autoencoder neural networks

Antonio Liguori[1,*], Romana Markovic[2], Jérôme Frisch[1], Andreas Wagner[2], Francesco Causone[3], Christoph van Treeck[1]
[1]E3D - Institute of Energy Efficiency and Sustainable Building, RWTH Aachen University, Aachen, Germany
[2]Building Science Group, Karlsruhe Institute of Technology, Karlsruhe, Germany
[3]Department of Energy, Politecnico di Milano, Milano, Italy
[*]Corresponding author: liguori@e3d.rwth-aachen.de

Abstract

This study explores the applicability of a deep learning-based approach for reconstructing missing room temperature data from different domains where relatively few training samples are available. For that purpose, the existing convolutional, long short-term memory (LSTM) and feed-forward autoencoders were combined with a suitable domain adaptation procedure. Eventually, the developed models were evaluated on data collected in four buildings with significant differences in thermal mass, design and location. The findings pointed out that the domain adaptation can be conducted efficiently by using a small data sample from the target domain. Additionally, the results showed that the proposed model can reconstruct up to 80 % of the missing daily room temperature inputs with RMSE accuracy of 0.6 °C.

Introduction

Heating, ventilation, and air conditioning (HVAC) systems account for more than 50 % of the total building energy use (Yang et al. (2016)). Since providing comfortable indoor environmental conditions is the key objective of HVAC systems, the control strategy together with the measured indoor environmental parameters have an explicit impact on the resulting energy consumption (Ge et al. (2013); Mba et al. (2016); Wu et al. (2017)). Particularly, the availability of precise indoor temperatures and indoor environmental quality (IEQ) data in general is of utmost importance for fault free and energy efficient HVAC control.

In modern buildings, HVAC control is implemented through advanced technologies, like building automation systems (BAS) (Ibrhim et al. (2020)). However, these systems are often characterized by sensors malfunctioning or network issues that could cause anomalies and information loss in the collected data (Liguori et al. (2021); Markovic (2020)). Missing data are, indeed, commonly observed in building energy management systems (Drgoňa et al. (2020)). As a consequence, these faulty measurements could determine a bad management of HVAC, leading to a degradation of the indoor environment conditions (Loy-Benitez et al. (2020)). In summary, all these issues lead to either higher than required energy consumption, lower thermal comfort or unstable conditions for real-time and model predictive HVAC control.

According to the international guidelines, missing data in BAS could be considered as "non annotated faults" and addressed by using methods for fault detection and diagnosis (FDD) (ASHRAE (2015); Markovic (2020)). Data-driven FDD models have a particular practical potential, due to their adaptability to complex systems (Mirnaghi and Haghighat (2020)) and since they do not require detailed physical system modeling. The major downside of the data-driven FDD approaches is, that they require extensive historical data that include normal and faulty conditions (Ebrahimifakhar et al. (2020)). This is currently one of the major burdens towards the inclusion of the data-driven FDD components for handling missing data in BAS. In order to overcome the challenge of high data requirements and computational costs of the modeling, researchers often rely on simplified techniques. However, these come to the cost of lower accuracy (Chong et al. (2016); Liguori et al. (2021)) and less reliable signal's dynamics representation.

Motivated by the need to address the problem, this paper proposes a deep learning-based method for indoor environmental quality data reconstruction. Methodologically, it relies on autoencoders neural networks, which proved to be able to capture the autoregressive properties of occupant behavior (OB) data (Liguori et al. (2021)), while the use of a suitable domain adaptation procedure assures the model's applicability in different target domains.

https://doi.org/10.26868/25222708.2021.30232

Related research

Autoencoder neural networks in context of building energy analytics

Autoencoder neural networks are unsupervised deep learning models trained to learn the inner representation of the input data (Qian et al. (2019)). Recently, these methods have started to receive great attention from the energy-related building research, due to their ability to detect and restore faulty sensor measurements (Loy-Benitez et al. (2020); Liguori et al. (2021); Liu et al. (2020); Fan et al. (2018); Kim and Cho (2019)). However, one of the main drawbacks of this approach is the need of massive training data (Loy-Benitez et al. (2020); Liguori et al. (2021)), which limits their practical integration into BAS.

Liguori et al. (2021) developed univariate autoencoder neural networks models to reconstruct missing IEQ data. The results proved that the missing indoor air temperature, relative humidity and CO_2 concentration data could be reconstructed with high accuracy. However, the available data set consisted of approximately 31,361 full days of observations per variable. Lower performance might be achieved by using a smaller data set. Loy-Benitez et al. (2020) exploited variational autoencoders in conjunction with convolutional layers to reconstruct missing indoor air quality (IAQ) subway data. The implemented models exploited different types of input variables that may be not always known and they were affected by the limited size of the used data set.

Domain adaptation

Consider two sets of data drawn from two different distributions A and B of a same variable (e.g. room temperature). Domain adaptation is the process of transferring knowledge from A to B (Glorot et al. (2011)). This technique may be particular useful in the event that the size of B is too limited to train a deep learning model entirely on it. In that case, part of the data from the target domain may be used to improve generalization on a pre-trained model.

As already pointed out by Markovic (2020), the research on domain adaptation for OB and energy consumption data in buildings is still too limited. Markovic et al. (2018) implemented a deep learning-based window opening model using data from a small group of offices in a commercial building. The proposed method was then adapted to other buildings, by running further iterations on the target domain data. The results proved, that the implemented algorithm was more accurate than other building-wise calibrated models. Zhang and Ardakanian (2019) developed an occupancy model based on recurrent neural networks (RNNs) on a specific data set and applied a domain adaptation technique to transfer the acquired knowledge to a building located in a different continent. The results led to almost the same performance of a model trained entirely on the target domain data. Based on the previous considerations, deep learning models were identified as a suitable approach to transfer the acquired knowledge from one domain to another (Glorot et al. (2011)).

Data set

In the scope of the initial model development study (Liguori et al. (2021)), the used data set consisted of 31,381 full days of observations distributed over four years for 84 offices in a mechanically ventilated building located in Aachen, Germany (Fütterer and Constantin (2014); Fütterer et al. (2013)). Models were trained using 30 % of available data (9,414 sequences) and validated using additional 10 % of data (3,138 sequences). The rest of the measurements were used for models' evaluation (18,829 sequences). Considering the relatively high dimension, data were downsampled from a minute-wise to a 30-minutes frequency.

In order to investigate the generalization capability of the original models to other domains, four additional data sets, collected in different buildings, were introduced. The first one consisted of a two year-long monitoring campaign of a naturally ventilated office building with passive cooling located in Frankfurt, Germany. The studied building has 17 single and double-occupied offices with operable windows facing either west or east (Schweiker et al. (2019)). The U-values of the walls range from 0.24 to 0.5 $\frac{W}{m^2K}$, whereas windows have U-values of 1.5 $\frac{W}{m^2K}$, solar transmittance less than 40 % and light transmittance less than 70 % (Schweiker et al. (2019)). For a detailed data set description, the reader is referred to Kleber and Wagner (2006) as well as to the publicly available data set repository (Schweiker et al. (2019)). The second data set is also publicly available (Langevin et al. (2015)) and it was recorded over one year in a mixed-mode office building in Philadelphia, USA. The building has private, shared private, cubicle and open desk offices. As for the observed occupants, 33 % were male while 67 % were female (Langevin et al. (2015)). The third open source data set was collected in a naturally ventilated open space area in a university building located in Vienna, Austria (Mahdavi et al. (2019)). The usable area includes single and double-occupied offices, a kitchen and a meeting room, all with operable windows (Mahdavi et al. (2019)). The fourth data set was recorded over approximately ten weeks in Karlsruhe, Germany. The data were collected on multiple workstations in an open space office, a meeting room and a two person office. The building in question is naturally ventilated and the data were collected between July 27th and October 6th 2020. The monitored spaces were typically occupied from Monday to Friday with flexible working hours. The maximum occupancy was limited to five people, due to the COVID-19 pandemic. The lengths of each data set were respectively 11,648,

7,761, 2,555 and 446 full workstation-wise days of observations. Data were preprocessed in the same way as described by Liguori et al. (2021), shuffled sequence-wise and sampled to the same frequency.

Table 1 summarizes descriptive statistics for each data set, while Figure 1 represents the indoor air temperature measurements distribution for each of the analyzed buildings. As presented in the latter figure, observed indoor air temperature measurements were normally distributed in three out of five data sets (Aachen, Frankfurt and Philadelphia), while the data had significantly different distribution in the data sets collected in Vienna and Karlsruhe.

Table 1: Descriptive statistics for each data set. Std stands for standard deviation.

	Aachen	Fran.	Phil.	Vienna	Karl.
Min [°C]	10.30	9.70	11.22	18.82	17.24
Max [°C]	32.20	30.80	32.22	35.14	33.46
Mean [°C]	22.67	22.59	22.64	23.81	25.34
Median [°C]	22.70	22.40	22.82	23.31	25.04
Std [°C]	1.11	1.88	1.82	2.29	2.83

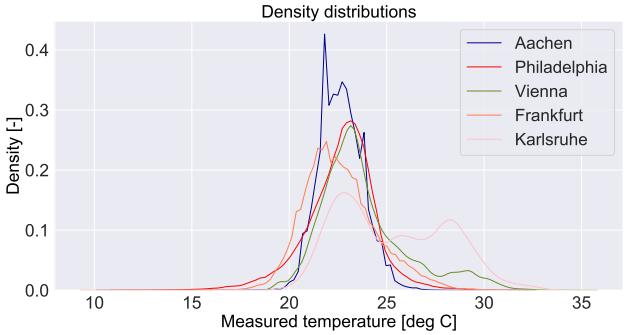

Figure 1: Indoor air temperature density distribution for each building.

Method

The proposed method consists of the following steps:

- Evaluation of the generalization capabilities of the existing models proposed by Liguori et al. (2021), using data from alternative buildings.
- Optimization, validation and evaluation of the same models, using data from alternative buildings and a suitable domain adaptation procedure.

Additionally, the experimental setting adopted in this paper is presented in Figure 2. Here, the term "adaptation set" refers to the portion of the data set, from the target domain building, used to run the domain adaptation procedure. In particular, domain adaptation is the process of transferring knowledge between two different distributions, as described in Section "Domain adaptation".

In the original paper from Liguori et al. (2021), three different denoising autoencoder architectures were used, namely feed-forward, convolutional and long short term-memory (LSTM). Convolutional and LSTM autoencoders were introduced to respectively capture the spatial and temporal dependencies in the input time-series. Each model was trained with daily indoor air temperature sequences, where an interval of data with random length was purposely set to zero to simulate the missing continues values. In particular, intervals' length ranged between 10 % and 90 % of the daily data. The occurrence of faulty measurements was also investigated at the end of each sequence, to represent the forecasting performance of the models.

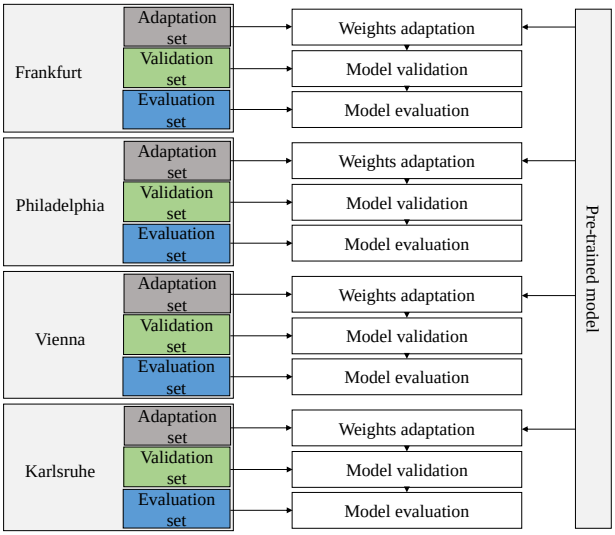

Figure 2: Overview of the proposed method.

Starting from the pre-trained weights of the original models, the adaptation process was performed by running further iterations on the target domain data, as proposed in Markovic et al. (2018). Further information related to the building characteristics was not explicitly encoded during the adaptation, but it could be learned by the model during the run of additional iterations. For Frankfurt, Philadelphia and Vienna, the length of the adaptation set ranged between 0 % (no adaptation) and 30 % (3,494, 2,328, 766 full workstation-wise days respectively) of the total data set, while the last 60 % (6,989, 4,657, 1,533 full workstation-wise days respectively) were used as evaluation set. For data collected in Karlsruhe, the adaptation set ranged between 0 % and 10 % (44 full workstation-wise days), due to the limited data availability (compared to the original study where significantly larger data sample was used for model development (Liguori et al. (2021)), while the last 80 % (357 full workstation-wise days) were used as evaluation set. Due to the stochastic nature of the training masking noise, adaptation was repeated ten times

		Aachen	Frankfurt				Philadelphia				Vienna				Karlsruhe	
	CR [-]	CONV	CONV				CONV				LSTM				LSTM	
			30 %	20 %	10 %	0 %	30 %	20 %	10 %	0 %	30 %	20 %	10 %	0 %	10 %	0 %
RMSE [°C]	0.20	0.31	0.40	0.40	0.41	0.45	0.55	0.54	0.54	0.49	0.36	0.36	0.40	0.82	0.50	1.26
	0.40	0.41	0.50	0.49	0.50	0.59	0.58	0.63	0.65	0.64	0.52	0.53	0.54	0.63	0.73	1.00
	0.60	0.50	0.56	0.56	0.58	0.68	0.70	0.72	0.73	0.80	0.59	0.61	0.64	0.79	0.88	1.14
	0.80	0.50	0.61	0.62	0.64	0.74	0.73	0.76	0.77	0.86	0.56	0.57	0.57	0.92	0.81	1.49
	Average	0.43	0.52	0.52	0.53	0.62	0.64	0.66	0.67	0.70	0.51	0.52	0.54	0.79	0.73	1.22
NRMSE [-]	0.20	0.21	0.17	0.17	0.17	0.8	0.27	0.26	0.26	0.25	0.14	0.15	0.17	0.34	0.10	0.26
	0.40	0.27	0.21	0.21	0.21	0.24	0.28	0.31	0.32	0.32	0.21	0.22	0.23	0.26	0.15	0.20
	0.60	0.33	0.24	0.24	0.24	0.27	0.34	0.35	0.36	0.40	0.24	0.25	0.27	0.33	0.18	0.23
	0.80	0.33	0.26	0.26	0.27	0.30	0.36	0.37	0.38	0.43	0.22	0.24	0.24	0.38	0.17	0.30
	Average	0.29	0.22	0.22	0.22	0.25	0.31	0.32	0.33	0.35	0.20	0.22	0.23	0.33	0.15	0.25

Table 2: Performance of the proposed method for filling sub-daily room temperature data gaps. "CONV" and "LSTM" stand for convolutional and LSTM denoising autoencoders.

and the best autoencoders were exported for further evaluation after being validated on 10 % of the data.

Normalization of the input data represented one of the main challenges in data preprocessing for domain adaptation. As highlighted by Markovic et al. (2018), the goal should be to include the actual distribution into the scale metrics. For each domain, the adaptation set was then normalized in the same way as proposed in Liguori et al. (2021) and fed to the models with zero mean and unit variance, while maintaining the initial form. The chosen computational environment consisted of Python 3.6.8, Tensorflow 1.12.0 and Keras 2.3.1.

Results

Performance evaluation metrics

In order to make reliable comparisons with the original study, it was opted for the same performance evaluation metrics used in Liguori et al. (2021), namely the root mean squared error (RMSE) and the normalized root mean squared error (NRMSE). The RMSE equation is given as follows (Candanedo et al. (2018); Ma et al. (2020)):

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}(X_i^{obs} - X_i^{inserted})^2}{n}}, \quad (1)$$

where X_i^{obs} are the $i{-}th$ real missing values, $X_i^{inserted}$ are the $i-th$ reconstructed missing values and n are the total number of missing data points. The NRMSE is obtained by normalizing the RMSE over the interquartile range (IQR) (Mahdavi and Tahmasebi (2017)), where IQR is the difference between the third and first quartile of the room temperature data.

Model performance

Table 2 summarizes the gap-filling performance of the proposed method for each of the observed buildings, by varying the adaptation and corruption rates (CR). Due to space constraints, only the best performing models for each two CR are represented. Since no clear guidelines about handling missing data in building control are available, the results of the pre-trained autoencoders on the Aachen data set are used as the only benchmark of this case. In this way, it is possible to quantify the performance drop of the original model, when applied to other domains. Further comparisons with other commonly adopted approaches can be found in Liguori et al. (2021).

On average, a domain adaptation with 10 % of the target domain data results in a better RMSE, namely from 4 % (Philadelphia) to 40 % (Karlsruhe) less than the no adaptation case. In particular, the proposed method can fill room temperature data gaps with an average RMSE of 0.62 °C (Aachen excluded), by using only 10 % of the target domain data (i.e. from 44 to 1,165 full workstation-wise monitoring days, depending on the size of the data set). However, for the convolutional model, the imputation performance can worsen for small gaps. By increasing the adaptation set size from 10 % to 30 % of the total data set, the average RMSE increases up to 0.55 °C.

The combined RMSE-NRMSE analysis points out that the most challenging data set to reconstruct is the one collected in Karlsruhe, with an average RMSE of 1.22 °C. However, despite the challenging room temperature distribution (Figure 1), the model performance on this data set is acceptable. By using 10 % of the target domain data for further optimization, the average RMSE drops to 0.73 °C, in view of a NRMSE of only 0.15 (48 % lower than Aachen). On the other hand, the worst performance is obtained on data collected in Philadelphia, where the NRMSE varies from 0.35 (no adaptation) to 0.31 (30 % adaptation).

Table 3 shows the performance of the proposed method for the forecasting of faulty real-time room temperature measurements. The masking noise is ap-

	PH [h]	Aachen		Frankfurt		Philadelphia		Vienna		Karlsruhe	
		LSTM	FP	LSTM	FP	LSTM	FP	LSTM	FP	LSTM	FP
RMSE [°C]	5.00	0.29	0.63	0.52	0.76	0.50	0.71	0.16	0.29	0.35	0.53
	9.50	0.46	0.74	0.63	0.73	0.80	0.88	0.32	0.58	0.47	0.54
	14.50	0.59	1.00	0.74	1.08	0.98	1.24	0.51	0.67	0.82	1.29
	19.00	0.63	1.04	0.81	1.54	1.07	1.71	0.58	0.72	0.88	1.22
	Average	0.49	0.85	0.68	1.03	0.84	1.14	0.39	0.57	0.63	0.90
NRMSE [-]	5.00	0.19	0.42	0.22	0.31	0.24	0.35	0.07	0.12	0.06	0.11
	9.50	0.30	0.49	0.27	0.30	0.39	0.43	0.13	0.24	0.08	0.11
	14.50	0.38	0.67	0.32	0.45	0.48	0.61	0.21	0.28	0.14	0.27
	19.00	0.41	0.69	0.35	0.64	0.52	0.84	0.24	0.30	0.16	0.26
	Average	0.32	0.57	0.29	0.42	0.41	0.56	0.16	0.24	0.11	0.19

Table 3: Performance of the proposed method for the forecasting of faulty real-time room temperature measurements. "LSTM" and "FP" stand for LSTM denoising autoencoder and forward propagation.

plied to the end of each input sequence and for this reason it is renamed as predictive horizon (PH). In line with the previous table, only the best performing models for each two PH are represented.

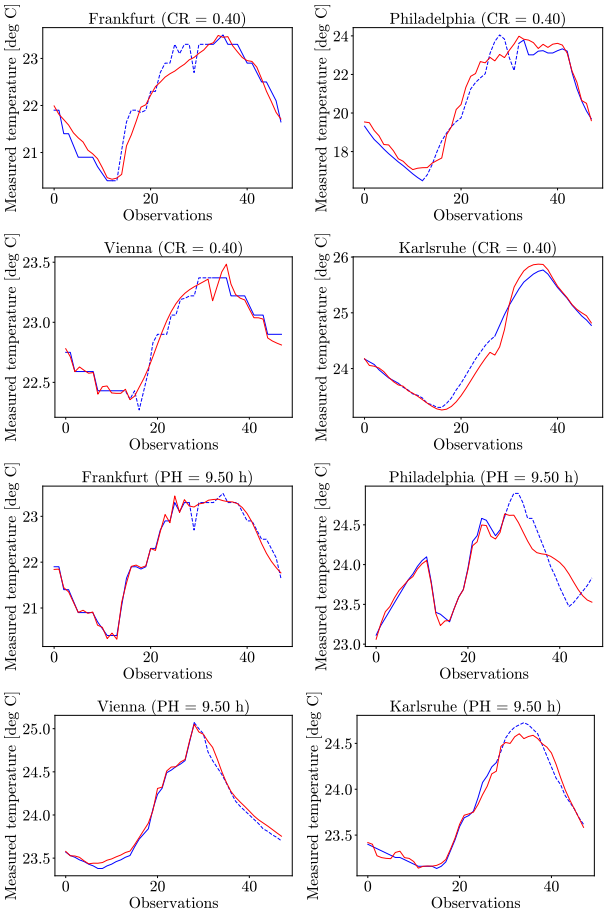

Figure 3: Gap-filling (CR) and forecasting (PH) of a random sequence with missing room temperature data from the test set. Blue colored line represents the real data. Hashed blue colored line represents the missing data. Red colored line represents the reconstruction of the whole sequence with the adopted model.

Furthermore, since the previous analysis did not point out a significant reconstruction improvement by varying the adaptation rate from 10 % to 30 %, only 10 % of the target domain data are used for models' optimization. The results are benchmarked to the performance of the original models and to the forward propagation method (FP). This method consists in reconstructing missing data with the last known real values (Woolley et al. (2009)). The choice of the last approach as an additional baseline for this case is based on the indication provided by the international guidelines on how to operate BAS devices in the event of network communication failure (Markovic (2020); ASHRAE (2015)).

For each data set, the forecasting of faulty room temperature measurements is always more accurate by using the LSTM model.

On average, a domain adaptation with 10 % of the target domain data results in a better RMSE than the forward propagation technique, namely from 26 % (Philadelphia) to 34 % (Frankfurt) less. In particular, the overall forecasting RMSE is 0.63 °C. Therefore, the accuracy of the proposed method is almost unchanged, with respect to the gap-filling case. However, the models' performance is more affected by the length of the missing interval.

Figure 3 shows exemplary room temperature data reconstruction over one random day from the test set, by using only 10 % of the target domain data for weights adaptation. Both the gap-filling (CR) and forecasting (PH) cases are presented. In order to represent the average reconstruction performance of the proposed method for both small and large gaps, it is opted for a masking noise of 40 %. It is possible to observe that, despite the small amount of used data from the target domains, the proposed method generalizes very well to multiple buildings. However, the forecasting of data from Philadelphia represents the main issue for the optimized model, confirming once again the findings in Table 3.

Discussion

The aim of this study was to propose a new approach to restore faulty or missing room temperature measurements. In particular, the focus was on the implementation of an advanced method that could be easily integrated independently of the quality, size and distribution of the available data set. For that purpose, an existing deep learning-based technique developed in the scope of a related study was analyzed and tested. The original models proposed by Liguori et al. (2021) were trained, validated and evaluated using an extensive data set of about 31,381 full days of observations, they relied only on the autoregressive properties of the indoor air temperature data, and they already had sufficient generalization capability. The additional buildings, introduced in the scope of this work, had clear differences in terms of thermal mass, design and location. Furthermore, the collected data sets consisted of various sizes, distribution of measured values and quality.

The performance of the original models were first evaluated on each system without further optimization. The worst performance, in terms of RMSE, was registered on data sets with non-normal room temperature distribution, namely Vienna and Karlsruhe. These different distributions, compared to the original training set, represented the highest complexity for the pre-trained models.

In the second set of experiments, a batch of data from the target domain was taken to adapt the pre-trained weights and let them vary from 10 % to 30 % of the total data set. For the convolutional and feed-forward models, a deterioration of performance was registered for small gaps that had size of up to 20 % of the sequences' length. This behavior could be explained with the use of batch-normalization in these configurations. While batch-normalization is intended to reduce conditions of internal covariate shift (Ioffe and Szegedy (2015)), it appears to affect the training of the previous models when a change in data distribution takes place externally. The influence on the small gap-filling accuracy can be traced back to the fact that batch-normalization was applied only to the non-corrupted data. Despite that, the overall performance of the convolutional autoencoder was higher than the LSTM configuration for both Frankfurt and Philadelphia, while the inverse occured for the rest. The reason behind that could be, once again, the use of batch-normalization in the convolutional model together with the completely different room temperature distributions in Vienna and Karlsruhe. Finally, no significant models' performance drop was registered by varying the adaptation rate from 10 % to 30 %. This means that the proposed gap-filling method could be effectively applied by using only a small percentage of the target domain data. The use of between 44 and 1,165 sequences to adapt the pre-trained weights is, indeed, a major improvement compared to the 9,414 daily observations used to train the model from scratch in Liguori et al. (2021).

In the last set of experiments, the proposed method was also predicting room temperature measurements over a time span between 5 and 19 hours. The previous gained knowledge about gap-filling accuracy was exploited by fixing the adaptation set size to 10 % on the available data. The LSTM configuration gained considerable advantage with respect to the convolutional model, confirming the findings in the original study. Namely, the temporal dependencies of the input data ("day-to-day" data points) were more relevant than the spatial correlations (data points "within the same day"), when a forecasting method was researched. Despite the performance variability of the proposed method, the reconstruction accuracy was always higher than the last-known method suggested by the international guidelines. Based on the previous considerations, the proposed models could be used as a back-up option in case of sensor failure in the real-time building control, by using only a small percentage of the target domain data for further optimization.

The main limitation of this method is that it builds on models trained to fill missing room temperature sub-daily sequences. Accordingly, continuous gaps larger than one day or next-day values cannot be reconstructed.

An analysis on the performance variability of the proposed method with respect to a different time discretization was not carried out neither in this nor in the original study. However, it has been already pointed out that, provided a sufficient large training set of data is available, there are no significant performance drops by modeling OB data ranging between minute-wise and hourly time resolutions (Markovic et al. (2019)). Re-training the original models with an other frequency of data may, therefore, only leads to a different required space complexity for models' evaluation as well as a different space and time complexity for models' training.

Finally, as an extension of this work, future research should focus on how to perform domain adaptation with as little data as possible.

Conclusion

This study explored the generalization capabilities of a model for reconstructing missing data streams from building automation. For that purpose, an earlier proposed autoencoder neural network was tested for reconstructing indoor temperature streams from four different office buildings. Furthermore, the domain adaptation was introduced as a performance optimization step and the suitable domain adaptation settings were analyzed. Based on the conducted experiments, the key findings may be summarized as

follows:

- The model could reconstruct up to 80 % of the missing daily values in significantly different buildings with RMSE of 0.6 °C, given a single shot domain adaptation.
- The domain adaptation is a required step for a satisfying model performance, due to the differences in shapes of temperature density distributions obtained in distinct data sets. This is of particular importance when transferring the model from mixed-mode to naturally ventilated buildings.
- The use of batch normalization in case of both convolutional and feed-forward denoising autoencoders led to lower performance, when compared to domain adaptation with no batch normalization step.
- The temporal dependencies (data points "day-to-day") of the input data (LSTM denoising autoencoder) are more relevant than the spatial correlations (data points "within the same day") (convolutional denoising autoencoder), when a forecasting method is applied.

In summary, the proposed method was successful in reconstructing indoor air temperatures in four buildings. Furthermore, the results of an earlier study (Liguori et al. (2021)) showed, that the model is also well applicable for producing the time-series of alternative room automation data streams, such as indoor CO_2 concentration and relative humidity. Based on these comprehensive results, the developed model shows a significant practical potential for the use in indoor climate monitoring, real-time room control as well as model predictive control.

Acknowledgements

Part of this work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – TR 892/4-1 – and by the German Federal Ministry of Economic Affairs and Energy (BMWi) as per resolution of the German Parliament under the funding code 03EN1002A. Simulations were performed with computing resources granted by RWTH Aachen University under project rwth0622. This paper benefited greatly from discussions with members of the IEA EBC Annex 79.

References

ASHRAE (2015). Guideline 13-2015 specifying direct digital control systems. *American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Atlanta, GA*.

Candanedo, L. M., V. Feldheim, and D. Deramaix (2018). Reconstruction of the indoor temperature dataset of a house using data driven models for performance evaluation. *Building and Environment 138*, 250–261.

Chong, A., K. P. Lam, W. Xu, O. T. Karaguzel, and Y. Mo (2016). Imputation of missing values in building sensor data. *Proceedings of SimBuild 6*(1).

Drgoňa, J., J. Arroyo, I. C. Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, et al. (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control*.

Ebrahimifakhar, A., A. Kabirikopaei, and D. Yuill (2020). Data-driven fault detection and diagnosis for packaged rooftop units using statistical machine learning classification methods. *Energy and Buildings 225*, 110318.

Fan, C., F. Xiao, Y. Zhao, and J. Wang (2018). Analytical investigation of autoencoder-based methods for unsupervised anomaly detection in building energy data. *Applied energy 211*, 1123–1135.

Fütterer, J. and A. Constantin (2014). *Energy concept for the E. ON ERC main building*. Hochschulbibliothek der Rheinisch-Westfälischen Technischen Hochschule Aachen.

Fütterer, J., A. Constantin, M. Schmidt, R. Streblow, D. Müller, and E. Kosmatopoulos (2013). A multifunctional demonstration bench for advanced control research in buildings—monitoring, control, and interface system. In *IECON 2013-39th Annual Conference of the IEEE Industrial Electronics Society*, pp. 5696–5701. IEEE.

Ge, F., X. Guo, H. Liu, J. Wang, and C. Lu (2013). Energy performance of air cooling systems considering indoor temperature and relative humidity in different climate zones in china. *Energy and buildings 64*, 145–153.

Glorot, X., A. Bordes, and Y. Bengio (2011). Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*.

Ibrhim, H., S. Khattab, K. Elsayed, A. Badr, and E. Nabil (2020). A formal methods-based rule verification framework for end-user programming in campus building automation systems. *Building and Environment 181*, 106983.

Ioffe, S. and C. Szegedy (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.

Kim, J.-Y. and S.-B. Cho (2019). Electric energy consumption prediction by deep learning with state explainable autoencoder. *Energies 12*(4), 739.

Kleber, M. and A. Wagner (2006). Results of monitoring a naturally ventilated and passively cooled office building in Frankfurt Main, Germany. In *Proceedings of EPIC 2006 AIVC Conference: Lyon, France*.

Langevin, J., P. L. Gurian, and J. Wen (2015). Tracking the human-building interaction: A longitudinal field study of occupant behavior in air-conditioned offices. *Journal of Environmental Psychology 42*, 94–115.

Liguori, A., R. Markovic, T. T. H. Dam, J. Frisch, C. van Treeck, and F. Causone (2021). Indoor environment data time-series reconstruction using autoencoder neural networks. *Building and Environment 191*, 107623.

Liu, Y., Z. Pang, M. Karlsson, and S. Gong (2020). Anomaly detection based on machine learning in IoT-based vertical plant wall for indoor climate control. *Building and Environment 183*, 107212.

Loy-Benitez, J., S. Heo, and C. Yoo (2020). Imputing missing indoor air quality data via variational convolutional autoencoders: Implications for ventilation management of subway metro systems. *Building and Environment 182*, 107135.

Ma, J., J. C. Cheng, F. Jiang, W. Chen, M. Wang, and C. Zhai (2020). A bi-directional missing data imputation scheme based on lstm and transfer learning for building energy data. *Energy and Buildings*, 109941.

Mahdavi, A., C. Berger, F. Tahmasebi, and M. Schuss (2019). Monitored data on occupants' presence and actions in an office building. *Scientific data 6*(1), 1–5.

Mahdavi, A. and F. Tahmasebi (2017). On the quality evaluation of behavioural models for building performance applications. *Journal of Building Performance Simulation 10*(5-6), 554–564.

Markovic, R. (2020). *Generic occupant behavior modeling for commercial buildings*. Ph. D. thesis, RWTH Aachen University.

Markovic, R., E. Grintal, A. Nouri, J. Frisch, and C. van Treeck (2019, Sep). Right on Time : Exploring Suitable Time Discretization for Occupant Behavior Co-Simulation. In *Proceedings of Building Simulation 2019 : 16th Conference of IBPSA : Rome, Italy, 2-4 September 2019 / edited by V. Corrado and A. Gasparella*, pp. 2099–2106. 16th IBPSA Conference, Rome (Italy), 2 Sep 2019 - 4 Sep 2019.

Markovic, R., E. Grintal, D. Wölki, J. Frisch, and C. van Treeck (2018). Window opening model using deep learning methods. *Building and Environment 145*, 319–329.

Mba, L., P. Meukam, and A. Kemajou (2016). Application of artificial neural network for predicting hourly indoor air temperature and relative humidity in modern building in humid region. *Energy and Buildings 121*, 32–42.

Mirnaghi, M. S. and F. Haghighat (2020). Fault detection and diagnosis of large-scale HVAC systems in buildings using data-driven methods: A comprehensive review. *Energy and Buildings*, 110492.

Qian, S., G. Li, W.-M. Cao, C. Liu, S. Wu, and H.-S. Wong (2019). Improving representation learning in autoencoders via multidimensional interpolation and dual regularizations. In *IJCAI*, pp. 3268–3274.

Schweiker, M., M. Kleber, and A. Wagner (2019). Long-term monitoring data from a naturally ventilated office building. *Scientific data 6*(1), 1–6.

Woolley, S. B., A. A. Cardoni, and J. W. Goethe (2009). Last-observation-carried-forward imputation method in clinical efficacy trials: review of 352 antidepressant studies. *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy 29*(12), 1408–1416.

Wu, X., Y. Liu, G. Liu, F. Wang, and Z. Wang (2017). Effect of supply air temperature on indoor thermal comfort in a room with radiant heating and mechanical ventilation. *Energy Procedia 121*, 206–213.

Yang, Q., J. Zhu, X. Xu, and J. Lu (2016). Simultaneous control of indoor air temperature and humidity for a chilled water based air conditioning system using neural networks. *Energy and Buildings 110*, 159–169.

Zhang, T. and O. Ardakanian (2019). A domain adaptation technique for fine-grained occupancy estimation in commercial buildings. In *Proceedings of the International Conference on Internet of Things Design and Implementation*, pp. 148–159.