


Accepted Manuscript

Learning short-term past as predictor of window opening-related human behavior in commercial buildings

Romana Markovic, Jérôme Frisch, Christoph van Treeck

PII: S0378-7788(18)32586-6
DOI: <https://doi.org/10.1016/j.enbuild.2018.12.012>
Reference: ENB 8941



To appear in: *Energy & Buildings*

Received date: 14 August 2018
Revised date: 8 November 2018
Accepted date: 13 December 2018

Please cite this article as: Romana Markovic, Jérôme Frisch, Christoph van Treeck, Learning short-term past as predictor of window opening-related human behavior in commercial buildings, *Energy & Buildings* (2018), doi: <https://doi.org/10.1016/j.enbuild.2018.12.012>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Learning short-term past as predictor of window opening-related human behavior in commercial buildings

Romana Markovic^{a,*}, Jérôme Frisch^a, Christoph van Treeck^a

^aE3D - Institute of Energy Efficiency and Sustainable Building, RWTH Aachen University, Mathieustr. 30, 52074 Aachen, Germany

Abstract

This paper addresses the question of identifying the time-window in short-term past from which the information regarding the future occupant's window opening actions and resulting window states in buildings can be predicted. The addressed sequence duration was in the range between 30 and 240 time-steps of indoor climate data, where the applied temporal discretization was one minute. For that purpose, a deep neural network is trained to predict the window states, where the input sequence duration is handled as an additional hyperparameter. Eventually, the relationship between the prediction accuracy and the time lag of the predicted window state in future is analyzed. The results pointed out, that the optimal predictive performance was achieved for the case where 60 time-steps of the indoor climate data were used as input. Additionally, the results showed that very long sequences (120-240 time-steps) could be addressed efficiently, given the right hyperparameters. Hence, the use of the memory over previous hours of high resolution indoor climate data did not improve the predictive performance, when compared to the case where 30/60 minutes indoor sequences were used. The analysis of the prediction accuracy in the form of *F1* score for the different time lag of future window states dropped from 0.51 to 0.27, when shifting the prediction target from 10 to 60 minutes in future.

Keywords: neural networks, stacked input vectors, sequence modelling, building automation systems, occupant behavior, window opening

1. Introduction

Occupant behavior (OB) has been identified to be one of the principal factors influencing the energy consumption in commercial buildings [1], [2], [3], [4]. Due to that, developing an accurate model that predicts human actions would be beneficial for achieving higher indoor comfort or optimization of energy consumption. Additionally, there have been a number of studies that addressed modeling the OB for its inclusion in building automation systems (BAS) [5], [6], [7], [8], [9], [10], [11].

According to the current research, OB in buildings is often defined as a discrete sequence in the temporal domain [5], [12], [13], [14], [15], [16], [17], [18]. As such, not only is it necessary to identify which variables lead to occupants' actions, but also in which temporal range

do the changes of variables in question occur, which motivated a number of studies on time-series modeling of OB. Liao et al. [14] presented a probabilistic graphical model for depicting the time-series of occupancy data. Fritsch et al. [12] presented the model that generates time series of window opening angles with the same statistics as the measured openings for the heating period. Dong and Andrews [5] proposed occupancy pattern recognition using semi-Markov models. Youngblood and Cook [13] introduced a hierarchical model for controlling the smart environment based on occupants' activities and concluded that learning algorithms built on Markov models experience performance issues when scaled to large problems. Wilke et al. [19] modeled residential activities based on time-dependent probabilities to start activities and their corresponding duration distributions. Cali et al. [20] concluded that the Markov chain models have the clear advantage of properly including the time dependency of the process. Additionally, they pointed out that a logistic regression analysis –as an alternative to the proposed Markov model allows modeling based on more variables, such as the indica-

*Corresponding author. Tel.: +49-241-80-25541 ; fax: +49-241-80-22030.

Email address: markovic@e3d.rwth-aachen.de (Romana Markovic)

tors of the indoor air quality. As already shown by Dong et al. [5], a parallel can be drawn between speech recognition problems and OB modeling. Resultantly, the recent findings on efficient speech recognition modeling may be projected in the field of OB. For instance, graphical models and in particular Hidden Markov Models (HMM) were widely used for representing the temporal variability of speech [21] since more than 30 years [22]. Although the HMMs are very efficient at solving a number of problems, their significant drawback is that probabilities need to be hard coded, which requires a new problem formulation for each task in question. On the other side, the deep neural networks do not require domain knowledge encoded into transition probabilities, since they can learn the feature mapping, given enough data that contains the underlying information. As a result, they outperformed graphical models in solving speech recognition tasks, usually by a large margin [21], [23].

The problem of modeling sequences using neural networks in the field of OB and energy efficient buildings has already been addressed by a number of studies [24], [25], [26], [27], [28], [29], [15], [30]. Coelho et al. [26] proposed a deep learning method for forecasting the sequences of energy consumption in micro grids. Wang et al. [24] proposed a single layer recurrent neural network (RNN) for predicting the occupancy on the following time-steps, where the input consists of a WiFi signal, while Zhao et al. [25] investigated multi-layer RNNs for predicting the occupancy based on physical measurements. Zhang et al. [28] proposed HVAC control using deep reinforcement learning (DLR) with a feed-forward neural network. The simulation results pointed out, that the DLR approach lead to 15 % lower heating energy consumption while perceiving thermal comfort. Wang et al. [29] presented a long-short term memory (LSTM) recurrent neural network for HVAC control. The results on the training and validation set pointed out, that the use of LSTM-RNNs could lead to improved energy efficiency and thermal comfort.

However, there is little work available on the duration of the time-window from which the information regarding occupants' short-term future actions can be retrieved. For instance, Wang et al. [24] proposed modeling the occupants' presence using five past time-steps of a WiFi signal, where time-step duration was 30 seconds. Hence, the research question regarding the time-window of physical measurements of indoor climate required for algorithmically efficient formulations of predictive OB models remained open.

In the scope of this study, the time-window that contains the information about future OB actions is investigated.

In particular, the following research questions are addressed:

- what is the sequence duration in the short-term past, from which the information regarding the occupants' future action can be retrieved?
- what is the relationship between the time lag between the current time-step and the occupants' action, and the accuracy of the resulting predictive model?
- the learning progress and problem formulation in case of stacked sequences of OB data will be analyzed.

For that purpose, a deep neural network is developed to predict the window states as the results of OB actions on future time-steps, where the input consists of indoor and outdoor climate information from the previous time-steps. The learning progress is analyzed, and eventually, the predictive performance was quantified. The remaining part of the paper is organized as follows: section 2 describes the hypothesis and experimental setting; Section 3 includes the results in particular, it addresses an optimal problem formulation, learning progress and resulting predictive performance. Eventually, the gained knowledge and limitations of this study are elaborated and summarized in Sections 4 and 5.

2. Method

The problem was formulated and the *hyperparameters* that lead to the optimal model formulation were investigated. Here, hyperparameters were settings used to control the behavior of the learning algorithm [31], [32], such as the number of neurons, the number of hidden layers, the learning rate and the regularization. One particular emphasis was put on the usage of *regularization*. In this context, the regularization was defined as any modification applied to improve the generalization capabilities of the model [32]. Once the optimal hyperparameters were defined, the training progress was analyzed through the changes in the *loss function* over *training iterations*, where loss was defined as overall measure of loss incurred in taken any of the available decisions or actions [31]. Training iteration was defined as a single optimization step during model training [32], [31]. Eventually, the learned content was interpreted using domain knowledge by analyzing the weights in the first hidden layer. Finally, the predictive accuracy was quantified and analyzed.

2.1. Experimental setting

The starting hypothesis is, that OB actions are influenced by the changes in the indoor climate that occurred in the short-term past. Here, the short-term past was defined as a time span over several hours. For the following experiments, the short-term past was analyzed in a range between 30 minutes and 4 hours. For instance, an increase in indoor air temperature or presence of additional occupants (and resulting increase in indoor CO_2 levels) may carry the information regarding future window openings. For that purpose, the model input consisted of weather and climate data from the current time-steps, and additionally the indoor air temperature, indoor air humidity and CO_2 concentration measured at every time-steps from the short-term past. The short-term past was considered as time-window used as the input for the developed model. The first objective of this paper was to investigate the size of this time-window, by treating it as a hyperparameter during model training. An optimal time-window size was analyzed between 30 minutes and 4 hours in the 30-minutes steps. The input sequence consisted of between 30 and 240 minute-wise time-steps, which resulted in problem dimensionality between 101 (21 features at time-step t plus 3 sequences of 30 time-steps each) and 741 (21 features at time-step t plus 3 sequences of 240 time-steps each). Even though the RNNs are widely used neural networks for the sequence modeling, an alternative approach had to be applied for the case of addressed lengths of OB data sequences. Although there is no theoretical limit on maximal sequence duration for conventional RNNs [35], the experimental proofs showed that the successful back propagation through time (BPTT) is possible on up to 5000 time-steps in cases where the gradient clipping and regularization were applied [36]. However, based on the results from the recent empirical studies on modelling the computer vision and speech recognition problems ([37], [38]), the maximal lengths of successfully modeled sequences were in the range between 15-30 time-steps. Hence, exploring OB sequences in the range of several hundreds of data points may be unfeasible using the above mentioned approaches. As a result, either a more complex learning method should be applied (gated RNNs, echo networks or memory networks), or the problem may be simplified by "un-rolling" the sequence through stacking the features for a feed forward neural network. The latter concept found application in related fields ([39], [40]), and it was also applied in the scope of this study. Resultantly, the input time-window was stacked as additional features for a feed-forward neural network (Figure 1), which resulted in relatively simple training, in comparison to RNNs.




Figure 1: Learning procedure overview. The sequences of measured indoor climate and the data from the last time-step t are defined as the model inputs. The duration of the input sequences i (red color symbol) is defined as additional hyperparameter. Eventually, the relationship between the predictive performance and the time lag of prediction l (red color symbol) is analyzed.

The time-step t was defined as step on which the prediction of the future window states was made. The input features were defined as vector X , which consists of $X_{outdoor}$ - measured outdoor climate and temporal information at time-step t , and the indoor climate information X_{indoor} on all time-steps between $t-i$ and t :

$$X = [X_{outdoor,t} \ X_{indoor,t} \ X_{indoor,t-1} \ \dots \ X_{indoor,t-i}]. \quad (1)$$

The window state was predicted for the future time-step, namely $t+l$. Here, the window state was the output based on the input features X and learned feature mapping $g(\cdot)$ through the hidden layers 1 to N :

$$y_{t+l} = g^{(n)}(\dots g^{(1)}(W^{(1)}X)), \quad (2)$$

where W^n , $n \in N$ were the tuned weights. Since it was aimed to perceive the indoor climate information at a high resolution, the temporal discretization was fixed at 1 minute steps, while even the lower time resolution (which would hence result in information loss) may hardly be explored. An additional argument for keeping the high resolution inputs was that the problem dimensionality in question still remained low, when compared to alternative problems addressed by similar methods, such as simplified computer vision problems (for example an image that consists of 32×32 pixels, where each pixel was defined as an RGB-color value).

It was considered that not every minute in the time-window carried information regarding the future OB actions. Also, it was considered that the members of the sequence that carry information are highly variant. As a result, the aim was that a neural network could learn which time-steps are irrelevant. Resultantly, a suitable architecture should be given enough hidden neurons to learn the dependencies and a possibility to exclude surplus neurons and input features by setting them to zero. For that purpose, a regularization using $L1$ norm was

introduced. Namely, [41] presented a theoretical proof that $L1$ results in efficient feature selection for the case similar to the problem addressed in this paper – where data sets with a potentially large number of irrelevant features could be present. The $L1$ factor defines a subset of weights in the hidden layers to have zero as optimal value, which results in sparse weight matrices [34]. As a result, $L1$ found a wide application as a feature selection method, which was also applied in the scope of this study.

Once the optimal input sequence duration was identified and the optimal model formulation was defined, the relationship between the time lag of prediction in future and the predictive performance was explored. Here, the time lag was defined as the time between the time-step of prediction t and the time-step for which the prediction was made. The investigated cases include the time lags between 10 and 60 minutes in the 10 minute-steps. For that purpose, the neural network with the hyperparameters as defined previously was retrained using the window state from $t+10$ to $t+60$ as a target variable.

2.2. Neural network training

An optimal neural network architecture in terms of number of neurons per hidden layer and the number of hidden layers was investigated for each duration of input sequence. The number of hidden layers was investigated in the range between 2 and 9 hidden layers, while the optimal number of neurons per hidden layer was searched in range between 1 and 400. $L1$ regularization is introduced in order to avoid over-fitting, by setting some of the hidden neurons to zero. Here, the $L1$ penalty was tuned in range between $1e-1$ and $1e-5$. An optimal learning rate was searched in the range between 0.2 and 0.01. It is opted for an adaptive learning rate. The rectified linear units (ReLU) were used as activation functions. The number of training iterations for different duration of training sequences was set based on the changes in the learning process between 10k and 90k. The batch size was kept fixed at 128, due to cache memory constraints. Resultantly, between 2 and 18 epochs were required for each training process.

2.3. Data set

The data were collected on an university building in Aachen, Germany. The building has operable windows and mechanical ventilation available in all offices from the data set

The HVAC system and the overview of the controlling strategy was described by Fütterer and Constantin [42] and is summarized bellow. The heating and cooling is

provided through concrete core activation as a base load system, while the mechanical air conditioning is used additionally to cover the peak heating and cooling loads. The volume flow of the mechanical ventilation is controlled by the indoor air quality, where a volume flow can be selected from three possible levels. The maximal air volume flow represents a air exchange rate of 3 [1/h], while the lower ventilation rate is assigned for energy saving reasons in case the office remained unoccupied. Due to non-disclosure agreements, the detailed controlling strategy for the mechanical ventilation is not available to the authors.

The available data consist of monitoring data from 52 offices in a period between January 1st, 2014 and October 1st, 2015, hence the monitoring campaign covered all seasons. The logging frequency was 1 minute for all indoor climate data, while the weather data were logged every 10 minutes. The data set includes the weather data that were measured on a nearby weather station by the Physical Geography and Climatology Group at RWTH Aachen University [43], [44] and the buildings monitoring data. In total, the data set consisted of features measured at the weather station (air temperature, air humidity, ground temperature, precipitation, wind speed, wind direction, maximal wind speed, avg. pressure, global and diffuse radiation), features for each office from the buildings monitoring data (presence, indoor CO_2 , humidity, two thermostat set points, air temperature, window state, outdoor temperature at building site) and the temporal information (Unix time, hour of day, day of week). For additional information on the used data, the reader is referred to [33], while the short overview of the measured indoor climate conditions is presented in Table 1.

Table 1: Mean values of the indoor climate and window opening data for the used data set.

variable	unit	mean value
indoor CO_2	ppm	514
T_{indoor}	$^{\circ}C$	22.9
indoor air humidity	%	38.5
window opening actions	1/h	1.07
proportion of window openings	-	0.07

The training set consists of approximately 600k data points collected on three single/double offices. The model suitability is optimized on the validation set that consists of 4 millions data points. Eventually, the remaining 15 millions data points were used for model evaluation.

2.4. Computational environment

The models were developed using the Tensorflow [45] library for Python 3.5 and for Python 3.6. The hyperparameters were tuned using computational resources from the RWTH compute cluster (CPU only), the GPU-Cluster at RWTH Aachen (parallel access to max. 8 nodes à 2 GPUs, out of 57 available Nvidia Quadro 6000), and a single PC for mixed CPU/GPU computations (CPU- Intel Core i7-6900K (3.2 GHz) , GPU- Nvidia GeForce GTX 1080). All computational resources were running on Linux-based operating systems.

3. Results

3.1. Input sequence duration

A qualitative analysis of the impact of the input sequence duration on the model's predictive performance is conducted, with the aim to identify an optimal model formulation. Since the input sequence duration is handled as an additional hyperparameter, the accuracy as a function of the input duration is analyzed on the validation set. The results of the hyperparameter search for each investigated input sequence duration are presented in Figure 3 in terms of accuracy for each class. The performance of the models with the varied input sequence duration is compared in Figure 2. A convex hull was created around the resulting data points representing the TPR/TNR combination for each input sequence duration. An entry towards the upper right corner of the diagram represents a better performance due to the higher TPR in combination with a higher TNR. As presented in Figure 2, the resulting true positive rate (TPR) and true negative rate (TNR) pointed out, that the TPR rate was slightly lower for longer input sequences, when compared to 30 and 60 minutes input sequence durations. As a result, input sequence durations longer than 60 minutes did not lead to any performance improvement by incorporating "memory" regarding the occupants' actions and resulting window states over the previous 1-4 hours. Hence, a satisfying performance was also achieved for multiple hyperparameter combinations where longer input sequences were used.




Figure 2: Qualitative representation of the relationship between the input sequence duration and validation performance.




Figure 3: Validation performance for each investigated case of input sequence duration. The true negative rate is plotted over the true negative rate.

Based on the results of the hyperparameter search, a subset of suitable hyperparameter combinations is identified, and the training is repeated 50 times for each combination with the randomly chosen initial weights. The mean validation results over repeated trainings are summarized in Table 2. The results pointed out, that

no significant difference could be observed in the validation performance, in cases where the input sequences over 30 and 60 minutes were used. Accuracy was in the same range for each sequence duration, while the proportion of correctly identified opened windows dropped for 7 percent points, where the sequences longer than 120 minutes were used.

Table 2: Validation results for 50 repeated model trainings using optimal hyperparameters for varied input sequence duration.

input sequence duration in minutes	ACC [-]	TPR [-]	TNR [-]	F1 [-]
30	0.91	0.41	0.95	0.56
60	0.90	0.41	0.94	0.55
90	0.89	0.39	0.93	0.53
120	0.89	0.39	0.93	0.54
180	0.90	0.32	0.95	0.46
240	0.88	0.32	0.92	0.45

3.2. Optimal model formulation

The developed main graph, and the structure of one hidden layer is visualized in Figure 4. The resulting network architecture consisted of five hidden layers with the following number of fully connected hidden neurons: 227-314-394-34-26. The $L1$ regularization coefficient was 0.01. The input sequence used was 60 minutes, where time-steps of 1 minute are selected. The model was trained for 50k iterative steps with the batch size of 128 data points (approximately 11 epochs) with the learning rate of 0.03.

3.3. Analysis of the learning progress

During the model training, the loss function formulated as the cross-entropy was minimized on the training set over multiple training epochs. Eventually, the training event at which the resulting validation error was minimal was identified as an optimal training duration. The training loss for each mini-batch and as a smoothed function is presented in Figure 5. Based on these results, the optimal number of training iterations was set to 50k with a learning rate of 0.03. Additionally, at the convergence of the loss function, the learning rate was reduced by factor 10 and eventually by factor 100, and additional 2×10^4 training iterations were conducted in order to reduce the variance in resulting predictive performance.




Figure 4: Visualization of the developed neural network. The graph notation corresponds to notation as used in Tensorboard [45].




Figure 5: Resulting training loss function (per mini-batch and smoothed) and validation loss.

The impact of the defined regularization on the size and complexity of the trained neural network is addressed by analyzing the tuned number of neurons per hidden layer and used weights in each layer. Figure 6 presents the proportion of weights set to zero in each hidden layer. Namely, the input feature selection was

conducted by setting 97 % of the weights to the first layer to zero, which resulted in a sparse weight matrix. Additionally, the results of the learning progress resulted in later updates in the last hidden layer, which correspond to the improvements of accuracy presented in Figure 5.




Figure 6: Fraction of the weights set to zero during model training.

In order to detect the relevant input features that were used to learn the mapping in hidden layers, the learned weights from the input layer to the neurons in the first hidden layer are presented in Figure 7. The weights from the input features sorted as indoor- and outdoor climate on the time-step of prediction and sequences of CO_2 measurements over last 60 minutes, indoor air humidity indoor and air temperature are presented on the x-axis, while the available neurons in the first hidden layers are listed out on y-axis.

Based on the learned information and the tuned weights, the neurons may be classified in two groups. Firstly, a number of neurons learned the relevant information from the last time-step (colored bars on the left side of the diagram). A separate group of neurons learned the sequence of the changes in indoor climate together with the measurements from the last time-step. Namely, there were neurons responsible to learn densely represented CO_2 over time as the main prediction for window openings. Simultaneously, the sparse course of the indoor air temperature over last 60 minutes was identified as an additional predictor of window states. In addition, most of the weights in the matrix from indoor air humidity were set to zero. Resultantly, it may be interpreted that the neural network did not use a significant information from the changes in the indoor air humidity as a

predictor of the future window states.




Figure 7: Magnitude of the weights from the input layer (x-axis) to the neurons first hidden layer (y-axis). For the visualization purpose, the bars representing the weights of higher absolute magnitude (colored fields) are enlarged vertically by the factor 6.

The analysis of the learned weights from the non-sequential features to the first hidden layer showed that all input features besides the volume of precipitation contributed to the prediction of future window states (Figure 8)¹. Additionally, weights coming from the variables "day of the week" and "outdoor air temperature" had a larger magnitude, compared to alternative features. This may be interpreted that these variables have strong impact on the predicted window states. The sequence of the indoor air temperature over the past 60 minutes was learned to have impact on the future window states. However, the weights from the indoor air temperature on the previous step to the first hidden layer did not have large magnitude, when compared to the alternative features. Based on these results, the trained neural network did not identify the indoor air temperature at the last time-steps as one of the main predictors of the future window states.

¹Rain* referred to the volume of the rain droplets. T_{out} refers to the air temperature at the weather station 1.3 km away from the building in question, while T_{out}^{**} refers to outdoor air temperature at the building site- there was up to 10 % difference between the both values.




Figure 8: Magnitude of the weights from the input features that refer to non-sequential data, to the neurons in the first hidden layer. For more detailed description of the input features, the reader is referred to [33].

3.4. What did the neural network learn?

In the following Section, the domain knowledge will be applied to exemplarily interpret the correctly predicted window states. The set of correctly identified window states consists of approximately 400k data points for the 10 minute time lag and around 200k data points for the 60 minute time lag. Since the elaboration on each of these data points would be unfeasible, 4 cases are extracted, where the network correctly predicted window states as open. The information that the trained neural network learned may be summarized into the following phenomena:

- events - the developed neural network learned that some occupants open the windows after the arrival (Figure 9 (A.1-A.3))
- past actions - the network could predict correctly that the window(s) will be opened again in 60 minutes, in case the short window opening occurred during the time-window used as the input sequence (Figure 9 (B.1-B.3))
- indoor climate and thermal comfort - the opened window state could be correctly predicted in case there was a steep rise in the indoor air temperature. This is of particular importance in the winter months, since a portion of the window opening caused by overheating could be correctly identified (Figure 9 (C.1-C.3))

- representations that could not be directly interpreted (Figure 9 (D.1-D.3))

A set of rules was defined, that could describe the above mentioned phenomena (events, past actions and overheating), in order to quantify the models' performance for these subcases. The arrival into an unoccupied office was defined either as a) the change in the occupancy status from unoccupied to occupied, or b) the CO₂ concentration that rise more than 50 ppm/hour from the starting concentration that was lower than 500 ppm. In this paper, a rate of the 50 ppm/hour was chosen to correspond to the presence of at least one occupant in the office with two work places [47].

Table 3: Confusion matrix of the predicted and measured window states for the subcases presented in section 3.4 (events, past actions and overheating). The results are presented as the absolute numbers of data points.

A. Events

opening probability: 14.4 %

		measured	
		opened	closed
predicted	opened	14792	35624
	closed	42476	256855

B. Past actions

opening probability: 10.9 %

		measured	
		opened	closed
predicted	opened	572	711
	closed	1520	8984

C. Overheating

opening probability: 4.1 %

		measured	
		opened	closed
predicted	opened	1015	1323
	closed	15584	38726




Figure 9: Changes in the indoor climate prior over one hour, where the window states were correctly predicted to be opened.

Past actions were defined as window opening and subsequent closing that both occurred within a 60 minutes input window. The overheating was described as the rise of indoor air temperature greater than 1.1 degree Celsius over one hour during a day where heating is activated (i.e. outdoor temperature under 12 degrees Celsius). The temperature rise was defined based on ASHRAE 62 [47], while the heating day was defined based on DIN EN 12831:1-2017 [48]. The model's performance in case of these subcases was summarized in Table 3.

3.5. Predictive performance

The results of 100 consecutive model trainings and evaluations on the test set using the same hyperparameters are presented in Table 4. The mean accuracy scored was 0.88, while the mean TPR and TNR were 0.37 and 0.92 respectively. The models' performance with respect to the data set's imbalance is quantified by the means of the *F1* score, which had the mean value of 0.51. In addition to the performance mean values over repeated training with random initial weights, the extreme values (minimal and maximal performance for each metrics) and the 25% and 75% quantiles were analyzed. Since these results showed low variance, the randomness of the weights initialization had no significant

impact on the models' predictive performance.

Table 4: Prediction performance of the investigated models for window opening.

	ACC	TPR	TNR	F1
	[-]	[-]	[-]	[-]
min	0.84	0.30	0.87	0.44
25 % quantile	0.87	0.35	0.91	0.49
mean	0.88	0.37	0.92	0.51
median	0.88	0.36	0.92	0.50
75 % quantile	0.89	0.38	0.93	0.52
max	0.90	0.46	0.94	0.58

The absolute performance results were analyzed using metrics proposed by [46] and the results are summarized in Table 5. The deviation between the observed and predicted proportion of time where windows were opened was 1 %. The prediction results overestimated the number of opening actions per day by the factor 2.4. Resultantly, the durations of sequences where windows did not change the state (both open and close) were predicted to be lower, when compared to the measured durations of window openings and closings.

Table 5: Relative predictive performance.

	25 %	median	75 %	IQR²
opening duration [hrs]				
observed	0.15	0.51	1.63	1.48
predicted	0.05	0.18	0.67	0.62
closing duration [hrs]				
observed	2.30	12.30	25.76	23.46
predicted	0.07	0.52	11.15	11.08
open state [-] actions [1/d]				
observed	0.07		1.07	
predicted	0.08		2.10	

²Distance between 25 % quantile and 75 % quantile.

The results showed, that 25 % of the window opening actions could be correctly identified on the evaluation set. It is defined that the window opening action occurred, in case the transition from the closed to the open window state took place at least one time-step after the time-step of observation, and no later than at the time-step of the time lag in future. Additionally, the prediction is handled as correct, in case the model predicted that the transition from the closed state, to the open window state will occur in the same time window.

3.6. Time lag between observation and prediction and comparison to alternative methods

The optimal models with the 60-minute input sequences were re-trained for the case where the prediction target was between 20 and 60 minutes after the time of observation. Based on the training progress, the training was conducted for 20k more iterations (around 3 epochs using mini batches of 128 data points) for the cases where the time lag was equal or greater than 30 minutes. Training with randomly chosen initial weights was repeated 100 times for each model and for each time lag duration.

The predictive performance for varied time lag between the time of observation and the time of prediction for the optimal models with 60 minutes input sequences is presented in Figure 10 (red-colored lines). The performance on the under-represented class (open windows) for both models was decreasing by 5-10% for each 10 minutes of delay between the time of observation. Due to this decreased performance on the under-represented class, the TNR was slightly increasing in case of longer time lag.

The model's predictive performance over a varied time horizon was compared to logistic regression, since this approach was the preferred model and was previously evaluated by a number of studies [49], [50], [51], [52]. The logistic regression uses the variables indoor and outdoor air temperatures that were identified using expert knowledge in the previous studies as the key drivers for window openings and closings. In total, three model implementations were analyzed in the scope of model comparison. Case A addresses the logistic regression with the parameters proposed by Haldi and Robinson [51]. In the case "B", the two variable model proposed by Haldi and Robinson was recalibrated using the training set described in section 2.3. The case "C", the above mentioned logistic regression model was re-trained using the training set from the section 2.3, while the hyperparameters were chosen based on the performance on the validation set. In case of models B-C initial experiments were conducted in order to make sure that the model does not overfit on a large training set. Since the validation accuracy in terms of TPR and TNR did not drop with the increased number of training samples, all 600k training samples were used. In case of model C the chosen solver was saga, due to the large training set size [53]. Additionally, the features were scaled in range between 0 and 1. The analyzed weighting of the classes include balanced weight, analysis of the overweighting of open windows with the factors in range from 2-15 as well as weighting based on priors.

No regularization was used, due to low problem dimensionality. The comparison results are summarized in Figure 10.




Figure 10: Performance comparison to alternative methods.

4. Discussion

The duration of the model training was around 0.5 CPU seconds per 10k iterations, which resulted in 2-5 seconds for training a model with a single hyperparameter combination. However, the introduction of the input sequence duration as an additional hyperparameter resulted in an extensive grid search. In total, approximately 15k core hours of computations were conducted for the hyperparameter search, repeated training with random initial weights, analysis of the relationship between the time lag and predictive accuracy and model

evaluation.

The optimal input sequence duration has been identified to be 60 minutes, based on the predictive performance on the validation set for 100 consecutive model trainings with random initial weights. This resulted in a 201-dimensional problem formulation. Although the predictive accuracy in a similar range could be achieved in case where a longer input sequences were used, there was no performance improvement by incorporating the information regarding the indoor climate over periods between 60 and 240 minutes.

The results from this empirical study are beneficial when choosing a suitable modeling approach for addressing the time-series of OB data for model predictive control, since the duration of sequences and the number of resulting time-steps used as model input are important for choosing optimal model architectures for further optimization and model development. Namely, the use of conventional RNN may result in an unstable model training due to exploding gradients due to long sequences that would be back-propagated through time (BPTT). As a result, the special case of RNN that include gated structures such as leaky units or long-short-term memory (LSTM) could be a suitable approach for further model development.

Based on the trained weights from the input layer to the first hidden layer, the sequence of the indoor CO_2 measurements was identified as a significant predictor of future window states. Interestingly, it carried more information regarding the future OB actions, when compared to the current air temperature or the sequence of the indoor air temperature. Consequently, a high resolution in the temporal domain of the CO_2 measurements should be one of the inputs for the predictive models of OB in commercial buildings.

Additionally, the sparse weight matrix formulation resulted in a low number of activated neurons per hidden layer and resultantly computationally efficient model evaluation. Also, the regularization applied on input features resulted in a lower problem dimensionality. Namely, the indoor air humidity is identified to have low impact on predictions regarding future window open states. The results may be interpreted carefully in case of mechanically ventilated buildings, as it was the case for the building in question. Namely, the humidity could be constantly low in the monitored offices due to mechanical ventilation control strategy, so that no conclusion regarding the impact of humidity on the actions could be made.

This work referred to the analysis of the impact of the short-term past on the OB, but the knowledge gained could be used as a basis for analyzing the long-term

impacts on OB such as acclimatization. In particular, a high temporal resolution (minute-wise) of indoor climate data would not be required for the measurements occurred more than 60 minutes in the past. Additionally, a similar implementation of the neural network could be applied for exploring longer time durations in lower resolution. For instance, introduction of multiple hidden layers could lead to accuracy improvement, and the similar regularization strategy could be applied.

The purpose of this work is to explore empirically the input sequence duration required for modeling the human actions in commercial buildings. This is conducted on the sub-case of modeling the window states. This is motivated by the necessity to further explore the human behavior as a sequential problem in temporal domain. The direct applications of gained knowledge, which include models inclusion in BAS for energy saving and comfort optimization were not addressed in the scope of this work, but they should be addressed in the future by the means of building performance simulation, laboratory studies or *in-situ* studies on the buildings in operation.

Some of the possible applications of this work are the use of predicted future window states as an input for model predictive control (MPC). For example, the future window states could be used as information to save energy used for heating, in case a window opening was predicted to occur. Alternatively, a window opening that was predicted to occur despite the running mechanical ventilation could be used as an input to save the energy used for mechanical ventilation in the time span until the window opening. However, no conclusions regarding the energy savings due to window opening related MPC could be drawn without their practical implementation in a field study and qualitative and quantitative energy consumption evaluation.

The absolute performance evaluation pointed out, that the neural network could accurately estimate the proportion of time where the window state was open, since the deviation between the measured and predicted open state was 1%. The duration of sequences where windows state did not change were analyzed through the variables "opening duration" and "closing duration". The predictions underestimated the duration of both states (open/closed) where no change occurred. This may be caused by the lack of knowledge incorporated regarding the transition probabilities. Although the transition probabilities over the maximal input sequence duration of sixty minutes could be learned, the transition probabilities over longer time-series were not addressed. In order to achieve the end-to-end learning of transition probabilities over whole duration where the win-

dow state did not change, the used model must address the sequences' full duration. Since this may be inefficient using the proposed temporal discretization or proposed method, the introduction of the varied discretization in time domain could be a promising approach. Alternatively, models such as memory networks may be suitable for the problem in question.

The models' predictive performance over varied time lag was compared to the logistic regression model proposed by [51] and its variations. The results showed, that the use of coefficients that were fitted on the original building leads to the significant overestimation of the proportion of the data points where windows were opened. Additionally, the model recalibration using identical settings did not lead to improved results. This was mainly caused by a large data imbalance of the opened windows in the data set used in the scope of this study. In the evaluation case C, the domain knowledge regarding the significant variables (namely indoor and outdoor temperature) and the modeling algorithm was incorporated based on the findings from previous studies. Additionally, the model implementation was adjusted in terms of the used solver, regularization, data preprocessing and class weights. Eventually, the model used in case C resulted in slightly lower predictive performance over time lags under 30 minutes, and significantly better performance over longer time lags. As a result, the predictive results showed that the suitable implementation of the logistic regression-based model proposed by Haldi and Robinson [51] could lead to competitive performance, when addressing longer time lags.

5. Limitations

The model was developed using data from an university office building in Germany. Validation and evaluation were conducted on a considerably large data set (training: 600k data points, validation: 4 millions data points, evaluation: 13 millions data points). Since no experiments and additional evaluation were conducted using the data from different building types of different locations, no conclusions regarding the large scale model applicability across different building types or geographic locations could be raised. Due to that, independent model evaluations using additional data sets or round robin studies should be conducted in future.

The data were obtained from a building with a mixed mode ventilation (mechanical ventilation and operable windows), which is a very common setting in commercial buildings in Europe. Here, an additional challenge was to correctly identify the window openings, where

occupants had multiple control options on their disposal (window opening and the use of mechanical ventilation).

In case of learning from OB in a setting with only mechanically ventilated building, the target variable window opening would be replaced by the alternative occupants' actions, such as adjusting the thermostat or use of local climatization systems. Based on the results from this study, a starting hypothesis could be set, that the time-window that contain the relevant information is similar in case of fully mechanically ventilated system, but the learning progress and the learned prediction variables differ from the case tested in this study. However, these questions were not analyzed in the scope of this empirical study, and they should be addressed in future work.

6. Conclusion

The goal of the presented work was to identify the time window in the short-term past where the information regarding the future OB actions was stored and to analyze the learning progress of deep learning driven predictive model of OB. The key findings may be summarized into the following:

- addressing more than 60 minutes of high resolution indoor climate data did not lead to improvements in the models' predictive performance,
- training of the deep neural network resulted in highly interpretable content, in terms of the weights from the input features, to the first hidden layer,
- a dense representation of sequential indoor CO_2 measurements was identified as one of the main predictors of window states,
- the predictive performance dropped with the increased forecasting horizon; resultantly, minute-wise predictions for more than 30 time-steps in the future could be made only with poor performance.

A significant contribution of the proposed method is learning end-to-end information regarding the drivers for future actions. Furthermore, the first hidden layer consisted of neurons specialized for phenomena that were interpretable and that were in accordance with the findings from the related studies based on domain knowledge. However, this can by no means be a replacement for research on understanding human actions

in buildings. It may rather be considered as a supporting method, which allows researchers to extract the knowledge from a high volume of data, where conventional approaches may be unfeasible.

7. Acknowledgements

Simulations were performed with computing resources granted by RWTH Aachen University under project nova0015. The authors appreciate the financial support of this work by the German Federal Ministry of Economics and Energy (BMWi) as per resolution of the German Parliament under the funding code 03ET1289D. We thank Mark Wesseling, Davide Cali and Felix Nienaber from EBC Institute, E.ON ERC at RWTH Aachen University for providing the monitoring data, as well as to the Physical Geography and Climatology Group for providing weather data. This paper benefited greatly from discussions with members of IEA EBC Annex 79.

References

- [1] Wagner, A., O'Brien, W., Dong, B. (Eds.). (2017). Exploring occupant behavior in buildings: methods and challenges. Springer.
- [2] Azar, E., Menassa, C. C. (2014). A comprehensive framework to quantify energy savings potential from improved operations of commercial building stocks. *Energy Policy*, 67, 459-472.
- [3] Hong, T. and Lin, H. Occupant behavior: impact on energy use of private offices. Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA (US).
- [4] Hong, T., D'Oca, S., Turner, W. J., Taylor-Lange, S. C. (2015). An ontology to represent energy-related occupant behavior in buildings. Part I: Introduction to the DNAs framework. *Building and Environment*, 92, 764-777.
- [5] Dong, B., Andrews, B. (2009, July). Sensor-based occupancy behavioral pattern recognition for energy and comfort management in intelligent buildings. In *Proceedings of building simulation* (pp. 1444-1451).
- [6] Dong, B., Lam, K. P. (2014, February). A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting. In *Building Simulation* (Vol. 7, No. 1, pp. 89-106). Springer Berlin Heidelberg.
- [7] Dobbs, J. R., Hency, B. M. (2014). Model predictive HVAC control with online occupancy model. *Energy and Buildings*, 82, 675-684.
- [8] Yang, R., Wang, L. (2013). Development of multi-agent system for building energy and comfort management based on occupant behaviors. *Energy and Buildings*, 56, 1-7.
- [9] Majumdar, A., Setter, J. L., Dobbs, J. R., Hency, B. M., Albonesi, D. H. (2014, November). Energy-comfort optimization using discomfort history and probabilistic occupancy prediction. In *Green Computing Conference (IGCC), 2014 International* (pp. 1-10). IEEE.
- [10] Peng, Y., Rysanek, A., Nagy, Z., Schlöter, A. (2018). Using machine learning techniques for occupancy-prediction-based cooling control in office buildings. *Applied Energy*, 211, 1343-1358.

- [11] Mirakhorli, A., Dong, B. (2016). Occupancy behavior based model predictive control for building indoor climate—A critical review. *Energy and Buildings*, 129, 499-513.
- [12] Fritsch, R., Kohler, A., Nygård-Ferguson, M., Scartezzini, J. L. (1990). A stochastic model of user behaviour regarding ventilation. *Building and Environment*, 25(2), 173-181.
- [13] Youngblood, G. M., Cook, D. J. (2007). Data mining for hierarchical model creation. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37(4), 561-572.
- [14] Liao, C., Lin, Y., Barooah, P. (2012). Agent-based and graphical modelling of building occupancy. *Journal of Building Performance Simulation*, 5(1), 5-25.
- [15] Li, Z., Dong, B. (2017). Investigation of A Short-term Prediction Method of Occupancy Presence in Residential Buildings. IBPSA, San Francisco, CA.
- [16] Li, Z., Dong, B. (2018). Short term predictions of occupancy in commercial buildings Performance analysis for stochastic models and machine learning approaches. *Energy and Buildings*, 158, 268-281.
- [17] Ardakanian, O., Bhattacharya, A., Culler, D. (2018). Non-Intrusive Occupancy Monitoring for Energy Conservation in Commercial Buildings. *Energy and Buildings*.
- [18] Pereira, P. F., Ramos, N. M. (2018). Detection of occupant actions in buildings through change point analysis of in-situ measurements. *Energy and Buildings*.
- [19] Wilke, U., Haldi, F., Scartezzini, J. L., Robinson, D. (2013). A bottom-up stochastic model to predict building occupants' time-dependent activities. *Building and Environment*, 60, 254-264.
- [20] Calí, D., Wesseling, M. T., Müller, D. (2018). WinProGen: A Markov-Chain-based stochastic window status profile generator for the simulation of realistic energy performance in buildings. *Building and Environment*, 136, 240-258.
- [21] Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A. R., Jaitly, N., ... Kingsbury, B. (2012). Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6), 82-97.
- [22] Baker, B. R. (1986). Using images to generate speech. *Byte*, 11(3), 160-168.
- [23] Graves, A., Mohamed, A. R., Hinton, G. (2013, May). Speech recognition with deep recurrent neural networks. In *Acoustics, speech and signal processing (icassp)*, 2013 IEEE international conference on (pp. 6645-6649). IEEE.
- [24] Wang, W., Chen, J., Hong, T., Zhu, N. (2018). Occupancy prediction through Markov based feedback recurrent neural network (M-FRNN) algorithm with WiFi probe technology. *Building and Environment*, 138, 160-170.
- [25] Zhao, H., Qi, Z., Wang, S., Vafai, K., Wang, H., Chen, H., Tan, S. X. D. (2016, May). Learning-based occupancy behavior detection for smart buildings. In *Circuits and Systems (ISCAS)*, 2016 IEEE International Symposium on (pp. 954-957). IEEE.
- [26] Coelho, I. M., Coelho, V. N., Luz, E. J. D. S., Ochi, L. S., Guimarães, F. G., Rios, E. (2017). A GPU deep learning meta-heuristic based model for time series forecasting. *Applied Energy*, 201, 412-418.
- [27] Kraipeerapun, P., Amornsamankul, S. (2017, February). Room occupancy detection using modified stacking. In *Proceedings of the 9th International Conference on Machine Learning and Computing* (pp. 162-166). ACM.
- [28] Zhang, Z., Chong, A., Pan, Y., Zhang, C., Lu, S., Lam, K. P. (2018). A Deep Reinforcement Learning Approach to Using Whole Building Energy Model For HVAC Optimal Control. 2018 Building Performance Modeling Conference and Sim-Build co-organized by ASHRAE and IBPSA-USA
- [29] Wang, Y., Velswamy, K., Huang, B. (2017). A Long-Short Term Memory Recurrent Neural Network Based Reinforcement Learning Controller for Office Heating Ventilation and Air Conditioning Systems. *Processes*, 5(3), 46.
- [30] Li, Z., Dong, B. (2017). A new modeling approach for short-term prediction of occupancy in residential buildings. *Building and Environment*, 121, 277-290.
- [31] Bishop, C. M., Mitchell, T. M. (2006). *Pattern Recognition and Machine Learning*.
- [32] Goodfellow, I., Bengio, Y., Courville, A., Bengio, Y. (2017). *Deep learning (Vol. 1)*. Cambridge: MIT press.
- [33] Markovic, R., Grintal, E., Wölki, D., Frisch, J., van Treeck, C. (2018). Window Opening Model using Deep Learning Methods. arXiv preprint arXiv:1807.03610.
- [34] Goodfellow, I., Bengio, Y., Courville, A. (2017). *Deep learning*. MIT Press.
- [35] Mozer, M. C. (1995). A focused backpropagation algorithm for temporal. *Backpropagation: Theory, architectures, and applications*, 137.
- [36] Pascanu, R., Mikolov, T., Bengio, Y. (2013, February). On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning* (pp. 1310-1318).
- [37] McLaughlin, N., Martinez del Rincon, J., Miller, P. (2016). Recurrent convolutional network for video-based person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1325-1334).
- [38] Chung, J., Gulcehre, C., Cho, K., Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *NIPS 2014 Deep Learning and Representation Learning Workshop*.
- [39] Wu, Z., Valentini-Botinhao, C., Watts, O., King, S. (2015, April). Deep neural networks employing multi-task learning and stacked bottleneck features for speech synthesis. In *Acoustics, Speech and Signal Processing (ICASSP)*, 2015 IEEE International Conference on (pp. 4460-4464). IEEE.
- [40] Abdel-Hamid, O., Mohamed, A. R., Jiang, H., Deng, L., Penn, G., Yu, D. (2014). Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing*, 22(10), 1533-1545.
- [41] Ng, A. Y. (2004, July). Feature selection, L1 vs. L2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning* (p. 78). ACM.
- [42] Fütterer, J., Constantin, A. (2014). Energy concept for the E.ON ERC main building. Hochschulbibliothek der Rheinisch-Westfälischen Technischen Hochschule Aachen.
- [43] Schneider, C. and Ketzler, G. (2014). Klimamessstation Aachen-Hörn - Monatsberichte, RWTH Aachen, Geographisches Institut, Lehr- und Forschungsgebiet Physische Geographie und Klimatologie ISSN 1861-4000".
- [44] Schneider, C. and Ketzler, G. (2015). Klimamessstation Aachen-Hörn - Monatsberichte, RWTH Aachen, Geographisches Institut, Lehr- und Forschungsgebiet Physische Geographie und Klimatologie ISSN 1861-4000".
- [45] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C. and others. (2016). Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint arXiv:1603.04467.
- [46] Mahdavi, A., Tahmasebi, F. (2017). On the quality evaluation of behavioral models for building performance applications. *Journal of Building Performance Simulation*, 10(5-6), 554-564.
- [47] ASHRAE (2010). Guideline 62.1-2010, Ventilation for acceptable indoor air quality. Atlanta, GA, USA: American Society of Heating, Ventilating, and Air Conditioning Engineers.
- [48] DIN EN 12831 (2017). Energy performance of buildings - Method for calculation of the design heat load. DIN Deutsches

Institut für Normung.

- [49] Rijal, H. B., Tuohy, P., Humphreys, M. A., Nicol, J. F., Samuel, A., Clarke, J. (2007). Using results from field surveys to predict the effect of open windows on thermal comfort and energy use in buildings. *Energy and buildings*, 39(7), 823-836. ISO 690
- [50] Rijal, H. B., Tuohy, P., Nicol, F., Humphreys, M. A., Samuel, A., Clarke, J. (2008). Development of an adaptive window-opening algorithm to predict the thermal comfort, energy use and overheating in buildings. *Journal of Building Performance Simulation*, 1(1), 17-30.
- [51] Haldi, F., Robinson, D. (2009). Interactions with window openings by office occupants. *Building and Environment*, 44(12), 2378-2395. <https://doi.org/10.1016/j.buildenv.2009.03.025>
- [52] Schweiker, M., Haldi, F., Shukuya, M., Robinson, D. (2012). Verification of stochastic models of window opening behaviour for residential buildings. *Journal of Building Performance Simulation*, 5(1), 55-74.
- [53] Defazio, A., Bach, F., Lacoste-Julien, S. (2014). SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in neural information processing systems* (pp. 1646-1654).