A Fully Automated and Scalable Approach for Indoor Temperature Forecasting in Buildings Using Artificial Neural Networks

Jakob Bjørnskov – University of Southern Denmark, Denmark – jabj@mmmi.sdu.dk
Muhyiddine Jradi – University of Southern Denmark, Denmark – mjr@mmmi.sdu.dk
Christian Veje – University of Southern Denmark, Denmark – veje@mmmi.sdu.dk

Abstract

Improving the performance of buildings is a core pillar to attaining future energy and environmental goals in different countries, considering that the building sector is a major contributor in terms of both energy consumption and carbon emissions. These ambitious goals and the call for smarter, energy-efficient, and flexible buildings have called for innovative and scalable energy and indoor thermal comfort modeling and prediction approaches. This work presents a fully automated and scalable solution using Artificial Neural Networks to forecast indoor room temperatures in buildings. A case study of an 8500 m² university building in Denmark was considered for testing and evaluating the proposed approach. An extensive dataset was constructed with sensor data from 76 rooms that contain both readings on indoor temperature, CO_2 concentrations, and actuating signals on radiator valves and dampers, as well as outdoor ambient conditions. Using this dataset, a well-performing architecture is identified, which provides accurate temperature predictions in the various rooms of the building for prediction horizons of 24 hours.

1. Introduction

Buildings are widely regarded as one of the major contributing sectors in terms of both energy consumption and CO_2 emissions. Furthermore, future energy systems with high fractions of Renewable Energy Sources (RES) depend on high demand-side flexibility. Therefore, there is a clear need for increasing not only the performance but also the flexibility of buildings. However, to achieve feasible and intelligent operation strategies for both cost minimization and flexibility services implementation without compromising the indoor comfort levels of buildings, reliable and accurate forecasting of building indoor thermal behavior is vital. In terms of indoor temperature forecasting, Artificial Neural Network (ANN) models have shown great potential in capturing the dynamics with high prediction accuracy (Alawadi, et al., 2022). In addition, these models can also be easily adapted and scaled up to different building cases.
This work presents an ANN-based approach that requires no prior specifications for the modeled building and can achieve accurate indoor temperature predictions for long prediction horizons of 24 hours or more. The models developed generalize well enough to be used for scenario planning and what-if analyses, e.g., to test the impact of custom setpoint and shading schedules on indoor temperature.

2. Case Definition

The building under consideration in this work is an 8500 m² highly energy-efficient university building from 2015. It is located in Denmark, and it primarily consists of space types such as classrooms, study zones, corridors, and offices. In each of these spaces, indoor air temperature T and CO_2 concentration C are measured through installed sensors. In addition, each space contains space heaters of specific capacities with equipped mechanical valves that control the water massflow. The position of these valves $u_v \in [0,1]$ is managed centrally by the Building Management System (BMS) with $u_v = 0$ meaning fully closed with no massflow and $u_v = 1$ meaning fully open with maximum massflow. The supply water temperature is kept constant at approximately 60 °C.

Part of
Pernigotto, G., Patuzzi, F., Prada, A., Corrado, V., & Gasparella, A. (Eds.). 2023. Building simulation applications BSA 2022. bu,press.
https://doi.org/10.13124/9788860461919

The building is also equipped with a weather station that measures outdoor air temperature T_o, longwave solar irradiance Φ_L, shortwave solar irradiance Φ_S, and wind speed. The supply air temperature setpoint is constant, at either 21-22 °C, depending on each space. The supply and exhaust airflows are controlled in each space by the supply and exhaust damper positions $u_d \in [0,1]$, with $u_d = 0$ meaning fully closed with no airflow, and $u_d = 1$ meaning fully open with maximum airflow. These damper positions are also managed through the BMS with Demand Controlled Ventilation (DCV), aiming at keeping the measured CO_2 concentration below 600 ppm. Finally, each space is also equipped with shades that are controlled by the BMS through a position parameter $u_{sh} \in [0,1]$, with $u_{sh} = 0$ meaning fully exposed with no shading, and $u_{sh} = 1$ meaning fully enclosed with maximum shading. The shades are controlled based on outdoor and indoor illuminance with a safety roll-up mechanism that sets $u_{sh} = 0$ at wind speeds higher than 15 m/s.

3. Methodology

3.1 Model Architecture

This work makes use of a specific type of ANNs called Recurrent Neural Networks (RNN) to model the transient temperature dynamics of a room. Specifically, Long Short-Term Memory (LSTM) networks were chosen due to their numerous demonstrations of adaptability and robustness in time-series black-box modeling, including indoor environment modeling (Fang et al., 2021; Mtibaa et al., 2020). LSTM models are a specific kind of RNN that were originally developed to deal with the vanishing and exploding gradient problem of traditional RNN models. A detailed explanation of the LSTM model is provided in the references (Hochreiter & Schmidhuber, 1997; Van Houdt et al., 2020).

To properly account for all phenomena that can significantly influence the energy balance of the room, it is very important to choose appropriate dynamic inputs for the model. In this work, the inputs are determined by considering the following energy transfer mechanisms:

- *Heat transfer by conduction through external surfaces* is considered by including indoor air temperature T, and outdoor temperature T_o as input.
- *Heat transfer by radiation* is considered by including longwave solar irradiance Φ_L, shortwave solar irradiance Φ_S, and the position of the shades u_{sh} in the model input.
- *Internal heat gains through occupancy* are included indirectly by including measured CO_2 concentration C, and damper position u_d. These inputs can, to a certain extent, represent occupancy due to the direct correlation between CO_2 concentration, ventilation airflow, and occupancy presence (Franco & Leccese, 2020).
- *Heat added by the space heater* is considered by adding the measured valve position u_v as input, which represents the water massflow. The supply water temperature is constant and does not therefore contribute as input.
- *Heat transfer by ventilation* is considered by adding the supply and exhaust damper positions u_d as inputs, which represent the airflow rates. The supply air temperature is constant and therefore does not contribute as input.

The model architecture is seen in Fig. 1 with inputs and outputs of the model. As shown, the model consists of two sequential LSTM models, A and B. All previously mentioned weather and sensor inputs from the previous timestep are fed to LSTM A. In addition, the LSTM also receives the cell state vector $c_{A,t-1} \in \mathbb{R}^n$ and hidden state vector $h_{A,t-1} \in [-1,1]^n$, where n is a hyperparameter that determines the size of these vectors. These two state vectors are an integral part of LSTM models. which essentially dictate the state of the system modeled during a given timestep. LSTM A outputs $c_{A,t}$ and $h_{A,t}$, which represents the updated state vectors.

LSTM B has $n = 1$ and is only given three inputs, the cell state $c_{B,t-1} \in \mathbb{R}$, the hidden state $h_{B,t-1} \in [-1,1]$, and the hidden state vector $h_{A,t}$. LSTM B outputs $c_{B,t}$ and $h_{B,t}$, which represent the updated state vectors of LSTM B. The training task is then to find an optimal set of parameters in LSTM A and B to minimize the error between $h_{B,t}$ and the chosen prediction target over multiple sequences of data.

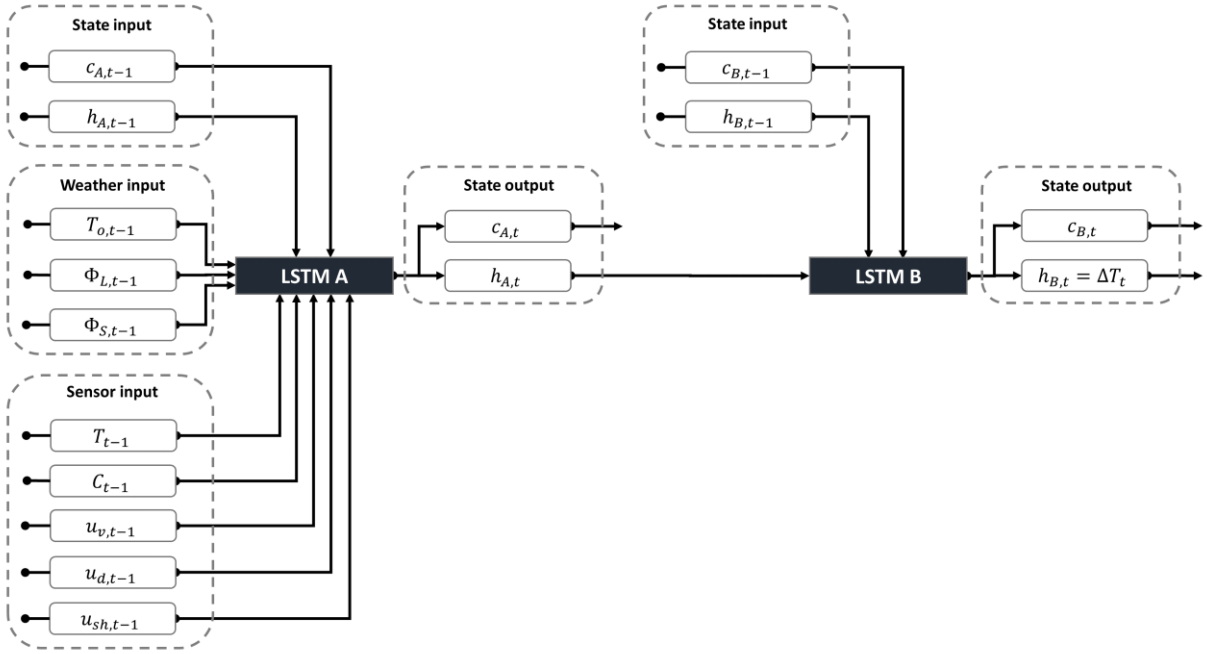

Fig. 1 – Model architecture with inputs and outputs of the sequential LSTM models, A and B

In the studies reviewed, the prediction target of the learning algorithms employed was indoor temperature in all cases. During initial testing, it was found that this configuration yielded good prediction results when testing with historical data as input. However, it was found that the models generalized poorly when fed with custom inputs instead of historical data, e.g., when using the model for setpoint control, as shown later in Section 4.2. When collecting operational data from a building, the actuation signals, e.g., valve positions, are typically directly correlated with indoor temperature through a thermostat with a simple control law, e.g., in the form of a Proportional (P) or Proportional Integral (PI) controller. Therefore, if the prediction target is temperature, the model will likely overfit the specific mode of operation that is reflected in the historical data used to train and test the model. We hypothesize that the model essentially learns to map the inverse control law of the thermostat instead of the actual thermal physics of the room. In this work, we are proposing that the model should predict the indoor temperature change ΔT instead of the actual temperature value, as this disrupts the direct correlation that is otherwise present between input and output.

3.2 Data Preprocessing

The dataset was constructed with all the weather and sensor readings introduced for 76 rooms in the case study building at a 10- minute interval for two years spanning January 1st 2018 to December 31st 2019. The raw data were pretreated and validated to ensure that proper and clean data were used. Following this, all inputs were min-max scaled between -1 and 1. After preparing the dataset, 24-hour sequences of 144 timesteps with no missing data were selected to form a collection of sequences for each space. In Fig. 2, the distribution of data sequences available among spaces is shown monthly. As shown, the number of sequences varies between rooms. The month with the most data is January, with a median of about 6000 sequences, while the month with the least data is August, with a median of about 3000 sequences.

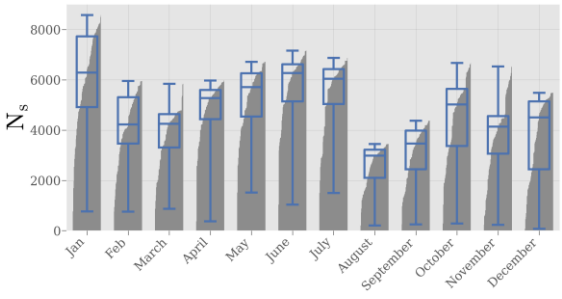

Fig. 2 – Number of sequences available per space for each month

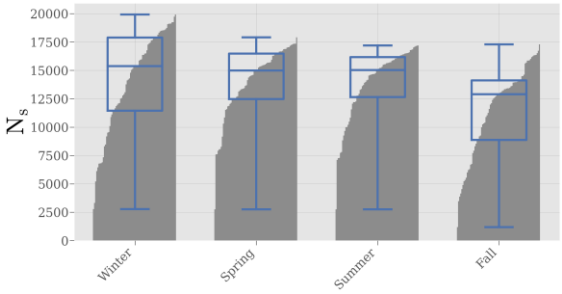

Fig. 3 – Number of sequences available per space for each season

In Fig. 3, the distribution is shown on a seasonal level. Here, the dataset generally appears to be more balanced with less variation. Hence, it is expected that the constructed datasets have enough diversity to cover most of the seasonal variance in the operation and thermal dynamics of the building.

3.3 Training and Testing Method

After preprocessing, the data is split into training, validation, and testing data sets with the splits [2/3, 1/6, 1/6], respectively. These data splits were carefully designed to ensure no overlap between sequences in the three data sets, while ensuring that seasonal and monthly variations are reflected in each dataset. For training, the machine-learning library Pytorch was used with Stochastic Gradient Descent (SDG) as optimizer, a momentum of 0.9, a batch size of 32, and a learning rate of 10^{-1}. The size n of cell state c_A and hidden state h_A was set to 20 for all space models.

During training, the model loss was evaluated on the validation dataset and saved at every 64th gradient update. At every 3000th gradient update, a copy of the model was saved, along with an average of the last 100 validation loss evaluations. After 100,000 gradient updates, the model with the lowest saved validation loss was selected. If the model had not converged after 100,000 gradient updates, the procedure was repeated up until 400,000 gradient updates. This approach was used to train models for all 76 spaces. To properly test the trained space models, two modes of operation were presented.

The first mode was aimed at assessing the prediction accuracy of the space models given historic inputs. Here, the developed models were employed in a closed-loop configuration, as shown in Fig. 4, where future temperature predictions were based on past predictions. Perfect forecasting was assumed by feeding historical data for all weather and sensor inputs, except for the indoor temperature.

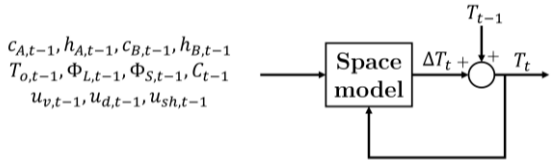

Fig. 4 – Closed-loop configuration designed to forecast indoor temperature for an arbitrary number of timesteps

The model then, for each timestep, predicted the temperature change to obtain the indoor temperature of the next time step, which was fed back to the model. This could be repeated as long as historical inputs were available. By repeating this for all timesteps in a simulation period, the produced temperature profile could then be compared with the actual measured temperature profile of the room to assess the prediction accuracy of the space model.

The second mode of operation was aimed at testing whether the developed space models generalize well enough to provide reasonable predictions under unseen operational conditions. Here, custom inputs were thus be fed to the model to observe the response. This was a very important property that made it possible to use the model for testing different operational strategies or what-if scenarios and their influence on indoor comfort in a safe environment.

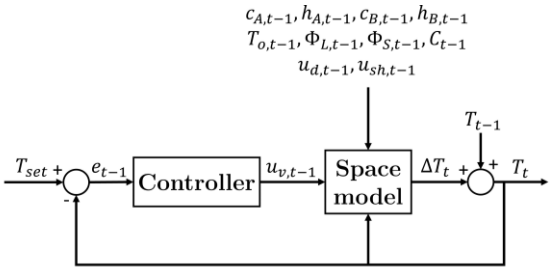

Fig. 5 – Closed-loop configuration designed for temperature set-point control through the valve position of the space heater

In this work, two model inputs were considered for this purpose; u_v and u_{sh}. In one simulation, the space heater valve position input u_v was constructed by implementing the space models in a closed-loop configuration for temperature setpoint control, as shown in Fig. 5. Here, all inputs were historical except for u_v and temperature T. The implemented controller was a simple proportional controller that each timestep scales the input signal $u_{v,t-1}$ proportionally to the error $e_{t-1} = T_{set} - T_{t-1}$, where T_{set} was the temperature setpoint in the room. In another simulation, the shades position input u_{sh} was constructed by implementing a simple predetermined schedule that operates based on the time of the day. All inputs except for u_{sh} and T were thus historical.

4. Results and Discussion

4.1 Quantitative Performance Assessment

First, the quantitative model performance was evaluated by using the first mode of operation as described in Section 3.3. Here, the Mean Absolute Error (MAE) between measured indoor temperature and predicted temperature was calculated for each space model across all 24-hour sequences in the test dataset. Fig. 6 shows the performance of each trained space model. Specifically, it shows the relationship between MAE, Standard Deviation of prediction targets $\sigma(\Delta T)$, and number of sequences N_s. The three marked space models were used for a qualitative performance assessment in Section 4.2. As seen in the figure, most of the space models (~86 %) achieve MAE values below 0.5 °C, which is lower than the required measurement accuracy of temperature sensors of ± 0.5 °C (ISO, 1998).

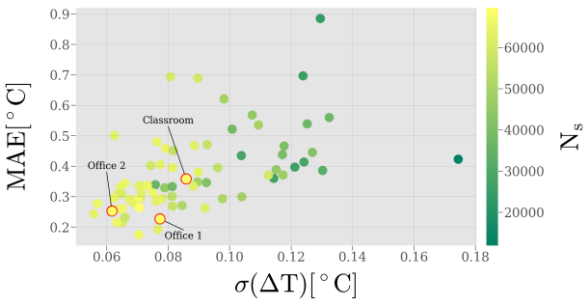

Fig. 6 – Relationship between Mean Absolute Error MAE, Standard Deviation of prediction targets $\sigma(\Delta T)$, and number of sequences N_s. The three space models chosen for qualitative assessment are marked

The best performing space model represents an office with $N_s = 66297$ and $\mathrm{MAE} = 0.17$, while the worst performing model is of a classroom with $N_s = 24210$ and $\mathrm{MAE} = 0.88$. As seen from the colormap in Fig. 6, there seems to be a negative correlation between the prediction error and the number of sequences available, agreeing with the general notion in machine-learning, that more data yields lower prediction error and better model generalization. Therefore, it is expected that the poor-performing space models could attain similar performance with more data. Furthermore, there seems to be a positive correlation between prediction error and the variation observed for the prediction target ΔT. This means that the prediction error will be higher for datasets that have a more fluctuating temperature profile. This is to be expected, as a fluctuating indoor temperature is generally harder to predict than a steady temperature.

4.2 Qualitative Performance Assessment And Applications

To provide a qualitative performance assessment of the models developed, three case study spaces were selected, one classroom and two offices, as also marked in Fig. 6. For each of these space models, a winter period and a summer period of 24 hours were chosen to evaluate how the space models perform under different ambient conditions. The models were first simulated for these periods using the first mode of operation, as explained in Section 3.3.

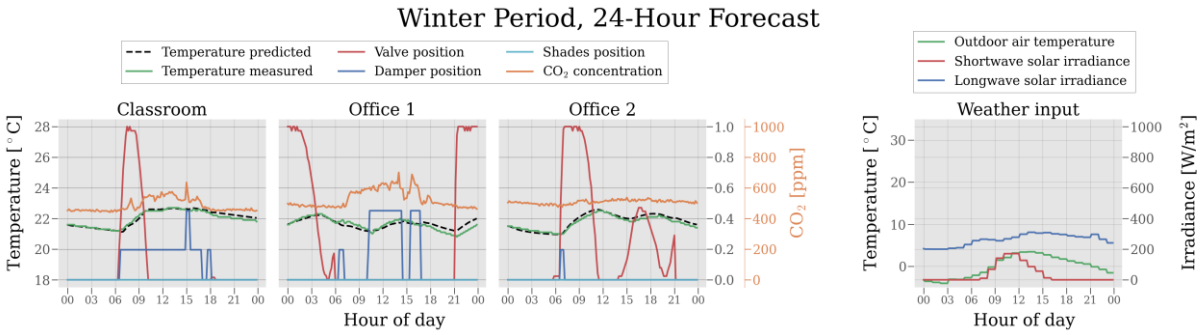

Fig. 7 – 24-hour temperature forecast in a winter month compared with actual measured temperature. The weather input is shown on the plot furthest to the right, while the individual inputs for the three selected space models are shown on the plots to the left

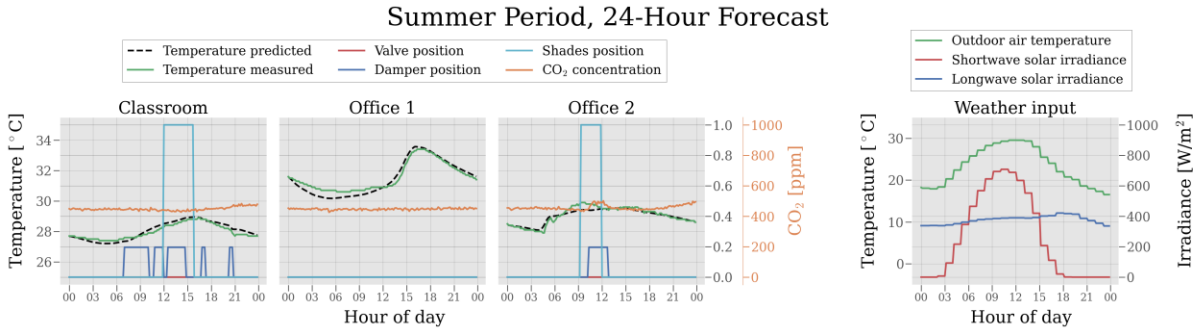

Fig. 8 – 24-hour temperature forecast in a summer month compared with actual measured temperature. The weather input is shown on the plot furthest to the right, while the individual inputs for the three selected space models are shown on the plots to the left

For the winter simulation, the results are shown in Fig. 7, with the weather inputs in the plot furthest to the right, while the results for each of the three chosen rooms can be seen on the left. As seen, all three space models accurately predict the indoor temperature, although Office 1 seems to slightly overestimate the temperature during the last 6 hours. The space heater valve position has a clear significance in all three spaces, where the general trend is that $u_v = 0$ results in decreasing room temperature, while $u_v > 0$ results in increasing room temperature. The models are also able to account for heat gains and heat losses associated with occupancy and ventilation. This is mostly seen in the Classroom and Office 1 in the period from 09:00 to 15:00, where the CO_2 concentration rises above 600 ppm and the dampers are positioned at around 50 %. Here, the model correctly predicts that the indoor temperature increases, although the space heater is not in operation. The shades are all rolled up ($u_{sh} = 0$) as they have no desirable effect during winter.

Moving to the simulation results for the summer period, the results are shown in Fig. 8. The weather data inputs are again seen on the plot furthest to the right, where the ambient temperature and irradiance levels are much higher compared with the winter period. As seen, this has a significant impact on the predicted and actual temperatures in the three spaces, especially for Office 1, where temperatures are above 30 °C during the whole 24 hour-period. The cause for large differences in both shape and peaks of the temperature profiles was different orientation, shading, and geometrical properties of the spaces. As seen, all three space models have learned to correctly account for these properties and provide accurate predictions for all 24 hours, although Office 1 slightly underestimates the temperature during the first 12 hours of the period. From the CO_2 levels, which have a very constant profile of around 450 ppm, the occupancy appears to be close to zero. However, this is expected during July and August, where the students are on summer leave. The shades are, to some extent, utilized in the Classroom and Office 2. However, it is expected that the high temperatures could be mitigated even more, by increasing the duration of the shades being rolled down ($u_{sh} = 1$). This will be investigated further in the following analysis.

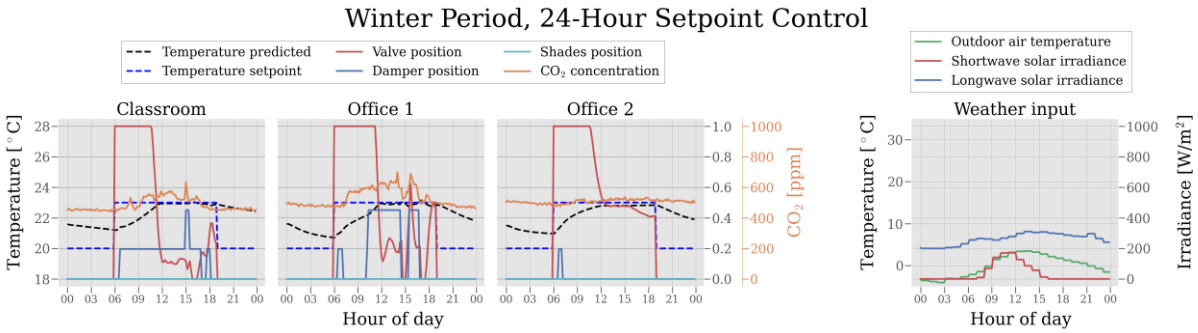

Fig. 9 – 24-hour setpoint control in a winter month. The weather input is shown on the plot furthest to the right, while the individual inputs for the three selected space models are shown on the plots to the left along with the predicted temperature and the setpoint

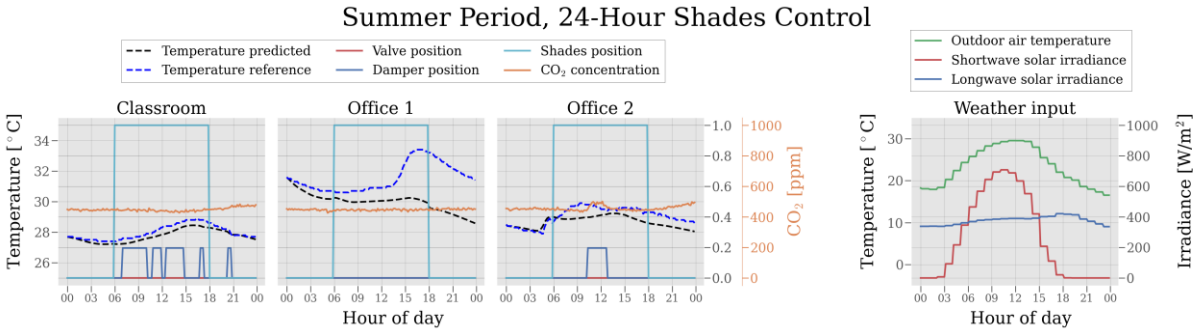

Fig. 10 – 24-hour shades control in a summer month. The weather input is shown on the plot furthest to the right, while the individual inputs for the three selected space models are shown on the plots to the left along with the predicted temperature. The original measured temperature in this period is shown as a reference

As just demonstrated, the models perform well when predicting temperatures under historical conditions. However, to demonstrate the applicability of the models, the second mode of operation, as presented in Section 3.3, was employed. The temperature setpoint control was implemented for the same winter period as shown in Fig. 7, while the shades control was implemented for the same summer period as shown in Fig. 8.

The results for the temperature setpoint control are shown in Fig. 9. As seen, the setpoint was varied, depending on the time of day. From 06:00 to 18:00, the setpoint was 23 °C, while from 18:00 to 06:00, the setpoint was 20 °C. The controller adapted to these setpoint signals by varying the valve position accordingly. For all three spaces, the valve was closed during the night, where the temperature was allowed to decrease. At 06:00, when the setpoint was raised to 23 °C, the valve was opened, and the indoor temperature increased until around 12:00, where the indoor temperature in all three rooms reached the specified setpoint. At this point, the valve position was operated between 0 and 1 in an attempt to keep the indoor temperature at the setpoint. For Office 2, it was noticed that the temperature and valve position profile was smoother compared to the Classroom and Office 1, which had small fluctuations in temperature. However, as seen, the Classroom and Office 1 also had more disturbances in the form of varying CO_2 levels and ventilation airflows. Despite these disturbances, the simple controllers managed to keep the temperature at 23 °C in all three space models until 18:00, where the setpoint was again decreased to 20 °C. Here, the valve was again shut, and the temperature started to decrease.

Moving on to shades control during the summer period, the results are shown in Fig. 10. Here, the shades were rolled down ($u_{sh} = 1$) from 06:00 to 18:00 and rolled up ($u_{sh} = 0$) from 18:00 to 06:00. This had a significant effect on the predicted indoor temperature when compared with the original measured temperature (temperature reference), where the duration of shading was very limited. The

effect was very clear for Office 1, where the peak at 17:00 shifted from around 33 °C to 30 °C. The Classroom and Office 2 also had a reduced temperature response, although not as significant as Office 1.

5. Conclusion

In this work, a fully automated and scalable approach for temperature forecasting in buildings was presented and assessed. The presented method relied on ANNs in the form of two sequential LSTM models to predict temperature change within a room, given weather data and sensor inputs such as space heater valve positions, damper positions, shades positions, and CO_2 concentration. Hence, the developed approach needed no prior information about the building such as geometry, material properties, design data, etc.

The methodology developed was implemented considering 76 rooms of an 8500 m^2 university building in Denmark with 86 % of the rooms achieving a Mean Squared Error of less than 0.5 for 24-hour forecasting. The difference in prediction performance between space models was explained by differences in the amount of data available. However, more work is needed to identify more robust criteria concerning the amount, type, and quality of data that is needed to obtain accurate space models. The applicability of the models was demonstrated by implementing three selected models in a closed-loop setpoint control configuration for a 24-hour winter period. Additionally, different shading schedules were also explored to show their impact during a 24-hour summer period.

In line with the emerging initiatives toward digitalization of the building sector, building digital twins has promising technical and economic impacts. In this context, a fully scalable and automated energy modeling approach is vital, so that these twins can provide a robust, generic, and effective solution for various applications in the building sector. The modeling approach proposed in this study serves as a core for future building digital twin development and could be used as a backbone for various automated services, including performance monitoring, scenario assessment, and operational management.

Acknowledgement

This work was carried out under the 'Twin4Build: A holistic Digital Twin platform for decision-making support over the whole building life cycle' project, funded by the Danish Energy Agency under the Energy Technology Development and Demonstration Program (EUDP), ID number: 64021-1009.

References

Alawadi, S., D. Mera, M. Fernández-Delgado, F. Alkhabbas, C. M. Olsson, and P. Davidsson. 2022. "A comparison of machine learning algorithms for forecasting indoor temperature in smart buildings." *Energy Systems* 13: 689–705. doi: https://doi.org/10.1007/s12667-020-00376-x

Fang, Z., N. Crimier, L. Scanu, A. Midelet, A. Alyafi,, and B. Delinchant. 2021. "Multi-zone indoor temperature prediction with LSTM-based sequence to sequence model." *Energy and Buildings* 245: 111053. doi: https://doi.org/10.1016/j.enbuild.2021.111053

Franco, A., and F. Leccese. 2020. "Measurement of CO2 concentration for occupancy estimation in educational buildings with energy efficiency purposes." *Journal of Building Engineering* 32: 101714. doi: https://doi.org/10.1016/j.jobe.2020.101714

Hochreiter, S., and J. Schmidhuber. 1997. "Long Short-term Memory." *Neural computation* 9(8): 1735-1780.

ISO. 1998. *ISO 7726 - Ergonomics of the Thermal Environment - Instruments for Measuring Physical Quantities*.

Mtibaa, F., K.-K. Nguyen, M. Azam, A. Papachristou, J.-S. Venne, and M. Cheriet. 2020. "LSTM-based indoor air temperature prediction framework for HVAC systems in smart buildings." *Neural Computing and Applications* 32: 17569–17585. doi: https://doi.org/10.1007/s00521-020-04926-3

Van Houdt, G., C. Mosquera, and G. Nápoles. 2020. "A Review on the Long Short-Term Memory Model." *Artificial Intelligence Review* 53(8): 5929-5955. doi: https://doi.org/10.1007/s10462-020-09838-1