

Short-term HVAC load prediction using Temporal Graph Convolutional Networks

Dennis Grauer¹ , Miguel Guzman-Merino¹  and Nils Soeren Krause¹ 

¹Artificial Intelligence for sustainable construction, University of Rostock, Germany

E-mail(s): dennis.grauer@uni-rostock.de, miguel.merino@uni-rostock.de, nils.krause@uni-rostock.de

Abstract: Machine learning approaches for building performance prediction have gained popularity in recent years. Backed by the higher availability of Internet of Things sensor data, researchers have created well-performing data-driven prediction models for indoor climate conditions, energy usage and thermal load. These models use features highly correlated to their prediction target to score high accuracies, but this often leads them to overfit, struggling with unseen data and not producing outputs complying with physical laws. In this paper, we present a Temporal Graph Convolutional Network (TGCN) approach, seeking to overcome these limitations. The model is intended to predict short-term Heating, Ventilation and Cooling (HVAC) load for a office building. Instead of using few features, we incorporate a high number of sensors into the models underlying graph structure. We validate this approach on a three year data set of an office building. The data set contains data from the 2020 Covid-19 pandemic, where occupational patterns of the building change and the correlations governing the HVAC load shifts. The TGCN model is trained with pre-pandemic data and confronted with the pandemic situation. The results show, that the model is able to maintain its performance on the prediction task for out-of-distribution data.

Keywords: Graph neural networks, building energy, machine learning



DOI: 10.18154/RWTH-CONV-254882. Published in the conference proceedings of the 36. Forum Bauinformatik 2025, Aachen, Germany, © 2025 The copyright for this article lies with the authors. This publication, except for quotations and otherwise indicated parts, is licensed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.

1 Introduction

Although predictions about building performance traditionally relied on physical simulations, the wider spread of Internet of Things (IoT) sensors has encouraged researchers to develop data driven prediction models. While these data driven approaches offer a high accuracy and require less configuration and knowledge about the buildings characteristics, they fall short of physics-based models in two categories: (i) as these models are purely based on correlation, they tend to overfit on correlation and struggle with capturing causation, particular in unseen situations, (ii) they are considered black-box models, as the interpretability of their predictions falls short to physic-based models [1]. In this paper, we want to focus on the first problem.

Neural networks conducting their learning on graph structures, commonly referred to as Graph Neural Networks (GNNs), have received extensive attention in the field of building (energy) management.

When their graphs encode physical relations, they can be categorized as a kind of physics-informed machine learning.

We want to explore the potential of graph neural networks to deal with a large, physics based graph structure. Therefore, we build a Temporal Graph Convolutional Network (TGCN) for a real-world office building dataset, aiming to predict short-term HVAC load.

The graph is based on the physical relations governing the HVAC load and utilizes a high number of IoT sensors to make its predictions. We seek to find an approach, that performs well on unseen data by learning causation instead of correlation. In the following section, we will first look at the relevant literature. We then define our problem and the model. After introducing the data set and the graph representation shaping our description, we proceed with presenting the results. This paper ends with a discussion on the results and an outlook on potential future research directions.

2 State of the art

Data-driven performance prediction approaches started early with works like [2], [3] to utilize ML methods such as long short-term memory (LSTM) to predict buildings thermal loads [4], [5], [6], [7], [8] or indoor climate conditions [9], [10], [11], [12], [13]. While all of these models achieved accurate results, they are still subjected to the problem of purely data-driven approaches: they often fail to make predictions complying with physical laws or lack the ability to generalize [1]. These shortcomings have motivated the development of physics-informed machine learning models like GNNs. GNNs are able to include physical knowledge into the learning process, since they are based on graphs. A basic graph consist of a set of nodes, connected by edges. Extending nodes with corresponding node features and correctly capturing causal, physical or spatial relations with edges, these structures offer a powerful way of encoding physical laws.

Various research has been conducted on employing GNNs for load prediction tasks on buildings with HVAC systems. [14] use a framework combining LSTM and GNN technology to provide heat load forecasting. Outdoor temperature, radiant flux, outdoor humidity content, and wind speed are used as features [15] to build a set of control relationship graphs based on spatial structure, control characteristics and equipment attributes and employ a GNN combined with an Informer model to predict an office buildings cooling load over an entire cooling season. [16] use a spatial temporal graph attention network to predict heating load for multiple thermal zones of an office building using outdoor features like solar radiation, temperature, hour of the day, wind speed and also indoor features like occupancy temperature differences of zones and between zones and outside.

Most of these approaches use a mixture of IoT data, outside climate influence and building usage data encoded in features like week day or hour of the day. Although these features are highly correlated to the target variable, we want to enable models to also deal with new situations. For example when there are significant changes in occupants behavior, time features quickly decrease in importance. To overcome this limitation and make use of a larger number of available sensors, we develop a model based on thermal zone sensors as well as HVAC internal data and outdoor climate.

3 Methodology

3.1 Problem definition

We define our approach as a graph learning problem on a building graph representation \mathcal{G} that captures the thermodynamic relationships influencing the HVAC load.

Buildings thermodynamics vary over time, with some connections in the graph only being active in specific situations. Thus we use a temporal approach to capture the correct dynamics. In this approach, each timestep t_i has its own adjacency matrix A_{t_i} , which is derived from a timesteps graph $G_{t_i} = \{(\mathcal{V}_{stat}), (\mathcal{E}_{stat} + \mathcal{E}_{t_i})\}$. \mathcal{V}_{stat} and \mathcal{E}_{stat} are the static node and edge sets, shared across all timesteps and holding connections that do not change over time. \mathcal{E}_{t_i} is a timesteps dynamically computed modified edge weight. \mathcal{V}_{stat} denotes a set $\{v_1, v_2, \dots, v_i\}$ of nodes representing a subset of all available sensors $\mathcal{V}_{stat} \subseteq \mathcal{S}$ in the data set. \mathcal{E}_{stat} denotes a set of directed edges $\{e_1, e_2, \dots, e_i\}$, each edge $e_{i,j} = (v_i, v_j)$ representing a connection directed from node v_i to node v_j .

By appending a feature vector $X_t \in \mathbb{R}^P$ to each node, with P the number of previous measurements, we receive the feature matrix $X \in \mathbb{R}^{(N \times P \times T)}$ with the number of nodes N and the number of timesteps T . Using the previously constructed adjacency matrices A , we define our TGNN as:

$$\tilde{y}_t = f(X, A_{t_i}) = \sigma(\hat{A}_{t_i} \text{ReLU}(\hat{A}_{t_i} X W^{(0)}) W^{(1)}), \quad (1)$$

where $\hat{A}_{t_i} = \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}}$ is a preprocessing step, $\tilde{A}_{t_i} = A_{t_i} + I_N$ (with I_N the unity matrix) is a matrix that considers the features of the nodes for which the learning is conducted, \tilde{D} is a degree matrix, where $\tilde{D} = \sum_j \tilde{A}_{ij}$, $W^{(0)}$ and $W^{(1)}$ represent the weight matrix in the first and second neighbourhood, and σ and ReLU represent sigmoid and ReLU activation functions.

3.2 Dataset and Scenario

To verify our approach, we utilize the data set presented in [17]. The data set offers time series data, ranging from January 2018 to December 2020. The data records contain energy use data outdoor and indoor environmental data and HVAC operational data. We will focus in particular on the data associated with the HVAC and the thermal conditions.

All data comes from a four floor building called Building 59, located in Berkeley, California. It contains 10.400 m² of air conditioned office space. The building is separated into two parts, the northern and the southern part with both parts containing 57 overall thermal zones. Since there is more data available for the southern part, we will use this part to build our models on.

The southern part of Building 59 is conditioned by two roof top air handling units (RTUs), serving the 20 thermal zones located in the building part. The loads of both RTUs are summarized to the HVAC load for the southern part.

In our scenario, \mathcal{E}_{t_i} contains weighted connections between a thermal zone j 's temperature and its cooling and heating setpoint. Both setpoints only influence a zone's climate if the corresponding system is active: the cooling setpoint if the HVAC fan is running, the heating setpoint if the heating valve is opened. Since there is also a quantitative difference in the influences depending on the valve opening/ fan speed, we map these values as the edge weights between the setpoints and the

zones temperature and scale them w_j to values $0 \leq w_j \leq 1$, with $w_j = 1$ indicating a zones fan is running at maximum speed or its valve is fully opened and $w_j = 0$ indicating a closed valve/not running fan. Computing this per timestep and thermal zone, we end up with the set of non-binary adjacency matrices A_{t_i} .

An interesting aspect of the data range from 2018 to 2020 is that it includes the Covid-19 pandemic. This poses an interesting test scenario for our prediction model: by training it on the data of 2019 and additionally testing it on the pandemic-influenced data, we can evaluate its ability to deal with out-of-training scenario. In this time, the usage pattern of Building 59 changed: the buildings occupancy and electrical energy usage decreased, but due to the low outside temperatures and a higher requirement of ventilation [17], the HVAC load does not decrease significantly. Figure 1 shows how the pattern of electrical usage change, as the occupancy drops and the consumption does not surge during the usual working hours.

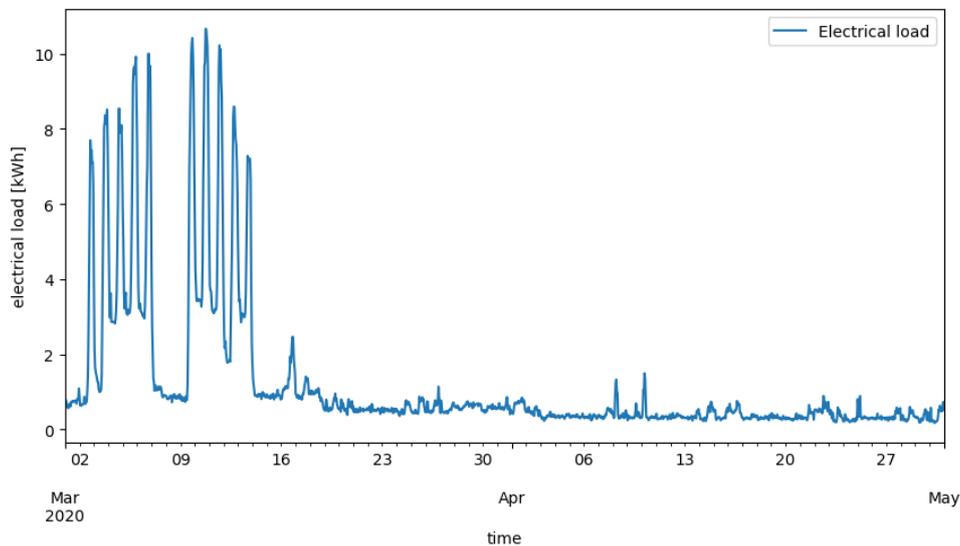


Figure 1: Hourly electrical energy usage at beginning of pandemic conditions

To additionally visualize that the pre-pandemic parts of the data set and the pandemic influenced parts differ and require a more causal model, we take a look at the Pearson correlation coefficients of features influencing HVAC load for both time series in Table 1. While outdoor environmental features like air temperature and relative humidity decrease in correlation, due to the higher ventilation requirements HVAC operational data gains influence on the load.

3.3 Graph representations and training

For the graph representation, we first encode the static part, consisting of nodes constantly contributing to HVAC load or influencing zone climate. In our scenario, this includes the sensors of both RTUs (003 and 004) as well as the fixed relations between the outside air temperature and the zones temperatures. Figure 2 shows a simplified visualization of the graph for RTU003.

The HVACs energy consumption is mainly governed by temperature difference between the supply air temperature and the return air and the fan speeds. We capture these relations in the graph. Using the various air temperature, setpoint and fan speed sensors inside the HVAC, we build the relations

Table 1: Pearson correlation coefficients for both parts of the data set

| Feature | Pre-pandemic | Pandemic part |
|------------------------------------|--------------|---------------|
| Outdoor air temperature | 0.795 | 0.585 |
| Outdoor relative humidity | -0.706 | -0.341 |
| RTU003 outdoor air damper position | 0.137 | 0.354 |
| RTU004 outdoor air damper position | -0.098 | 0.447 |
| RTU003 supply fan speed | 0.510 | 0.934 |
| RTU004 supply fan speed | 0.510 | 0.935 |

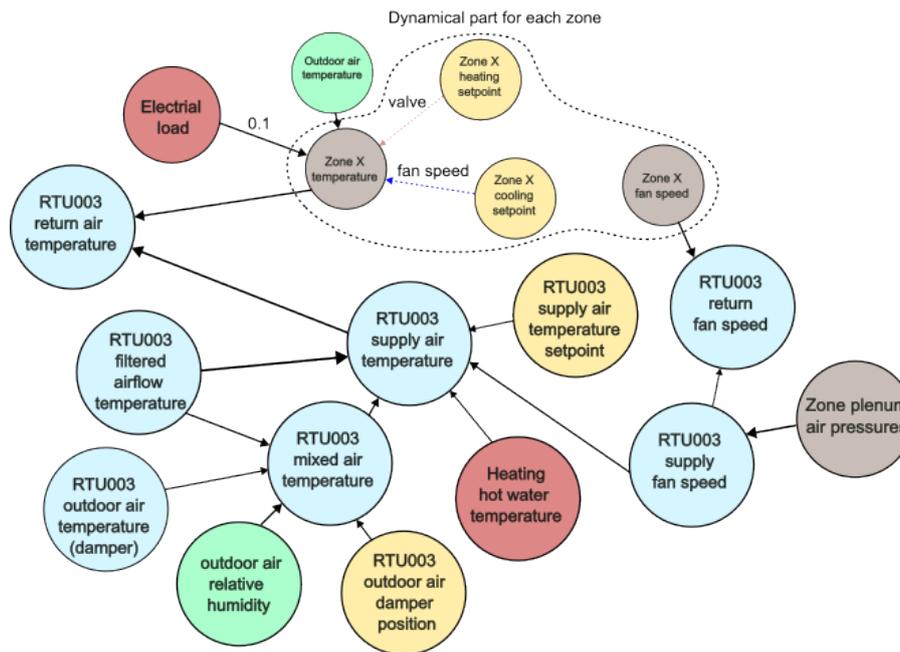


Figure 2: Graph representation for RTU003, including HVAC operational data (blue), setpoints (yellow), heat gains from electrical equipment and heating (red), thermal zone sensors (grey) and outdoor environmental data (green)

inside the HVAC. To include the influence of the thermal zones, we connect each zone to the return air temperature of its RTU. Because occupancy data is only available for short periods in the dataset, we use the electrical load as a substitution for internal heat gain. Since the internal gains do not influence a zone temperature in the same amount as other thermal relations, we weight this edge with a fixed value of 0.1.

Building up the static part of the graph for both RTUs and the thermal zones, we end up with 66 nodes connected by 107 edges. If all 20 thermal zones have opened valves and running fans, the number of edges can extend by 40 for a timestep. This approach gives us a solid representation of the thermophysics responsible for HVAC load.

The data is sampled in a 15 min interval. For training, we build a feature vector with $T = 8$ timesteps for each node, giving us a history of two hours. Our training data ranges from April 2019 to September 2019. Each feature vector is standardized using a standard scaler. We use a two layer TGCN for training, running with 50 epochs and optimized with the Adam optimizer, the learning rate is set to

0.01. A 80-20 % split is applied to separate training and test data (resulting in the pre-pandemic test data part).

4 Results

The results of our training are shown in Table 2. We use Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and R^2 as metrics to evaluate our models performance.

The model trained on the pre-pandemic data is able to maintain its performance, although decreasing in MAE and RMSE. This is also stated by the true versus predicted plot Figure 3. Although the correlations have shifted compared to the training, the model holds enough data points to not overfit on single features and anticipates the trend of the HVAC load data. It is especially remarkable, that the model accurately predicts load peaks, but struggles in periods with lower loads, underestimating the true values.

While R^2 score for the pre-pandemic data indicates a good regression of the problem, it significantly decreases on the out-of-distribution data.

Overall, the model demonstrates a strong ability to generalize from the high number of features it is given. Despite the decreasing performance on the pandemic data, it approximates the true values well and especially the trends and peaks. This indicates that the model has learned underlying patterns rather than overfitting to the training data or pure correlation, enabling it to adapt to previously unseen conditions.

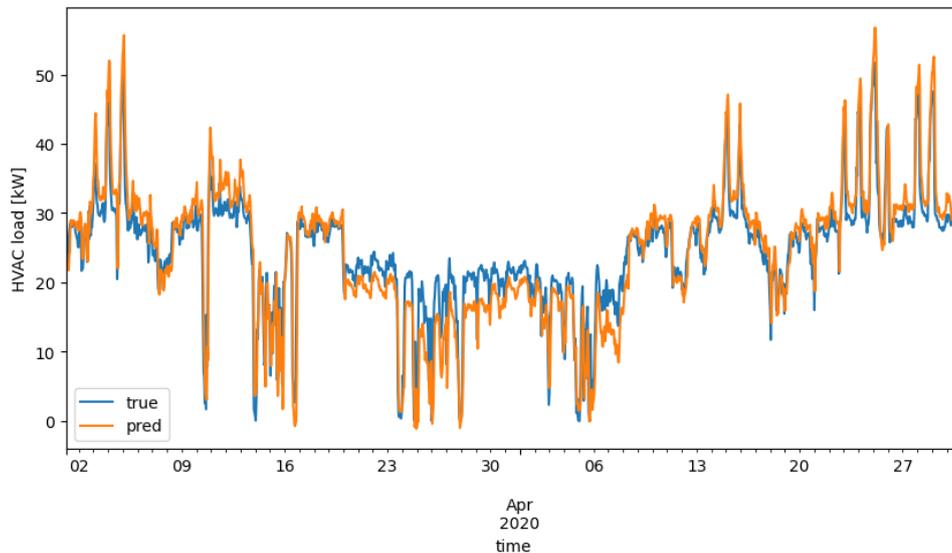


Figure 3: True vs predicted plot for pandemic data

Table 2: Model metrics for testdata and Covid 19 additional test data

| Metric | Pre-pandemic test data | Pandemic test data |
|--------|------------------------|--------------------|
| MAE | 2.223 | 2.818 |
| RMSE | 3.049 | 3.817 |
| R^2 | 0.941 | 0.776 |

5 Discussion and Outlook

The approach of incorporating a high number of IoT sensor into a temporal graph presented in this paper has shown to be resilient to changes in input data and maintain a high level of accuracy. Although the prediction accuracy falls short to other approaches presented in Section 2, the encoded physical knowledge lets the model stay on a similar level, even in unseen situations. Since machine learning approaches in practice often struggle to be accepted by decision-makers due to their black-box nature and their poor performance on unseen data [1], GNN/GCN approaches like the one presented in this paper are a step to more trustworthy predictions.

Despite the promising results, there is still a need for further research. As previously stated, interpretability of predictions also contributes to building trust in machine learning models. The approach presented here still remains a black-box, thus it is worth investigation the potential of explainability techniques for GNNs, like the ones presented in [18], [19].

Another area worth investigating is further automating the configuration of GNNs by automatically building the graph representation. Even though the approach presented here requires less knowledge about the buildings characteristics than a physical simulation, a thorough domain knowledge on building physics is required. There have been effort to automate configuration using semantic math parsers [20] or LLMs [21], but there is not yet a reliable framework for the process.

However, the approach benefits from the high number of features, but these are not always available. The approach is yet to prove that it can be applied to other datasets than the one presented here.

References

- [1] Z. Chen, F. Xiao, F. Guo, and J. Yan, "Interpretable machine learning for building energy management: A state-of-the-art review", *Advances in Applied Energy*, vol. 9, 2023.
- [2] A. E. Ben-Nakhi and M. A. Mahmoud, "Cooling load prediction for buildings using general regression neural networks", *Energy Conversion and Management*, vol. 45, no. 13-14, 2004.
- [3] Q. Li, Q. Meng, J. Cai, H. Yoshino, and A. Mochida, "Predicting hourly cooling load in the building: A comparison of support vector machine and different artificial neural networks", *Energy Conversion and Management*, vol. 50, no. 1, 2009.
- [4] J. Hu et al., "Thermal load prediction and operation optimization of office building with a zone-level artificial neural network and rule-based control", *Applied energy*, vol. 300, 2021.
- [5] S. Sholahudin and H. Han, "Simplified dynamic neural network model to predict heating load of a building using taguchi method", *Energy*, vol. 115, 2016.
- [6] Z. Wang, T. Hong, and M. A. Piette, "Building thermal load prediction through shallow machine learning and deep learning", *Applied Energy*, vol. 263, 2020.
- [7] A. Li, F. Xiao, C. Zhang, and C. Fan, "Attention-based interpretable neural network for building cooling load prediction", *Applied Energy*, vol. 299, 2021.
- [8] M. Zhou et al., "Issa-Istm: A new data-driven method of heat load forecasting for building air conditioning", *Energy and Buildings*, 2024.

- [9] F. Elmaz, R. Eyckerman, W. Casteels, S. Latré, and P. Hellinckx, “Cnn-lstm architecture for predictive indoor temperature modeling”, *Building and Environment*, vol. 206, 2021.
- [10] N. Attoue, I. Shahrour, and R. Younes, “Smart building: Use of the artificial neural network approach for indoor temperature forecasting”, *Energies*, vol. 11, no. 2, 2018.
- [11] B. Delcroix, J. L. Ny, M. Bernier, M. Azam, B. Qu, and J.-S. Venne, “Autoregressive neural networks with exogenous variables for indoor temperature prediction in buildings”, in *Building simulation*, Springer, vol. 14, 2021.
- [12] E. Arulmozhi, J. K. Basak, T. Sihalath, J. Park, H. T. Kim, and B. E. Moon, “Machine learning-based microclimate model for indoor air temperature and relative humidity prediction in a swine building”, *Animals*, vol. 11, no. 1, 2021.
- [13] N. Shakhovska, L. Mochurad, R. Caro, and S. Argyroudis, “Innovative machine learning approaches for indoor air temperature forecasting in smart infrastructure”, *Scientific Reports*, vol. 15, no. 1, 2025.
- [14] X. Jia, H. Song, X. Nan, and X. Cai, “Prediction of short-term heat load of office buildings based on gnn-lstm modeling”, in *2024 6th International Conference on Frontier Technologies of Information and Computer (ICFTIC)*, IEEE, 2024.
- [15] M. Zou, W. Huang, J. Jin, B. Hu, and Z. Liu, “Deep spatio-temporal feature fusion learning for multi-step building cooling load forecasting”, *Energy and Buildings*, vol. 322, 2024.
- [16] Y. Jia, J. Wang, M. Reza Hosseini, W. Shou, P. Wu, and C. Mao, “Temporal graph attention network for building thermal load prediction”, *Energy and Buildings*, vol. 321, 2024.
- [17] N. Luo et al., “A three-year dataset supporting research on building energy management and occupancy analytics”, *Scientific Data*, vol. 9, no. 1, 2022.
- [18] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, “Gnnexplainer: Generating explanations for graph neural networks”, *Advances in neural information processing systems*, vol. 32, 2019.
- [19] H. Yuan, H. Yu, S. Gui, and S. Ji, “Explainability in graph neural networks: A taxonomic survey”, *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 5, 2022.
- [20] A. Ba, K. Lynch, J. Ploennigs, B. Schaper, C. Lohse, and F. Lorenzi, “Automated configuration of heterogeneous graph neural networks with a semantic math parser for iot systems”, *IEEE Internet of Things Journal*, vol. 10, no. 2, 2022.
- [21] D. Grauer and J. Ploennigs, “Does chatgpt know building physics? exploiting foundation models for building performance prediction with gnns”.