A novel machine learning approach to predict short-term energy load for future low-temperature district heating

Thomas Ohlson Timoudas [a], *Yiyu* Ding [b], *Qian* Wang *[c, d]

[a] RISE Research Institutes of Sweden, Sweden, thomas.ohlson.timoudas@ri.se

[b] Department of Energy and Process Engineering, Norwegian University of Science and Technology (NTNU), Trondheim, Norway, yiyu.ding@ntnu.no

[c] Department of Civil and Architectural Engineering, KTH Royal Institute of Technology, Brinellvägen 23, Stockholm, Sweden, qianwang@kth.se

[d] Uponor AB, Hackstavägen 1, Västerås, 721 32, Sweden

Abstract.
In this work, we develop machine learning methods to forecast the day-ahead heating energy demand of district heating (DH) end-users in hourly resolution, using existing metering data for DH end-users and weather data. The focus of the study is a detailed analysis of the accuracy levels of short-term load prediction methods. In particular, accuracy levels are quantified for Artificial Neural Network (ANN) models with variations in the input parameters. The importance of historical data is investigated – in particular the importance of including historical hourly heating loads as input to the forecasting model. Additionally, the impact of different lengths of the historical input data is studied. Our methods are evaluated and validated using metering data from a live use-case in a Scandinavian environment, collected from 20 DH-supplied nursing homes through the years of 2016 to 2019. This study demonstrates that, although there is a strong linear relationship between outdoor temperature and heating load, it is still important to include historical heating loads as an input for prediction of future heating loads. Furthermore, the results show that it is important to include historical data from at least the preceding 24 hours, but suggest diminishing returns of including data much further back than that. The resulting models demonstrate the practical feasibility of such prediction models in a live use-case.

Keywords. Low-temperature district heating, short-term load prediction, machine learning, Scandinavian climate.
DOI: https://doi.org/10.34641/clima.2022.319

1. Introduction

District heating (DH) plays a vital role for the operation of building energy supply systems, which accounted for 35% of global final energy use and 38% of energy-related CO_2 emissions [1]. However, existing DH networks in many cold climates still use rather high supply temperatures, such as 75 °C or above [2]. In the face of green energy initiatives, increasing shares of low-energy buildings, and case examples in mild climates, there is a pressing need to transform the existing DH networks toward low-temperature DH (LTDH). Moreover, digitalization and the overall transition towards smart energy systems and cities are placing higher requirements on integration, communication, and cooperation with end-users (buildings) connected to such LTDH networks. As a result, future generations (4th and 5th) of LTDH networks will feature low operating temperatures, and greater integration with the end-users (buildings) and building-sized renewables. However, how to operate such integrations still rely fundamentally on a thorough understandings of heating loads.

On the other hand, digital solutions for measuring and controlling the network will allow for higher degrees of system optimization with intermittent renewables and heat pumps. This means that short-term predictions of heating loads are essential. But updating all the legacy monitoring facilities is a very costly and lengthy process. There is still a pressing need for more knowledge about what tools are available, and how well these methods can be utilized for load predictions in LTDH applications. At the same time, there is still room for improvement and solutions that can work on top of the existing DH systems, using existing metering data, during this transition period.

Large amounts of studies have developed different types of models to predict short-term energy demand in buildings in general. However, most of them are oriented towards electricity load predictions. For those reports that have been investigating DH load predictions, a great amount of methods are based on linear regression models, due to the strong linear relationships of heating load with respect to outdoor temperature. These existing methods commonly have not taken full advantage of using data-driven approaches, such as emerging machine learning (ML) models to perform such predictions. Even within those limited publications in the respective areas, it is still not clear what are the key advantages of using such ML approaches, and to what extent the accuracy levels can be quantified, given limited dataset inputs. This study provides a practice of the above raised challenges.

1.1 Previous studies

As mentioned, great amounts of studies exist with respect to building load predictions. These studies commonly utilize knowledge and experience gained from buildings' electricity demand data research, and transferable to heating demand analysis and uncertainties (i.e., weather forecasting) [3]. Among them, two types of models, namely, autoregressive multiple linear regression (MLR) and autoregressive multiple non-linear regression (MNLR), were firstly built to predict the DH load profiles, which can further aggregate them into district levels [4]. The Artificial Neural Network (ANN)-based MNLR shows better performance than MLR in the application of high load variations within 1-day time series [4]. Additionally, three supervised learning methods, namely, Support Vector Machine (SVM), deep neural network (DNN, ANN with two or more hidden layers), and extreme gradient boosting (XGBoost), were exercise to predict loads on a multi-step basis. These methods were also practiced with direct and recursive strategy [5]. By feeding day-before influential factors, all the above methods yield rather accurate forecasting for the day-ahead DH load at the given climate. However, how to use these heating load forecasting to optimize operation for DH system were not explored [5]. Some studies applied Gaussian mixture model (GMM) clustering enables defining four typical DH operation patterns in office buildings in a semi-arid climate (with cold and dry winters) [6]. Both outdoor temperature and occupant behavior data were considered. It is reported that by combining the clustering with regression and ANN models, the qualities of hourly heating load forecasting are improved by 38.7-75.7% in the given climate. However, it showed challenges to predict the peak heating loads during night-to-daytime periods, which can highly reply on random operation behaviors [6]. Forecasting model based on convolutional neural network long-short term memory (CNN-LSTM) outperformed other data-driven methods when solving thermal inertia problems in DH networks. This is reasoned by mainly owing to its integration of CNN's feature extraction ability and LSTM's two-dimensional space ability [7]. But this model requires large amounts of sensors, large data storage, and daily retraining, which place high pressure on the data management infrastructure [7]. Some other studies compared different data-driven methods, such as SVM and nonlinear autoregressive exogenous recurrent neural network (NARX-RNN) in the content of DH applications [8]. It was found that the NARX-RNN exceeds the SVM regarding the quality indicators and computation time. However, the overfitting tendency of NARX-RNN needs further study [8].

1.2 Objective

In this work, we develop ML methods to forecast the day-ahead heating energy demand of DH end-users in hourly resolution, by using existing metering data for DH end-users and weather data. The importance of historical data is investigated – in particular the importance of including historical hourly heating loads as input to the forecasting model. Additionally, the impact of different lengths of the historical input data is studied. The feasibility of such models, and their accuracy, are evaluated using data from a live use-case in Scandinavian environment. A detailed analysis of the accuracy levels of short-term load prediction methods are in focus.

2. Methodology

The study applies combinations of a two-step approach:

Step 1. A thorough understanding of the DH network and building load on annual basis, namely load profiles. This provides an overall view and boundary conditions of DH networks.

Step 2. Based on the definitions of DH load profile, day-ahead prediction models are developed. The model is rooted as an Artificial Neural Network (ANN) model, varying the input parameters, and trained and evaluated using the DH dataset.

To measure and evaluate the performance of the models, the mean squared error (MSE), and the mean absolute error (MAE), were both recorded for each model after training had been completed, using the 2019 test data (that had not been seen by the models during training).

2.1 Data inventory

The heating load was measured and collected for 20 separate nursing homes in Scandinavian climate, all located in the city of Trondheim, Norway. All of these buildings are connected to the same DH network, and the measurements were obtained directly from the measuring equipment of the network operator. The data contains the hourly heating loads for each of the buildings, spanning the entire time period from January 1, 2016 to December 31, 2019, obtained from the energy monitoring platform of Trondheim Municipality [9].

For the model construction and evaluation, the average heating load per square meter (W/m^2) was calculated across the 20 buildings for each hour. The data were supplemented with hourly outdoor temperature measurements obtained from the Norwegian meteorological station [10] in Trondheim, for the corresponding period.

2.2 Load profile development

The load profile was identified using an energy signature (ES) curve in the study. This method has been widely employed for planning and sizing purposes. An ES curve consists of a temperature dependent part, and a temperature independent part, which are divided by changing point temperature (CPT) or heating effective temperature, defined as:

If $T_t \leq$ CPT, $\quad P(T_t) = p_1 \cdot T_t + p_2 + \varepsilon \quad (1)$

If $T_t >$ CPT, $\quad P(T_t) = p_1 \cdot T_t + p_2 + \varepsilon \approx p_2 \quad (2)$

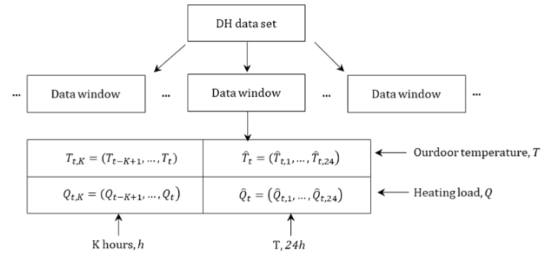

Fig. 1 The logic of short-term prediction model

where T_t is the outdoor temperature at time t, p_1 and p_2 are the coefficients of each ES curve model, and ε is the residual error. The heating demand follows the linear growth under the slope of p_1. Below the changing point temperature, it is the outdoor temperature dependent part and above the changing point temperature, it is the outdoor temperature independent part, when most of the heating needs go to domestic hot water (DHW) use.

For DH network monitoring, the load data are commonly aggregated as a combination of space heating and domestic hot water usage. Therefore, in the energy signature analysis, DHW load is extrapolated based on the existing studies [11], which has reported as a representative DHW profile for the given climate and resident types.

For modeling boundary conditions, daily heating degree hours (HDH) is calculated as the daily summation of the difference between balance temperature and hourly outdoor temperature, see below:

$$HDH = \sum_{t=t_0}^{t_0+23} \max(0, T_{bal} - T_t) \quad (3)$$

where t_0 is the first hour of the day, the heating balance temperature T_{bal} is assumed at 15°C and negative summands are set to zero. From this, high-heating season, mild-heating seasons and non-heating seasons can be identified in the ES curve.

2.3 The day-ahead prediction models

In this study, short-term prediction is defined as 24-hours (day-ahead) time horizon. ANN-based models were developed to predict the short-term heating load, starting from a given hour, for each hour of the following 24-hour period. As mentioned, this serves as a decision-supporting tool for the operation purposes in future LTDH transitions. All of these models used as input the forecasted outdoor temperature for the corresponding 24-hour period. To study the importance of historical data, and the performance impact of different measuring scenarios, nine differentiated ANN models were created and compared.

The models differed in what additional input data were used. One of them used no additional inputs, i.e., only the forecasted outdoor temperature. The other eight models were split into two main categories:

- Half of them were additionally supplied with the historical outdoor temperature,

- The other half were supplied, in addition to that, with the historical measured heating load.

For both cases, the historical data were given in the same hourly resolution. Within each category, the models were further differentiated based on the number of hours of historical data stretched back: 12, 24, 48, or 72 hours.

All of these models had one input layer (the number of inputs varied between the models), one hidden Rectified Linear Unit (ReLU) layer with 64 nodes (this number was determined through hyperparameter search), and one output layer. All the layers were densely connected. Mean squared error (MSE) was used as the loss function, and Adam was used for the parameter optimization, with the maximum number of epochs set to 100.

2.4 Mathematical description of the models

The logic of the developed model is presented in Figure 1. Let Q_t and T_t represent the measured heating load, and the measured outdoor temperature, at hour t, respectively; and let $\hat{Q}_{t,s}$ and $\hat{T}_{t,s}$ represent the predicted heating load, and the forecasted outdoor temperature, made at hour t for hour $t+s$ (defined for $s = 1, \dots, 24$), respectively. Let K be a parameter representing the number of hours of historical measured data to be used as input for the model. Introduce the shorthand notation as,

$$\hat{Q}_t = (\hat{Q}_{t,1}, \dots, \hat{Q}_{t,24}), \quad (4)$$

$$\hat{T}_t = (\hat{T}_{t,1}, \dots, \hat{T}_{t,24}), \quad (5)$$

$$Q_{t,K} = (Q_{t-K+1}, \dots, Q_t), \text{and} \quad (6)$$

$$T_{t,K} = (T_{t-K+1}, \dots, T_t). \quad (7)$$

That is, $\hat{Q}_t$, $Q_{t,K}$, $\hat{T}_t$, and $T_{t,K}$ represent, at the time instance t, the predicted heating load for the following 24 hours, the historical heating load for the preceding K hours (including t), the forecasted outdoor temperature for the following 24 hours, and the historical outdoor temperature for the preceding T hours (including t), respectively.

Each ANN model can then be expressed as either the function

$$\hat{Q}_t = f_K(\hat{T}_t, T_{t,K}), \quad (8)$$

if historical heating load is not an input to the model, or as

$$\hat{Q}_t = g_K(\hat{T}_t, T_{t,K}, Q_{t,K}), \quad (9)$$

if historical heating load is supplied, where f_K and g_K are abstract representations of our ANN models, and the parameter K takes either of the values 0, 12, 24, 48, and 72 (hours).

2.5 Training and evaluation of the models

As mentioned above, the different models were all trained and evaluated using the same dataset, introduced in Section 2.1. The dataset was created from the original data by first considering every possible consecutive 24-hour window of both the outdoor temperature and the heating loads, and then appending the preceding K-hour window to it, both for the outdoor temperature and the heating loads. In the cases that did not consider the historical heating load, that part of the window was simply discarded. Each window is therefore split into input and output, according to Fig. 1.

The data for the years 2016 and 2017 was used as the training set for the ANN, while the data for 2018 was used as the validation dataset, for the stopping criterion of the training. The resulting models were evaluated using the data for the entire year of 2019, the testing set, to ensure that the models were evaluated on a whole year of data.

Note that the model was evaluated using the actual measured outdoor temperature as the outdoor temperature forecast input. To improve statistical reliability, each model was trained from scratch ten times (using the same training data, but randomly initializing the weights each time), and the averages of these performance measures across the ten iterations were recorded.

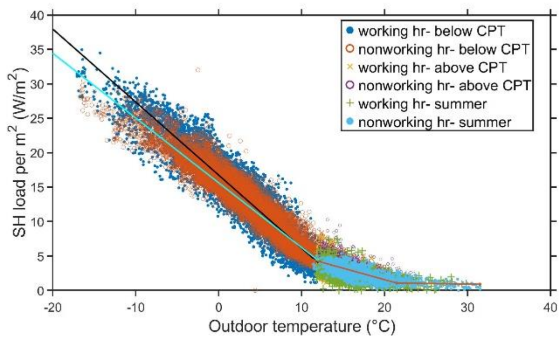

Fig. 2 – Energy signature (ES) curve for the district heating (DH) load profile

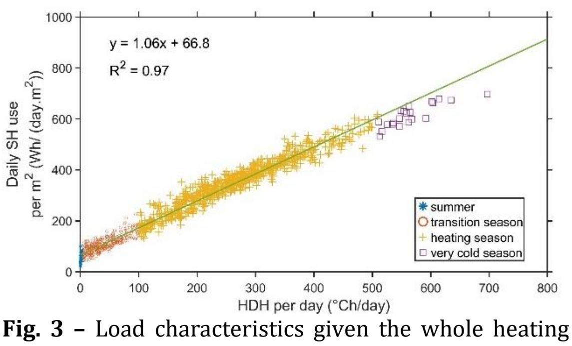

Fig. 3 – Load characteristics given the whole heating season, presented by heating degree hours (HDH)

3. Results

3.1 ES and load profile characteristics

Fig. 2 shows the ES of the DH network. Around 12°C was found as the changing point temperature for providing a proper piece-wise approximation. It is found that outdoor temperature that are above the changing point temperature consists of 22.4% of heating seasons. Fig. 2 also shows that space heating loads are less temperature dependent at the mild-heating season (constant slope), and these small loads can be described by one regression line regardless of working hours and non-working hours. The rest 77.6% of the time the outdoor temperature was below the changing point temperature, falls into high-heating season. Along the regression lines below the changing point temperature, there is a small region where non-working hour may need slightly higher space heating load than working hour under the same outdoor temperature (c.a. 10 - 12°C).

From the linear relationship between specific daily space heating and heating degree hours, as displayed in Fig. 3, it shows the daily space heating operation follows the daily heating degree hours, without influences from day types or manual false operation/intervention. These results are expected, given the rather high-temperature/conventional DH networks in the study. This also provides the boundary conditions that the day-ahead predictions will be constrained by the operation scenarios, instead of allowing the network temperature drift freely with load variations.

3.2 Accuracy levels of day-ahead prediction

Tab. 1 – Performance measures for the models, evaluated on the testing set of 2019

Model parameter	Mean squared error (MSE)	Mean absolute error (MAE)
No historical data, i.e., f_0		
K = 0	0.0824	0.2275
Only historical outdoor temperature, i.e., f_K		
K = 12	0.0790	0.2213
K = 24	0.0770	0.2183
K = 48	0.0753	0.2161
K = 72	0.0698	0.2086
Including historical heating load, i.e., g_K		
K = 12	0.0307	0.1299
K = 24	0.0219	0.1106
K = 48	0.0231	0.1133
K = 72	0.0221	0.1112

The evaluation error of the models are shown in Table 1. Recall that the evaluation of the models was performed on the dataset covering the entire year of 2019, and that this data had not been previously seen by the model (during the training stage). The results show a clear difference between the models g_K that use the historical load data, and the models f_K that do not. In particular, the impact of *only* including historical outdoor temperature data as input to the model is relatively small, even when longer periods of historical data are used, compared to also including the historical load data.

3.3 Example prediction results

The prediction performance of the two models f_{72}, g_{72}, and the models using 72 hours of historical data as input, are compared in both Fig. 4 and Fig. 5, showing results for different heating seasons. Recall that the model f_{72} does not consider the historical heating load as input, whereas g_{72} does. The top row shows the prediction results for f_{72}, and the bottom row the results for g_{72}. Each column shows the prediction for the 24-hour period following the time indicated above it, i.e., the prediction of the heating load $\hat{Q}_t = (\hat{Q}_{t,1}, \dots, \hat{Q}_{t,24})$ is plotted for the time instance t equal to the given date. Therefore, the performance of the two different models can be easily compared for the same time instant, by looking at each column.

The red crosses show the predicted heating load made by the model, while the blue boxes indicate the actual measured heating load. The dashed black line shows the forecasted outdoor temperature for the corresponding 24-hour period – in this case the actual outdoor temperature for evaluation purposes.

The three dates in Fig. 4 were randomly sampled from the 2019 testing data, to ensure statistical reliability. Since the first random sample did not cover the high-heating season (the period with the highest heating demand), we sampled and plotted the predictions for three additional dates from the months of December, January, and February, shown in Fig. 5, which falls in the high-heating season.

4. Discussion

A significant difference can be observed between the models those use both historical heating loads and outdoor temperature as inputs and the models those only use historical outdoor temperature. This difference is especially significant during the mild-heating season, when the heating load is dominated by domestic hot water, as evidenced by Fig. 4. This is likely due to the relatively weak relationship between outdoor temperature and the total heating load during that period, compared to the high-heating season, when space heating demand is the dominant component. Another reason could be due to thermal inertia and storage effects of the buildings, as well as suboptimal control of the heating loads, in which case the historical heating loads could be useful to model. This evidence provides a basis for how future LTDH should be transferred under different climate conditions, when heating loads fall more into the mild-heating season regime, with

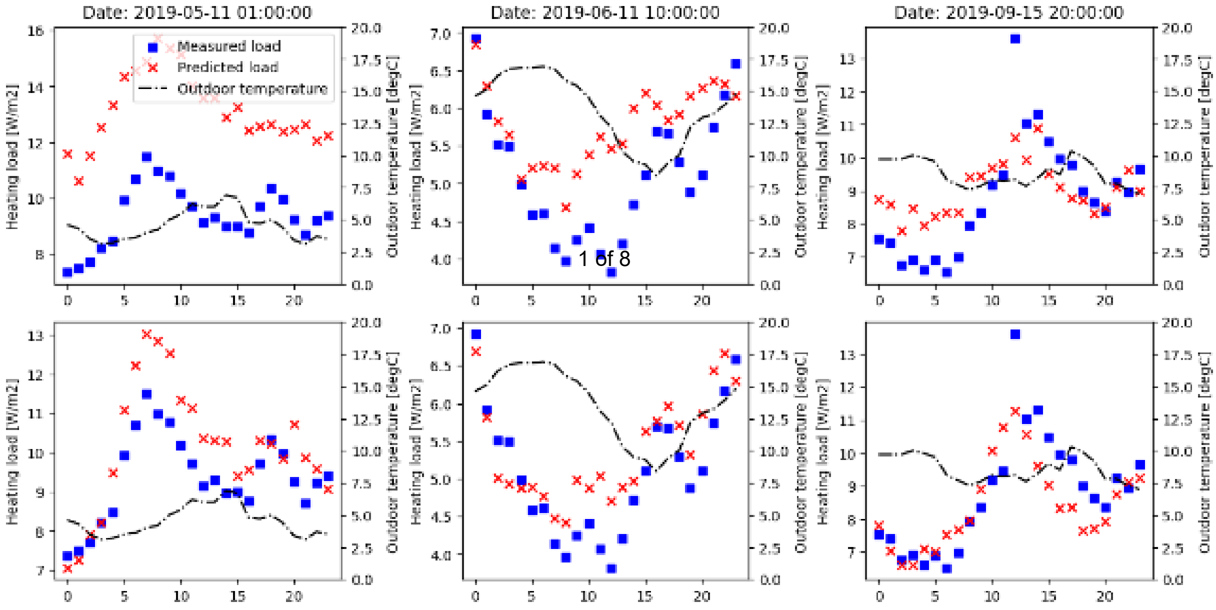

Fig. 4 – Plotted predicted heating load for the 24-hour period following the date indicated above each column, for the model f_{72} (top row), and the model g_{72} (bottom row). The dashed black line shows the outdoor temperature for the corresponding 24-hour period, the red crosses the predicted heating load for each hour, and the blue squares the measured (actual) heating load. Randomly sampled dates in mild-heating season.

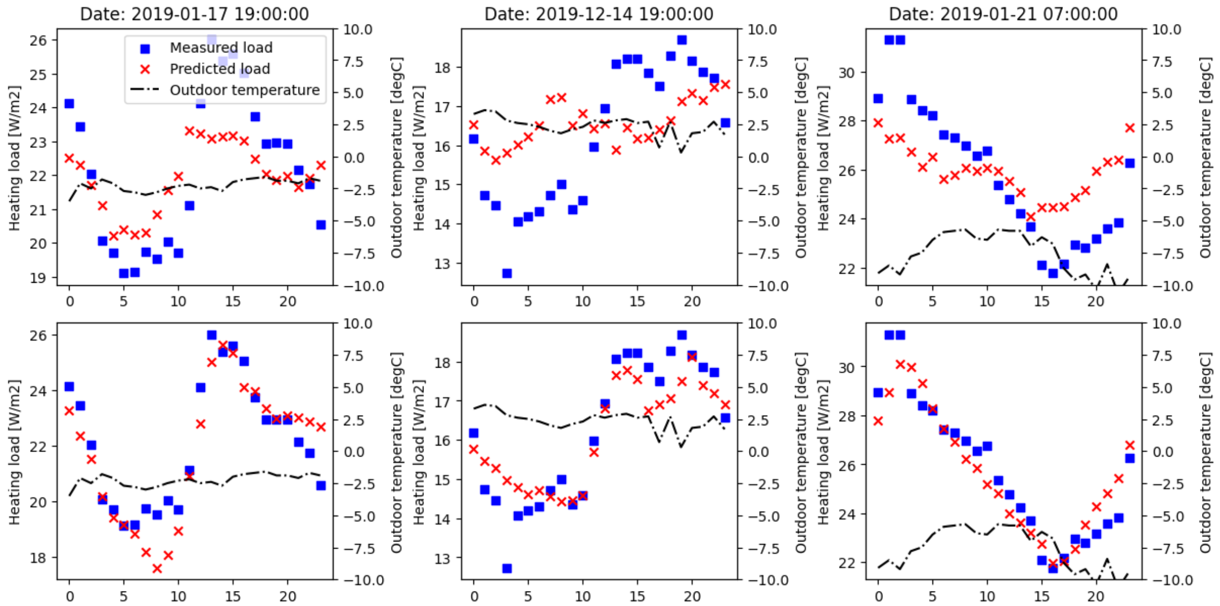

Fig. 5 – Plotted predicted heating load for the 24-hour period following the date indicated above each column, for the model f_{72} (top row), and the model g_{72} (bottom row). The dashed black line shows the outdoor temperature for the corresponding 24-hour period, the red crosses the predicted heating load for each hour, and the blue squares the measured (actual) heating load. Randomly sampled dates from December, January, and February, high-heating season.

perhaps only peaks fall into the high-heating season regime. These differences are also evident in the Table 1, which shows the average performance over the whole-year period. The results demonstrate the importance of making historical heating load available to heating load prediction models. Yet, while historical hourly outdoor temperature is often publicly available, historical heating loads are in many cases only available with large delays or low temporal resolution, if at all. The results additionally demonstrate the importance of using historical data from longer time periods, although they seem to suggest diminishing returns beyond the data for the previous 24 hours. This optimal cut-off period will likely differ between different building types, due to differences in thermal inertia.

It should be noted that the performance of the models was evaluated using the actual measured outdoor temperature as the forecasted outdoor temperature for the following 24-hour prediction. In practical applications, this forecast would typically be inaccurate. Such inaccuracies would lead to lower performance than observed in this study. As such, it is important that the base model is as accurate as possible, to reduce the propagation of such inaccuracies within the model. Moreover, further studies are needed to understand and quantify the effects of inaccurate forecasts, to ensure practical applications of these prediction models. Since thermal storage and intermittent renewables have own inaccuracies and uncertainties, it is especially important when they are integrated into the LTDH network perspective and optimized as a whole.

The mean average percentage error (MAPE), another common evaluation metric in the literature, was left out since the average heating load per square meter was close to 0 during certain times of the year (division by zero).

The limitation of this study is that it only considered a single type of model, an ANN, and that we did not perform a wide hyperparameter search (although we tried several different numbers of nodes and hidden layers, that all produced similar results). A further study should investigate if the results hold also for other types of prediction models and ANN architectures.

Moreover, we only considered two parameters: outdoor temperature (provided by a meteorological institute), and measured heating loads (provided by a municipality monitoring platform). It would be interesting to study the impact of additional variables, such as the integrated heat pumps and short-to-medium term thermal storage in such networks.

5. Conclusions

This study demonstrates that, although there is a strong linear relationship between outdoor temperature and heating load, it is still important to include historical heating loads as an input for prediction of future heating loads. Accuracy levels are quantified by using ANN models with input parameter variations. Furthermore, the results show that it is important to include historical data from at least the preceding 24 hours, but suggest diminishing returns of including data much further back than that. The models developed in this study were evaluated on actual measured data from a live use-case, demonstrating the practical feasibility of such prediction models.

6. Acknowledgement

This work is financially supported by Swedish Energy Agency with project No. 51544-1 and EU H2020 programme under Grant Agreement No. 101036656. Special thanks to Trondheim municipality, Norway, for providing data and user information.

7. References

[1] U. N. Environment, "2020 Global Status Report for Buildings and Construction." https://globalabc.org/news/launched-2020-global-status-report-buildings-and-construction (accessed Nov. 15, 2021).

[2] K. Lygnerud, S. Werner, and etc., "Low-Temperature District Heating Implementation Guidebook. IEA DHC Report; 2021," p. 206.

[3] M. Lumbreras *et al.*, "Data driven model for heat load prediction in buildings connected to District Heating by using smart heat meters," *Energy*, vol. 239, p. 122318, Jan. 2022, doi: 10.1016/j.energy.2021.122318.

[4] B. Talebi, F. Haghighat, and P. A. Mirzaei, "Simplified model to predict the thermal demand profile of districts," *Energy and Buildings*, vol. 145, pp. 213–225, Jun. 2017, doi: 10.1016/j.enbuild.2017.03.062.

[5] P. Xue, Y. Jiang, Z. Zhou, X. Chen, X. Fang, and J. Liu, "Multi-step ahead forecasting of heat load in district heating systems using machine learning algorithms," *Energy*, vol. 188, p. 116085, Dec. 2019, doi: 10.1016/j.energy.2019.116085.

[6] Y. Lu, Z. Tian, P. Peng, J. Niu, W. Li, and H. Zhang, "GMM clustering for heating load patterns in-depth identification and prediction model accuracy improvement of district heating system," *Energy and Buildings*, vol. 190, pp. 49–60, May 2019, doi: 10.1016/j.enbuild.2019.02.014.

[7] J. Song, L. Zhang, G. Xue, Y. Ma, S. Gao, and Q. Jiang, "Predicting hourly heating load in a district heating system based on a hybrid CNN-LSTM model," *Energy and Buildings*, vol. 243, p. 110998, Jul. 2021, doi: 10.1016/j.enbuild.2021.110998.

[8] D. Koschwitz, J. Frisch, and C. van Treeck, "Data-driven heating and cooling load predictions for non-residential buildings based on support vector machine regression and NARX Recurrent Neural Network: A comparative study on district scale," *Energy*, vol. 165, pp. 134–142, Dec. 2018, doi: 10.1016/j.energy.2018.09.068.

[9] "iEOS - Planning." https://www2.esave.no/Esave.nsf/iEOS_Hovedbilde.xsp (accessed Feb. 24, 2021).

[10] "Norsk Klimaservicesenter." https://seklima.met.no/observations/ (accessed Feb. 24, 2021).

[11] D. Ivanko, H. T. Walnum, and N. Nord, "Development and analysis of hourly DHW heat use profiles in nursing homes in Norway," *Energy and Buildings*, vol. 222, p. 110070, Sep. 2020, doi:

10.1016/j.enbuild.2020.11007

8. Data access statement

The datasets generated during and/or analysed during the current study are not publicly available because the Municipality holds the right of the datasets but the authors will make every reasonable effort to publish them in near future.