Toward Explainable Heat Load Patterns Prediction for District Heating

Authors

L. Minh Dang[1], Jihye Shin[2], Yanfen Li[3], Lilia Tightiz[4], Tan N. Nguyen[5], Hyoung-Kyu Song[1], and Hyeonjoon Moon[3,*]

Affiliations

1. Department of Information and Communication Engineering and Convergence Engineering for Intelligent Drone, Sejong University, Seoul, South Korea
2. Department of Artificial Intelligence, Sejong University, Seoul, South Korea
3. Department of Computer Science and Engineering, Sejong University, Seoul, South Korea
4. School of Computing, Gachon University, 1342 Seongnam-daero, Sujeong-gu, Seongnam-si, 13120, Gyeonggi-do, Republic of Korea
5. Department of Architectural Engineering, Sejong University, 209 Neungdong-ro, Gwangjin-gu, 05006, Seoul, Republic of Korea

corresponding author(s): Hyeonjoon Moon (hmoon@sejong.ac.kr)

Abstract

Heat networks play a vital role in the energy sector by offering thermal energy to residents in certain countries. Effective management and optimization of heat networks require a deep understanding of users' heat usage patterns. Irregular patterns, such as peak usage periods, can exceed the design capacities of the system. However, previous work has mostly neglected the analysis of heat usage profiles or performed on a small scale. To close the gap, this study proposes a data-driven approach to analyze and predict heat load in a district heating network. The study uses data from over eight heating seasons of a cogeneration DH plant in Cheongju, Korea, to build analysis and forecast models using supervised machine learning (ML) algorithms, including support vector regression (SVR), boosting algorithms, and multilayer perceptron (MLP). The models take weather data, holiday information, and historical hourly heat load as input variables. The performance of these algorithms is compared using different training sample sizes of the dataset. The results show that boosting algorithms, particularly XGBoost, are more suitable ML algorithms with lower prediction errors than SVR and MLP. Finally, different explainable artificial intelligence approaches are applied to provide an in-depth interpretation of the trained model and the importance of input variables.

Introduction

District heating (DH) has risen as a crucial energy supply infrastructure in order to effectively provide heat and cooling to consumers over the last few decades[1]. DH is superior in many aspects compared to other energy supply options, which include having a lower carbon footprint, the integration of multiple heat sources, and high energy throughput. The latest fourth and fifth generations of DH can utilize several heat sources, which include combined heat and power (CHP), gas boilers, water-source heat pumps (HPs), ground-source HPs, and solar energy-based HPs. The recent literature focused more on developing simulation frameworks and effective approaches in regards to designing and optimizing DH systems in terms of the economic and energetic factors, which is due to the fast development of DH technologies[2,3]. Storage technology is also a hot topic, because it helps decouple heat production and the demand to increase DH efficiency[4]. The following articles[1,5] were reviewed in order to obtain the latest information about DH networks.

The heat usage pattern analysis has become increasingly essential as the number of end-users increases, because it greatly impacts the entire network's efficiency. Variations in the heat usage behavior from the consumers' side lead to variations in the heat usage pattern of a single substation, which is a major matter for accurate and efficient DH management and operation[6]. For example, the substantial temperature difference between the summer and the winter significantly influences the users' heat demand. In addition, the hourly heat demand also varies between households, which causes heat demand variation at the substation[7].

An accurate heat demand prediction framework is imperative in order to effectively manage DH networks[8]. First, it facilitates the optimization of the overall heat production, minimizes the heat loss, and optimizes the operating costs. Second, the distribution temperature is provided at an appropriate range in order to predict the real-time heat usage using the heat demand forecast model. As a result, the number of studies proposed in regards to predicting the heat demand has been increasing. A heat demand analysis can generally be divided into model-based and data correlation categories[9]. The data correlation approach mainly depends on building functional correlations of the DH parameters in order to develop a heat usage profile for each substation or building. The model-based technique relies on machine learning (ML) algorithms in order to effectively learn the representative patterns using the historical heat load data[10]. The data correlation approach offers higher accuracy than the model-based approach, but it is time-consuming and laborious due to each building/substation having a unique heat usage profile that needs to be constructed. The performance of the model-based heat usage prediction algorithm has become significantly better, which is due to the huge advancements in artificial intelligence (AI) and big data over the past few decades[9,10].

The heat usage prediction, heat loss estimation, and abnormality analysis based on the energy signature (ES) have been increasingly investigated in recent years, which have shown promising results[11,12]. However, these studies mainly used outdoor temperature as the main feature in order to discover the heat demand pattern. Other studies focused on peak usage forecasting with the ultimate objective of optimizing the energy usage and DH management[13]. These studies, which are similar to the ES, failed to consider the meteorological data or the end-user behaviors. Potential influencers of the heat demand patterns can be divided into three main factors, which include meteorology, behaviors, and time[14]. Some common meteorological data that potentially affects heat demand are humidity, solar irradiation, outdoor temperatures, and the wind flow speed[15]. Time factor involves all time-related parameters, which include hours, days, months, and years. The social behaviors of the end-users are also a crucial influencer of the heat load variation, which can be affected by both meteorological and time factors[16]. These three main factors significantly influence the heat demand patterns.

There has been considerable interest in the research area of heat load forecasting for DH, as indicated by numerous recent studies. Idowu et al.[29] examined a range of supervised ML algorithms in order to perform heat load prediction up to 48 hours in advance. The experimental results revealed that conventional ML algorithms, such as SVM and linear regression, achieved the lowest normalized root mean square error when compared to other algorithms. In another study, Boudreau et al. found that ensemble models provided significantly better prediction accuracy than base ML models when it came to predicting peak power demand and next-day building energy usage[30].

Several studies have delved into specific aspects of DH systems. For example, Saloux et al. explored the application of ML algorithms for predicting the aggregated heating usage of a community. They concluded that the models' performance could be significantly enhanced by

considering other crucial factors, such as time of day, systematic variables, and temperature[31]. López et al. focused on the impact of specific days, such as holidays or festive periods, on the load curve, and determined that such events could considerably affect the heat usage pattern[32]. Moreover, a case study of a large DH network over several heating seasons revealed that the primary force of heat demand were the various operation settings during daytime (night shutdown and night temperature setback) and the outdoor temperature[33].

Despite the numerous issues addressed and methods discussed in existing literature on heat load prediction in DH networks, further research is needed to explore important external factors such as holiday and weather conditions, which could be utilized as input to improve the models' accuracy[6]. Additionally, while previous work has showed the high predictive performance of ML algorithms for heat demand, they have not provided a clear explanation of why the model achieved good performance, as well as which features are important and their correlation with the models[10].

This research is proposed in order to improve the heat usage prediction via an in-depth analysis of the dataset to figure out the potential factors that impact the heat demand. The main contributions include (a) performing a data analysis prior to the training process to help thoroughly understand the dataset, (b) training and comparing different ML models in order to obtain the best hourly heat load prediction model, and (c) offering detailed explanations about what features were imperative to the model prediction, which were overlooked in the previous studies.

The remainder of the manuscript is outlined as follows. Section “Dataset description” gives a detailed description of the proposed heat demand dataset. After that, the Section “Methodology” outlines all processes involved in heat demand prediction. Several experiments are performed in Section “Experimental results” to comprehensively assess the proposed framework. Next, the Section “Discussion” discusses the findings and provides a detailed analysis of the study. Finally, we conclude the study and offer future work in the Section “Conclusion”.

Dataset description

The dataset that is described in this research was the hourly heat demand from an eco-friendly liquefied natural gas (LNG)-based cogeneration plant in the Cheongju region, Korea. The plant produces around 76.5 Gigacalories (Gcal) of local heating to the distribution grid. Gcal is a common heat load unit, which measures the heat energy in the heating plants. The LNG-powered plant is more efficient and environmentally friendly for the generation of thermal energy, which has been reported to produce over 70% less emission than coal or oil sources.

The dataset introduced in this study includes the hourly heat usage from January 2012 to December 2020 of the residents from a region, which spans eight heating seasons from November to April. The heat usage profile suggests the amount of heat that is transmitted from the plant to the consumers at a specific duration, which mainly involves space heating (SH) and domestic hot water (DHW). The corresponding hourly historical weather data was also collected as an additional feature in order to discover the potential connections with the heat load patterns in addition to the heat load data. A holiday feature that indicates whether the day under consideration is a holiday is also added in order to investigate the end-user behaviors. The three main features that belong to the weather data include wind flow speed, humidity, and outdoor temperature. The collected heat usage dataset is used to study the hourly heat load patterns and provides some explanations for the model's predictions. The minimum, maximum, mean and standard deviation for each variable are described in Table 1.

Table 1: Description of important observations with possible values for the variables in the proposed dataset

Name	Minimum|Maximum	Mean| standard deviation	Unit
Date	01/01/2012 | 01/01/2022	-	-
Wind speed	0|8.7	1.47|0.93	m/s
Humidity	7|100	61.32|20.02	%
Outdoor temperature	-16.5|38.1	13.75|10.83	°C
Holiday	0 (normal day) | 1 (holiday)	0.32|0.46	-
Heat load	0|317	65.89|52.92	Gcal

In summary, 8760 hourly heat load profiles and their corresponding historical temperature data are obtained yearly. Therefore, a total of 87,672 entries, which include date and time, holiday, wind flow speed, humidity, and temperature, are used as the input variables, and the heat load profiles are used as the target variables. The data entries from 2012 to 2020 were used as the training set, whereas the hourly heat usage of 2021 was applied in order to test the model's performance.

Methodology

Figure 1 depicts the three components of the hourly heat usage prediction system, which are (a) data preprocessing, (b) pattern analysis and data partitioning, and (c) explainable heat load forecasting.

- Data preprocessing: There is a high possibility that the structured data may contain some common issues with data preprocessing, such as duplicate data, missing data, and negative data due to human errors, which can affect the system's performance. As a result, it is a prerequisite before the data analysis and training processes to fix all errors and standardize the data.
- Pattern analysis and data partitioning: Heat usage patterns play an important role in regards to enabling specialists to study consumer behavior. The distinctive patterns of the dataset are discovered in this section by using various data analysis approaches in order to thoroughly analyze the dataset before the training phase. The dataset is then divided into training and testing sets.
- Explainable heat load prediction: Different ML algorithms were trained in order to forecast the hourly heat usage. Some explainable artificial intelligence (XAI) approaches are finally implemented in order to interpret the model's predictions.

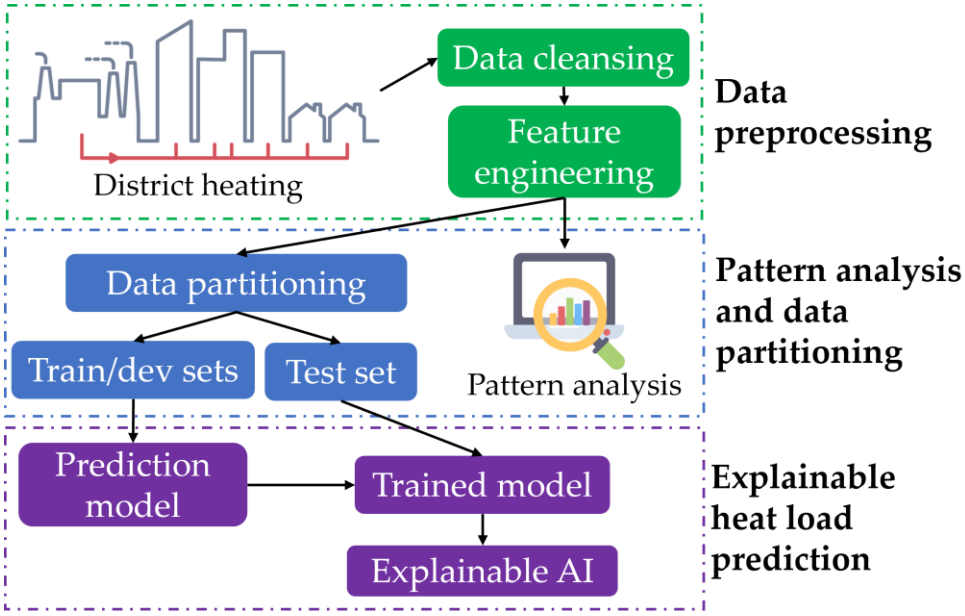

Figure 1: Description of the primary components of the heat usage patterns analysis framework

Data Preprocessing

Data cleaning

The structured data-related issues, such as missing and duplicated data are unavoidable during the data collection, and they can negatively affect the model's performance if not appropriately corrected. Data cleansing is therefore conducted in order to detect and fix error records in regards to the humidity, wind speed, outdoor temperature, and hourly heat usage data. There are various data cleaning processes, and the two main processes that were performed in this study include removing duplications and fixing the missing values. The dataset is loaded as a data frame using pandas, a famous data manipulation and analysis library. After that, data inconsistencies can be automatically detected using pandas-supported functions.

Standard techniques, such as moving average (MA) and imputation, are usually employed in order to correct the missing data. This study applied the exponential weighted moving average (EWMA) technique[17], which is an extension of the MA algorithm. EWMA considers the recent data points to be significantly important with a higher weight, whereas the data points in the further past receive an exponentially lower weight. Moreover, the EWMA method can be effectively applied due to the nature of the dataset, and the differences between the two consecutive data points are considered minor. The EWMA can be described as follows.

$$E_t = \alpha \times x_t + (1 - \alpha) \times E_{t-1} \quad (1)$$

where E_t indicate the computed value at time t based on the EWMA technique. x_t is the value of the series in the current period. E_{t-1} is the EWMA at the previous time period. Finally, α is the smoothing factor, which ranges between 0 and 1 and controls the influence of the current value x_t on the E_t. A larger α places more weight on recent observations and results in a more reactive EWMA, while a smaller α results in a smoother EWMA.

Feature engineering

Feature engineering is the process of selecting, extracting, and transforming relevant features or variables from raw data to enhance the performance of ML algorithms. The goal of feature engineering is to provide ML algorithms with informative and discriminative features that can help them better understand the underlying patterns and relationships in the data. Two main processes in the feature engineering process are standardization and feature transformation.

The regression model fitting and learned function can be negatively affected by structured data, and it eventually creates a bias when numerical features with different scales are fed into the model[18]. The normalization/standardization techniques therefore need to be implemented in order to normalize the input features. Min-max normalization and standardization are two common feature scaling approaches[19]. The heat usage dataset that is applied to fit the model contains peak heat load on some specific periods, which are outliers, and it has an essential role during the training process. The min-max normalization likely lowers the impact of those outliners by transforming all features into a range between 0 and 1. The standardization therefore scales the features in order to have a zero mean, and a standard deviation of 1 is implemented in this study.

Feature transformation is necessary for structured data in order to convert categorical inputs into numerical inputs, because most ML models work with numerical data. The holiday variable is categorical, because it has two distinctive values, which represent whether a particular day is a regular day or a holiday. As a result, one-hot encoding, which creates a binary representation of the categorical feature, is applied in order to transform the holiday feature[20]. For instance, when a specific day is a holiday, the value for the holiday binary variable is set to 1, and the regular binary variable is 0.

Pattern Analysis and Data Partitioning

Pattern analysis

Heat network during the summer season: The investigation of the heat network in the summer season, which spans from June to August, gives some exciting insights into the town's heat usage. Figure 2 illustrates the hourly heat demand distribution density for the summer months from 2012 to 2021. The average heat demand in the summer mainly involves the DHW consumption and the network heat losses. It can generally be seen that there was less heat demand in the distant past compared to the recent years. For instance, a roughly similar distribution can be observed for the following years, which include from 2012 to 2016, with the average heat demand being around 20 Gcal. However, the average heat demand increased to around 30 Gcal, which included the more recent years from 2019 to 2021, with some higher heat demands being related to particular heat usage patterns. Moreover, there has been a gradually increasing trend in the average heat usage of over 40 Gcal in recent years, and the year 2021 shows the highest density.

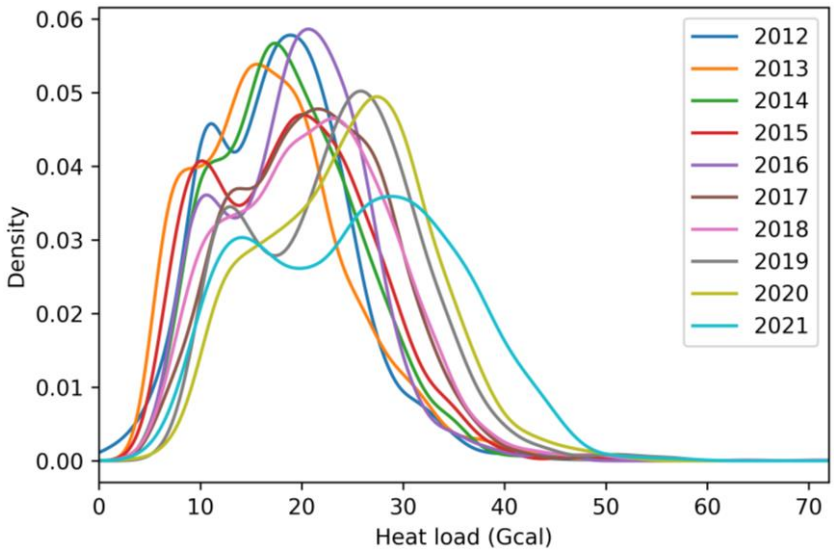

Figure 2: Distribution density plot of hourly heat demand during the summer season (Jun.-Aug.)

Heat network during the winter season: The chart in Figure 3 illustrates the network's energy consumption on an hourly basis during the winter season spanning from November to March. The chart depicts three distinct patterns for three different time periods: daytime (06:00 to 18:00), nighttime (22:00 to 05:00), and peak hours (19:00 to 21:00). The scatter plot reveals that the consumers tend to use more heat during the peak time at the same temperature level compared to the nighttime and daytime. Moreover, the lower the outside temperature, the higher the heat load that is required.

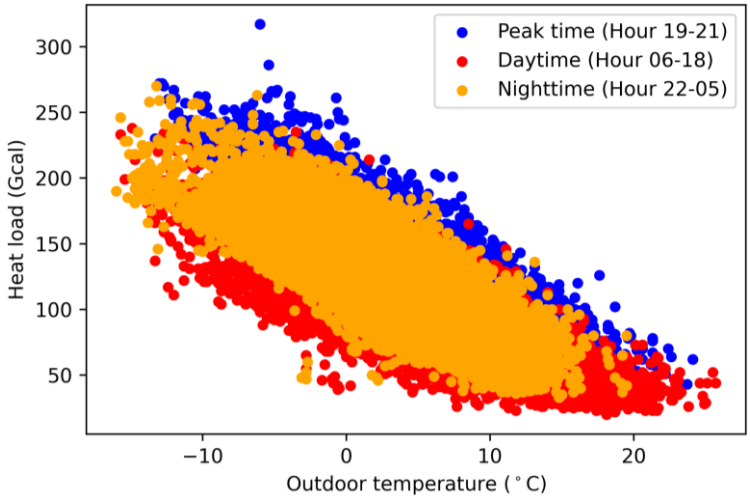

Figure 3: Scatter plot of the outdoor temperature and the heat usage during the winter season (Nov.-Mar.)

Some heat load patterns for each season of the year: A typical hourly heat load pattern for each season can be observed in Figure 4. The spring, fall, and winter seasons have similar variations in the hourly time scale, which is caused by the social behavior of the end-users. Reduced heat loads can be observed in the daytime, which is due to solar radiation that leads to higher daytime temperatures. The highest heat load during the daytime occurs around 8 am in order to prepare the space heating in offices and commercial buildings. The heat demand usually peaks between 19:00 and 21:00 because of the low temperature at night, which requires more heat for SH and DHW. DHW is a major part of the heat demand in the summer, when a tiny difference in the heat variation can be observed.

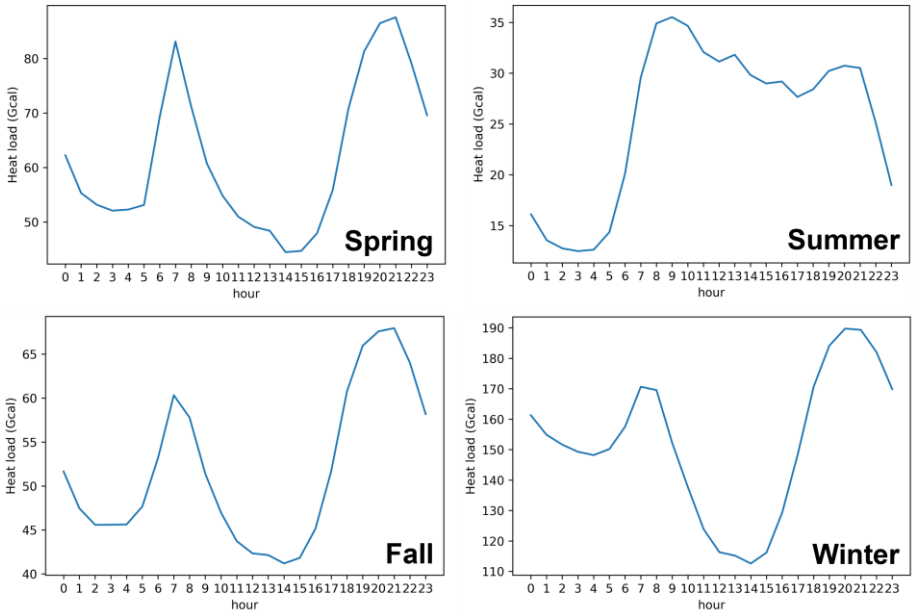

Figure 4: Average weekly heat load patterns during the four season periods

Data partitioning

Data partitioning is a fundamental step required before training and evaluating the model. After preprocessing, the data is split into two sets: the training set and the testing set. The training set is utilized to train and optimize the model, while the testing set is typically employed to assess the algorithms' performance across various scenarios. This study used the heat usage profiles between 2012 and 2020 as the training set, whereas the heat load profiles from 2021 were used for the testing. Each training or testing sample consists of day, hour, outdoor temperature, humidity, windspeed, and holiday as the input variables, while the output is the hourly heat usage corresponding to that particular input.

Explainable Heat Load Prediction

This section presents the main concepts behind boosting, support vector regression (SVR)[21], and multilayer perceptron (MLP) algorithms[22] that were implemented for the heat demand forecasting.

Boosting algorithms: Boosting algorithm belongs to the ensemble approach, which sequentially adds multiple weak learners. Each weak learner is added by using the learned information from its predecessor, and it tries to correct the errors that are predicted by them. A weak learner can be any learning algorithm that offers a slightly better performance than random guessing. Two standard boosting approaches are gradient boosting and adaptive boosting[23].

- Adaptive boosting: The adaptive boosting (AdaBoost) algorithm was proposed by sequentially adding weak learners, which involved using decision trees, and attempting in order to correct the wrongly predicted samples by applying a bigger weight to them during the training process of the latter weak learners. The AdaBoost model's final output is the weighted median.
- Gradient boosting: AdaBoost assigns new instance weights whenever a new weak learner is added, but gradient boosting aims to fit the new predictor to the residual errors that are caused by the prior predictor with the primary objective of minimizing a loss function[24]. Some popular gradient boosting algorithms include LightGBM and XGBoost.

XGBoost leverages the feature distribution across all data points to narrow down the search space of potential feature splits. The objective of the XGBoost algorithm can be expressed as:

$$objective = L + \mu \quad (2)$$

where the predictive ability of XGBoost is determined by the loss function L, while the regularization term μ is used to manage overfitting. μ is determined by the number of observers and their prediction threshold in the ensemble model. Since the problem in question belongs to regression analysis, the root mean squared error (RMSE) is used as the loss function L.

Support vector regression (SVR): Unlike typical regression algorithms that seek to minimize the sum of squared errors between actual and predicted values, SVR attempts to identify the optimal hyperplane within a user-defined threshold value. The threshold value is the distance between the boundary line and the hyperplane. Heat demand prediction is a complex non-linear topic, because it has multiple input variables. To address non-linearity in the initial feature space and treat it as a linear problem in the high-dimensional feature space, SVR requires the use of a non-linear kernel. The Gaussian Radial Basis kernel (RBF) was used in this study as the default kernel for SVR.

Multilayer perceptron (MLP): Multilayer perceptron (MLP) belongs to the feedforward artificial neural networks (ANN) category. MLP's fundamental structure consists of an input layer, one or more hidden layers with neurons, and an output layer that are stacked in sequence. The neuron is the primary computing component of MLP, and neurons from the current layers fully connect to neurons from the next layer. The inputs are added to the initial weights, fed into an activation function, and propagated to the next layer.

Experimental results

This section shows all experiments that were conducted to determine the most suitable algorithm for predicting heat usage. In addition, various XAI techniques were also conducted in order to provide an in-depth analysis of the trained models.

The heat load prediction models were constructed and trained on scikit-learn[28], a Python-based open-source ML library. Three main explainable AI libraries for analyzing the data include partial dependence plot[27] (PDP), which is a global and model-agnostic XAI algorithm, local interpretable model-agnostic explanations[39] (LIME), which create a local model approximation of the model around the prediction of interest, and shapley additive explanations[26] (SHAP), which employ a game-theoretic approach.

Evaluation Metrics

Three standard evaluation metrics were computed, which included the coefficient of determination (R^2), mean squared error (MSE), and mean absolute error (MAE) in order to evaluate the heat demand forecasting. MSE is computed by averaging the squared difference between the predicted values and actual values for all the training samples[25]. On the other hand, MAE is the average of the absolute differences between the predicted values and true values. While MSE measures the standard deviation of residuals, MAE calculates the average of the residuals in the dataset. R^2 is computed by determining the proportion of the dependent variable's variance predicted by the algorithm. The lower the MSE and MAE scores, the better the model's performance. However, a higher value of R^2 is considered better. The three metrics can be formulated as follows.

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \quad (3)$$

$$R^2 = 1 - \frac{\sum(y_i - \hat{y}_i)^2}{\sum(y_i - \bar{y})^2} \quad (4)$$

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i| \quad (5)$$

where N is the total number of training samples. y_i indicates the actual value, $\hat{y}_i$ means the predicted value of the ith profile, and $\bar{y}$ is the mean value of y.

Hyperparameter Fine-tuning

Five regression models were implemented in this study in order to perform the heat demand forecasting, which included SVR, AdaBoost, XGBoost, LightGBM, and MLP. Each model has its crucial hyperparameters that must be determined before the training. The hyperparameters control the training behavior of the learning algorithms, and they considerably influence the model's performance.

Table 2 shows the hyperparameters and the value range for each hyperparameter that is required by the five models. A grid search method was conducted next on the different combinations of the hyperparameters of each algorithm in order to explore the most suitable hyperparameter combination that helps the algorithm obtain the best performance.

Table 2: Initial hyperparameter value ranges and the optimal hyperparameter value for each algorithm

Model	Hyper parameter	Definition	Value ranges	Optimal value
AdaBoost	n	Number of estimators	50, 100, 150, 200	50
	σ	Learning rate	$10^{-3}, 10^{-2}, 10^{-1}$	10^{-1}
XGBoost	n	Number of estimators	50, 100, 150, 200	50
	d_{tree}	Max depth of a tree	3, 6, 9, 12, 15	9
	γ	Min loss reduction	0, 0.1, 0.2, 0.3	0
	subsample	Subsample ratio of the training instances	0.5, 1, 2	1
LightGBM	num_leaves	Max number of nodes per tree	21, 31, 41, 51	31
	σ	Learning rate	$10^{-3}, 10^{-2}, 10^{-1}$	10^{-1}
	n	Number of estimators	50, 100, 150, 200	100
	d_{tree}	Max depth of a tree	2, 3, 4, 5, 6	4
SVR	C	Regularization parameter	$10^0, 10^1, 10^2, 10^3$	10^0
	γ	Kernel coefficient	$10^{-6}, 10^{-3}, 10^{-1}$	10^{-3}
MLP	σ	Learning rate	$10^{-3}, 10^{-2}, 10^{-1}$	10^{-2}
	nh_i	Number of neurons in hidden layer ith	50, 100, 150, 200	150
	φ	Activation function	ReLU, tanh	ReLU
	B	Batch size	8, 16, 32, 64	32

Heat Usage Prediction Analysis

Figure 5 depicts the performance and scalability comparison of five different learning algorithms using the learning curves in order to show the effect of adding more samples during

the training process. The experiment involved randomly selecting samples from the training dataset. A training sample include date, outdoor temperature, windspeed, humidity, holiday, and hourly heat demand as the features.

It can generally be concluded that SVR and MLP were highly sensitive to the dataset size, because they widely fluctuated as more training samples were added. On the other hand, the boosting algorithms, which included AdaBoost, LightGBM, and XGBoost, showed their advantages and effectiveness with a bigger dataset. The three ensemble algorithms exhibited similar trends in variation; the error gradually decreased and eventually stabilized. Low MSE scores of less than 0.02 were obtained for the three boosting algorithms when the training dataset size was over 2000 samples. XGBoost achieved the lowest mean squared error of less than 0.01 among the three algorithms, and it showed its robustness when the number of training samples reached 7000. As a result, XGBoost was utilized as the primary model for the following experiments.

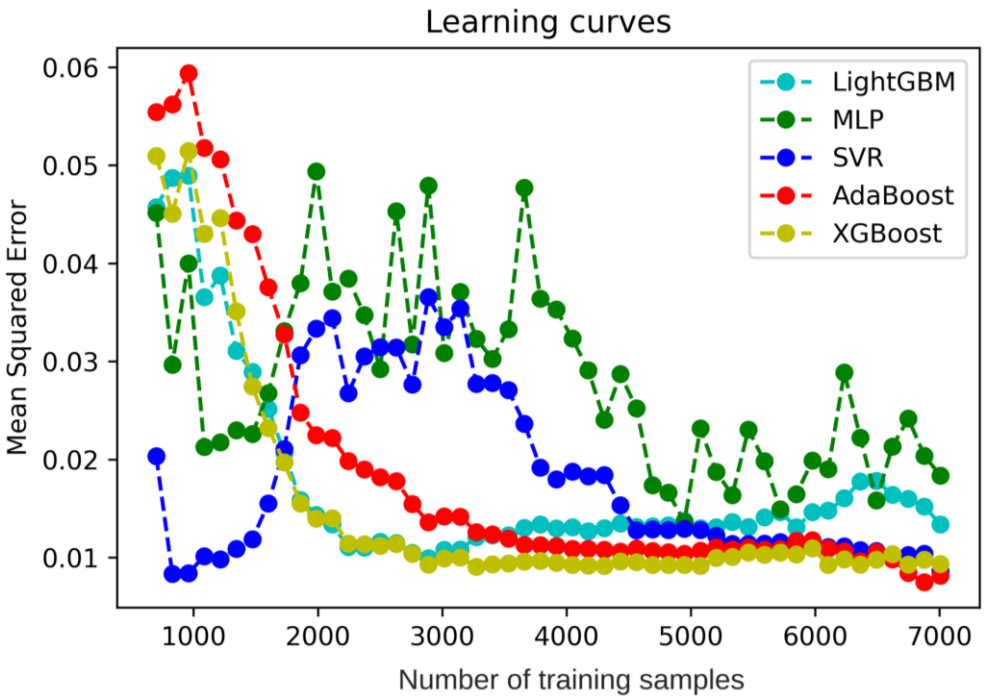

Figure 5: Heat demand forecasting performance using five different algorithms

Table 3 shows the heat demand forecasting performance using five ML algorithms on the test dataset. All the models generally obtained good performances on the dataset. The boosting algorithms performed better than SVR and MLP. The XGBoost algorithm achieved the highest R^2, MSE, and MAE at 0.95, 0.12, and 0.15, respectively. On the other hand, MLP showed the lowest heat usage prediction performance with an MSE value of 0.25 and R2 at 0.89.

Table 3: Hourly heat load prediction performance for the five ML algorithms on the testing dataset.

Model	MAE	MSE	R^2
AdaBoost	0.16	0.14	0.94
XGBoost	0.15	0.12	0.95
LightGBM	0.18	0.17	0.91
SVR	0.24	0.21	0.92
MLP	0.23	0.25	0.89

Figure 6 compares the actual and the predicted heat demand for 2021 using the XGBoost model. The heat usage values predicted by the model, which are illustrated by the red line, are roughly similar to the actual heat usage values, which are illustrated by the blue line. Moreover, each month's peak and bottom heat usage were accurately predicted. However, the model performance was significantly affected, which is due to some uncommon end-user's heat usage behaviors.

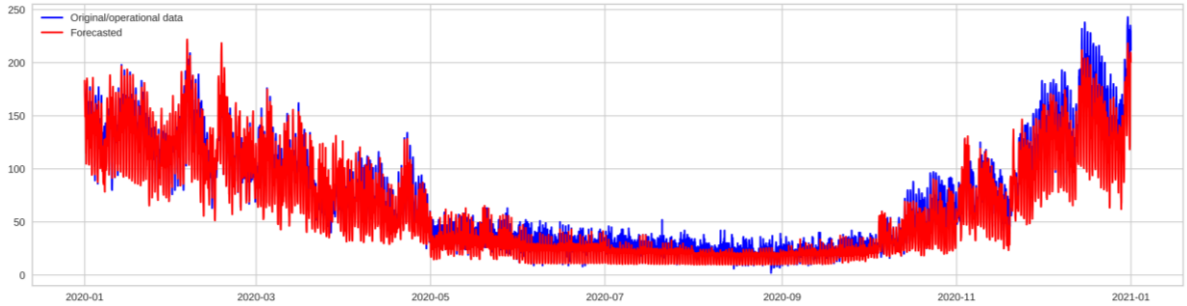

Figure 6: Daily heat load prediction results on the testing dataset

Explainable Heat Usage Prediction

The previous section discussed what model achieved the highest heat usage forecasting performance. However, it is challenging to reveal what features are influential and how they affect the model predictions. As a result, some interesting XAI approaches are implemented in this section in order to attempt to explain how ML models predict the outcomes.

Firstly, three different feature ranking techniques were implemented in order to evaluate each feature's importance in regards to predicting the output heat usage by the model, as displayed in Figure 7. Figure 7(a) calculates a feature's relative importance by examining the mean and standard deviation of impurity reduction across each tree. Figure 7(b) ranks the feature importance by computing the game's theoretically optimal shapley values[26]. The resulting shapley values provide a measure of the relative importance of each feature in the model prediction for a particular data point. It requires examining every possible feature combination and assessing the marginal impact of each feature on the prediction. Features with higher Shapley values are regarded as more significant. Ranking both approaches reveal that the temperature and month features are crucial, which is valid due to the end-users heat demand pattern being significantly affected by these two features.

Finally, Figure 7(c) visualizes the feature importance assessed by LIME. Positive weights indicate that a feature promotes a positive prediction, while negative weights indicate the opposite. The magnitude of the weight represents the importance of the feature. It is noticeable that a temperature of 4°C or lower (cold season) presses the model to output a higher heat usage.

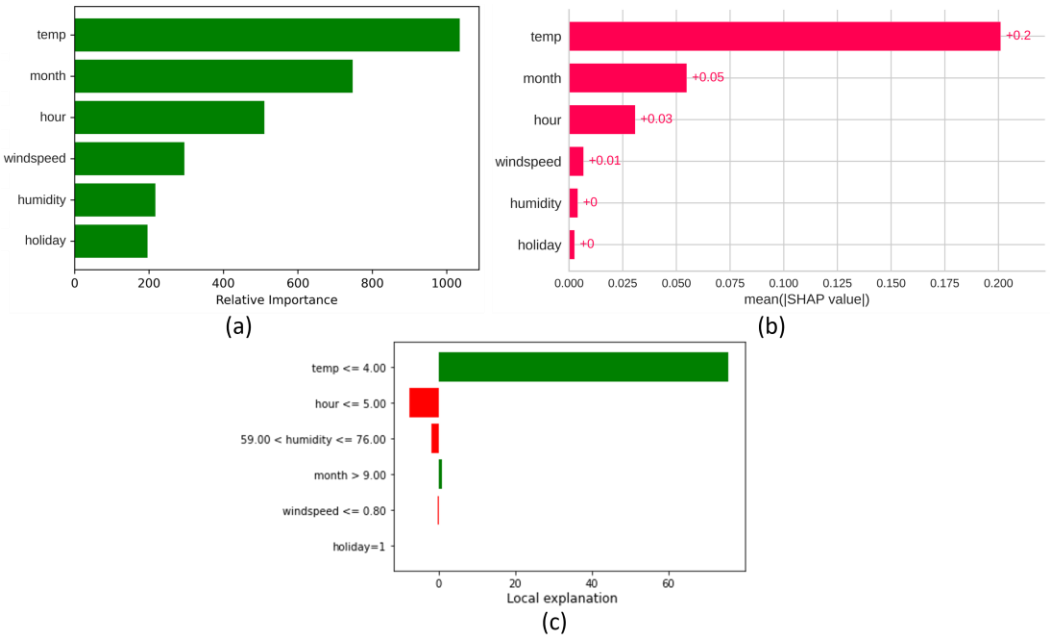

Figure 7: Feature importance analysis for the heat usage prediction model

The previous experiment indicated that the temperature and month features greatly impacted the model's predictions, but it did not explain exactly how the model was affected. As a result, PDP, was implemented in order to demonstrate a feature's marginal effect on the models' prediction.

Figure 8 shows how temperature and month together impact heat usage in the form of contour lines. Contour was proved to work best for analyzing the impact of continuous features in the PDP interaction plot[38].

The contour lines, ranging from 0.000 to 150.000, indicate how specific ranges of the two features affect heat usage. A higher value of the contour line implies a greater impact of the two features on heat usage. For example, during the summer season when the average temperature is above 22°C, the features have a negative influence on the model prediction, resulting in an average heat demand of less than 50 Gcal and a contour line value of under 25.000. On the other hand, contour line values greater than 125.000, corresponding to the winter season with an average temperature of fewer than 2°C, positively impact the model prediction leading to the average heat usage of over 120 Gcal.

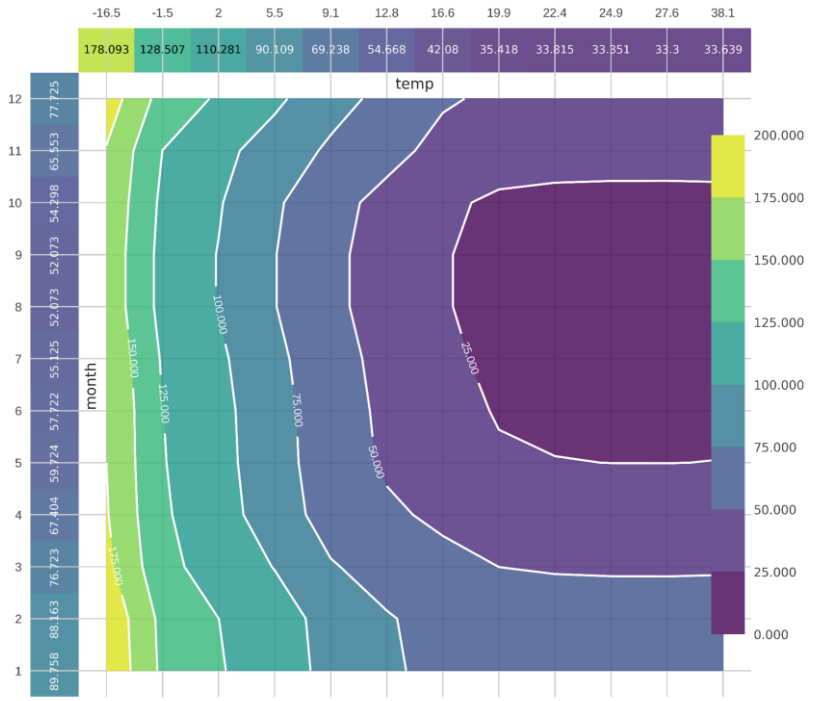

Figure 8: PDP interaction plot for the temperature and month features

Figure 9 illustrates how the temperature feature affected the heat demand through the distribution of the actual heat demand via fixed values of the temperature variable. It was observable that the hourly heat load achieved the biggest average value, which was approximately 150 Gcal, occurred when the temperature feature was between -16.5°C to -0.6 °C, indicating the winter season. Moreover, the hourly heat demand gradually dropped when the temperature rose. The lowest hourly heat demand, around 21 Gcal, was recorded when the temperature ranged from 26.9°C to 38.1°C , which corresponds to the summer season.

Based on the data, we can conclude that the hourly heat demand is directly proportional to the temperature. In the summer, DHW accounts for the majority of the heat demand. In contrast, both DHW and SH contribute to the heat demand during the winter. Additionally, the hourly heat demand is higher during the winter, with temperatures below 10°C, and lower during the summer, with temperatures above 26°C.

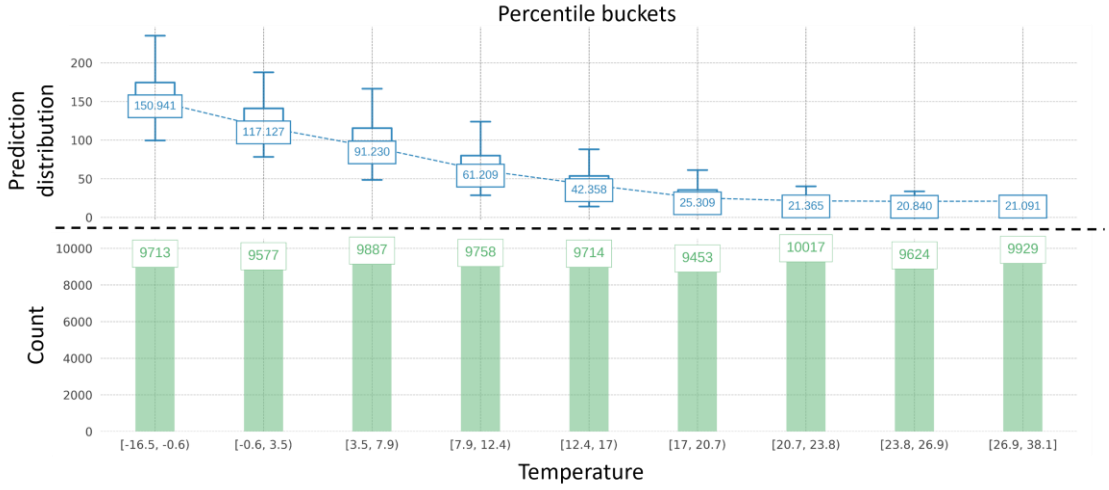

Figure 9: Actual predictions plot for the temperature variable. Distribution of the actual

prediction via different variable values

Comparison with Similar Studies

Numerous studies have been conducted in the past to predict and analyze DH head demand. However, direct comparisons with these studies are difficult due to differences in DH network designs, input data, and architecture implementations or experimental setups. We use operational data from DHS to predict heat usage patterns and compare our results using the XGBoost model, which exhibits the best prediction performance. The recorded MAE value from this study was 15%, which is smaller than the reported MAE of 18.07% by Huang et al[34]. In addition, the computed evaluation metrics are also superior to the following reseach[35,36]. Specifically, the proposed XGBoost model outperforms the study suggested by Ivanko et al[36] in terms of MSE and correlation coefficient, achieving 12% and 0.95 on the testing set, respectively, compared to MSE of 45.04% and a coefficient of determination of 0.81. In terms of the correlation coefficient, the XGBoost method also shows better hourly prediction performance than the ANN model proposed by Bünning et al[35], with a correlation coefficient of 0.95 for one hour compared to 0.88.

Discussion

This section provides a discussion based on our approach and the obtained results. Furthermore, a discussion about the interpretability of the study is also presented.

Model Performance

To establish the best heat demand prediction model, five different models were evaluated with varying sizes of training datasets. Then, three evaluation metrics (MSE, MAE, and R^2) were calculated. Figure 5 demonstrates the learning trend of these models as the number of training samples increases. When the training dataset size is less than 2000, MLP and SVR exhibit the highest accuracy. However, these models have drawbacks such as the need for sequential data and extended training times, making them more suitable for applications that can handle longer training periods. On the other hand, for larger training datasets (over 2000 samples), the accuracy of the three boosting algorithms is higher. Boosting algorithms, such as AdaBoost and XGBoost, are more appropriate for granular control and frequent updating due to their short training time, stability, and forecasting accuracy. Nonetheless, all models can generate predictions swiftly (within a second) after being trained. Hence, the time required for training and retraining the models is the primary constraint for their overall implementation.

Collinearity, which refers to the correlation between predictor variables, always exists in real-world data[24]. However, the impact of collinearity on prediction models varies due to differences in principles. Previously, several approaches have been introduced to address collinearity problems, such as pre-selection based on thresholds, clustering predictors, and regularization techniques. Regularization is a method used to reduce the complexity of the SVM model and prevent overfitting[14]. Similarly, boosting-based models like AdaBoost, XGBoost, and LightGBM can effectively handle multicollinearity problems by adjusting the number of variables sampled at each split[23], which acts as a regularization parameter. In contrast, MLP's ability to withstand collinearity is relatively weak, which may explain its relatively low accuracy.

The way in which heat is distributed varies greatly depending on the size of the DH network, and the proposed framework is appropriate for smaller networks where the behavior of customers has an impact on the load pattern. It is possible to apply the framework to other

small-scale DH networks, in order to anticipate the hourly heat demand, as long as records of the hourly heat demand and environmental factors such as wind speed, humidity and temperature are available.

Interpretability

Model interpretability for AI models refers to the ability to transform the training and testing processes into logical rules. The model's ability to display the significance and ranking of input variables[37] allows it to exhibit interpretability. The interpretability of a predictive model is crucial in evaluating the rationality of heat demands in a DH network. A lack of conformity to accepted principles in variable importance can indicate model instability or system malfunction[4]. Boosting-based methods are highly interpretable as they do not require the interpretation of tree structures by ML professionals, and each decision corresponds to a logical rule[14]. These models can output visual results of variable importance, with the weight and rank of variables differing depending on the model's inherent principles, as displayed in Figure 7. However, temperature and month were consistently the most influential variables, with humidity and holiday having a negligible impact, indicating the limited influence of these variables on heat usage.

On the other hand, SVR and MLP were less interpretable, with MLP being considered a black box method due to its difficulty in identifying the features extracted from each layer of the network. The use of a linear kernel function in SVR leads to a more interpretable model, but models with other kernels can be challenging to interpret[37].

Conclusion

Hourly heat demand forecasting is essential for heating providers to optimize heat production and heat supply operations. This research presents an hourly heat usage prediction system that is based on standard regression algorithms, and it systematically investigates the input features' influence on the models' outcomes.

First, additional weather information, which includes the outdoor temperature, wind flow speed, and humidity of the corresponding hourly historical heat demand, were extracted during the data collection process, and they were used as the input features. After that, various data preprocessing procedures were implemented in order to clean the dataset. The preprocessed dataset was utilized in order to thoroughly analyze the common heat demand patterns. Finally, the dataset was inputted into five well-known regression algorithms, namely SVR, MLP, XGBoost, AdaBoost, and LightGBM, in order to determine what model is the most suitable for the heat usage prediction task based on standard evaluation metrics.

The XGBoost model achieved the lowest MSE via various experiments, which was less than 0.01, and it was robust when the number of samples in the training dataset increased. Finally, various XAI methods, such as SHAP and PDP were applied in order to thoroughly analyze how the model gave a particular prediction. The results showed that temperature and time-related variables are the most critical features that contribute to the model's predictions.

More attention will be directed in the future toward novel heat load prediction techniques, such as multi-step ahead prediction. In addition, collecting a larger dataset with additional variables can improve the performance and efficiency of the model.

Data Availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Code Availability

The source code for all the analyses presented in this study can be found on these GitHub repositories: https://github.com/minhdl93/HeatLoadAnalysis

Acknowledgements

This work was supported by the Basic Science Research Program via the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2020R1A6A1A03038540) and by a grant (20212020900150) from "Development and Demonstration of Technology for Customers Bigdata-based Energy Management in the Field of Heat Supply Chain" funded by Ministry of Trade, Industry and Energy of Korean government and by the Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2022-0-00106, Development of explainable AI-based diagnosis and analysis frame work using energy demand big data in multiple domains).

Author contributions

H.M. and H.S. acquired the funding and supervised the study. Y.L., L.T., and T.N. performed the data collection, preprocessing, and experimental validation. L.D. wrote the original draft. J.S. revised the manuscript. All authors have read and agreed to publish the version of the manuscript.

Competing interests

The authors declare that there are no conflicts of interest to report in regards to this study.

References

1. Lund, H.; Østergaard, P.A.; Chang, M.; Werner, S.; Svendsen, S.; Sorknæs, P.; Thorsen, J.E.; Hvelplund, F.; Mortensen, B.O.G.; Mathiesen, B.V. The status of 4th generation district heating: Research and results. *Energy* **2018**, *164*, 147-159.
2. Barone, G.; Buonomano, A.; Forzano, C.; Palombo, A. A novel dynamic simulation model for the thermo-economic analysis and optimisation of district heating systems. *Energy Conversion and Management* **2020**, *220*, 113052.
3. Dorotić, H.; Pukšec, T.; Duić, N. Multi-objective optimization of district heating and cooling systems for a one-year time horizon. *Energy* **2019**, *169*, 319-328.
4. Guelpa, E.; Verda, V. Thermal energy storage in district heating and cooling systems: A review. *Applied Energy* **2019**, *252*, 113474.
5. Buffa, S.; Cozzini, M.; D'antoni, M.; Baratieri, M.; Fedrizzi, R. 5th generation district heating and cooling systems: A review of existing cases in Europe. *Renewable and Sustainable Energy Reviews* **2019**, *104*, 504-522.
6. Dang, L.M.; Lee, S.; Li, Y.; Oh, C.; Nguyen, T.N.; Song, H.-K.; Moon, H. Daily and seasonal heat usage patterns analysis in heat networks. *Scientific Reports* **2022**, *12*, 1-12.
7. Xue, P.; Jiang, Y.; Zhou, Z.; Chen, X.; Fang, X.; Liu, J. Multi-step ahead forecasting of heat load in district heating systems using machine learning algorithms. *Energy* **2019**, *188*, 116085.

8. Zhao, Y.; Zhang, C.; Zhang, Y.; Wang, Z.; Li, J. A review of data mining technologies in building energy systems: Load prediction, pattern identification, fault detection and diagnosis. *Energy and Built Environment* **2020**, *1*, 149-164.
9. Guo, Y.; Wang, J.; Chen, H.; Li, G.; Liu, J.; Xu, C.; Huang, R.; Huang, Y. Machine learning-based thermal response time ahead energy demand prediction for building heating systems. *Applied energy* **2018**, *221*, 16-27.
10. Wang, Z.; Hong, T.; Piette, M.A. Building thermal load prediction through shallow machine learning and deep learning. *Applied Energy* **2020**, *263*, 114683.
11. Nageler, P.; Koch, A.; Mauthner, F.; Leusbrock, I.; Mach, T.; Hochenauer, C.; Heimrath, R. Comparison of dynamic urban building energy models (UBEM): Sigmoid energy signature and physical modelling approach. *Energy and buildings* **2018**, *179*, 333-343.
12. Westermann, P.; Deb, C.; Schlueter, A.; Evins, R. Unsupervised learning of energy signatures to identify the heating system and building type using smart meter data. *Applied Energy* **2020**, *264*, 114715.
13. Luo, X.; Oyedele, L.O.; Ajayi, A.O.; Akinade, O.O.; Owolabi, H.A.; Ahmed, A. Feature extraction and genetic algorithm enhanced adaptive deep neural network for energy consumption prediction in buildings. *Renewable and Sustainable Energy Reviews* **2020**, *131*, 109980.
14. Ntakolia, C.; Anagnostis, A.; Moustakidis, S.; Karcanias, N. Machine learning applied on the district heating and cooling sector: A review. *Energy Systems* **2022**, *13*, 1-30.
15. Guelpa, E.; Marincioni, L.; Capone, M.; Deputato, S.; Verda, V. Thermal load prediction in district heating systems. *Energy* **2019**, *176*, 693-703.
16. Rouleau, J.; Gosselin, L. Impacts of the COVID-19 lockdown on energy consumption in a Canadian social housing building. *Applied Energy* **2021**, *287*, 116565.
17. Sukparungsee, S.; Areepong, Y.; Taboran, R. Exponentially weighted moving average—Moving average charts for monitoring the process mean. *Plos one* **2020**, *15*, e0228208.
18. Harrell, F.E. General aspects of fitting regression models. In *Regression modeling strategies*; Springer: 2015; pp. 13-44.
19. Munkhdalai, L.; Munkhdalai, T.; Park, K.H.; Lee, H.G.; Li, M.; Ryu, K.H. Mixture of activation functions with extended min-max normalization for forex market prediction. *IEEE Access* **2019**, *7*, 183680-183691.
20. Okada, S.; Ohzeki, M.; Taguchi, S. Efficient partition of integer optimization problems with one-hot encoding. *Scientific reports* **2019**, *9*, 1-12.
21. Zhang, F.; O'Donnell, L.J. Support vector regression. In *Machine Learning*; Elsevier: 2020; pp. 123-140.
22. Liu, Y.; Liu, S.; Wang, Y.; Lombardi, F.; Han, J. A stochastic computational multi-layer perceptron with backward propagation. *IEEE Transactions on Computers* **2018**, *67*, 1273-1286.
23. Azmi, S.S.; Baliga, S. An Overview of Boosting Decision Tree Algorithms utilizing AdaBoost and XGBoost Boosting strategies. *Int. Res. J. Eng. Technol.* **2020**, *7*.
24. Bentéjac, C.; Csörgő, A.; Martínez-Muñoz, G. A comparative analysis of gradient boosting algorithms. *Artificial Intelligence Review* **2021**, *54*, 1937-1967.
25. Chicco, D.; Warrens, M.J.; Jurman, G. The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation. *PEERJ. COMPUTER SCIENCE.* **2021**, *7*, e623.
26. Sundararajan, M.; Najmi, A. The many Shapley values for model explanation. In Proceedings of the International conference on machine learning, 2020, 9269-9278.
27. Moosbauer, J.; Herbinger, J.; Casalicchio, G.; Lindauer, M.; Bischl, B. Explaining Hyperparameter Optimization via Partial Dependence Plots. *Advances in Neural Information Processing Systems* **2021**, *34*, 2280-2291.
28. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. and Vanderplas, J. Scikit-learn:

682 Machine learning in Python. *Journal of machine Learning research* **2011**, *12*, 2825-
683 2830.
684 29. Idowu, S., Saguna, S., Åhlund, C. and Schelén, O. Applied machine learning:
685 Forecasting heat load in district heating system. *Energy and Buildings* **2016**, *133*,
686 478-488.
687 30. Bourdeau, M., qiang Zhai, X., Nefzaoui, E., Guo, X. and Chatellier, P. Modeling and
688 forecasting building energy consumption: A review of data-driven techniques.
689 *Sustainable Cities and Society* **2019**, *48*, 101533.
690 31. Saloux, E. and Candanedo, J.A. Forecasting district heating demand using machine
691 learning algorithms. *Energy Procedia* **2018**, *149*, 59-68.
692 32. López, M., Sans, C., Valero, S. and Senabre, C. Classification of special days in short-
693 term load forecasting: the Spanish case study. *Energies* **2019**, *12(7)*, p.1253.
694 33. Noussan, M., Jarre, M. and Poggio, A. Real operation data analysis on district
695 heating load patterns. *Energy* **2017**, *129*, 70-78.
696 34. Huang, Y., Yuan, Y., Chen, H., Wang, J., Guo, Y. and Ahmad, T. A novel energy
697 demand prediction strategy for residential buildings based on ensemble learning.
698 *Energy Procedia* **2019**, *158*, 3411-3416.
699 35. Bünning, F., Heer, P., Smith, R.S. and Lygeros, J. Improved day ahead heating
700 demand forecasting by online correction methods. *Energy and Buildings* **2020**, *211*,
701 109821.
702 36. Ivanko, D., Sørensen, Å.L. and Nord, N. Selecting the model and influencing
703 variables for DHW heat use prediction in hotels in Norway. *Energy and Buildings*
704 **2020**, *228*, 110441.
705 37. Minh, D., Wang, H.X., Li, Y.F. and Nguyen, T.N. Explainable artificial intelligence: a
706 comprehensive review. *Artificial Intelligence Review* **2022**, 1-66.
707 38. Greenwell, B.M. pdp: an R Package for constructing partial dependence plots. *R J.*
708 **2017**, *9(1)*, 421.
709 39. Visani, G., Bagli, E., Chesani, F., Poluzzi, A. and Capuzzo, D. Statistical stability
710 indices for LIME: Obtaining reliable explanations for machine learning models.
711 *Journal of the Operational Research Society* **2022**, *73(1)*, 91-101.