

Modelling Hourly Thermal Energy Demand: A Machine Learning Approach of Residential District Heating Substations in Turin

C. Monzani¹, G. Cerino Abdin¹, G. Montanari^{1,2}, A. Poggio¹

¹ Department of Energy Engineering
Politecnico di Torino
Corso Duca degli Abruzzi, 24 – 10129 Turin (Italy)

²IUSS PAVIA. IUSS - Istituto Universitario di Studi Superiori

Abstract.

This paper explores the application of the XGBoost machine learning model for forecasting the hourly thermal demand in District Heating Systems, aligning with the European Union's ambitious sustainability targets as outlined in the Renewable Energy Directive (RED) and the Energy Efficiency Directive (EED). Accurate forecasts of thermal demand are crucial for enhancing the efficiency of district heating systems through the integration of renewable energy sources and the adoption of waste heat recovery, thereby contributing significantly to achieving climate neutrality by the year 2050. This study presents a dual approach to forecasting: at the individual building level, and at an aggregated level by considering the average characteristics of the served building stock. Through a comprehensive case study of the Turin district heating system (Italy), which comprises hourly data from approximately 200 heat exchange substations across nine heating seasons, this research evaluates the comparative effectiveness of different forecasting approaches in terms of prediction accuracy and computational efficiency. The findings aim to guide district heating operators and planners in selecting the most suitable forecasting approach based on available input information, desired accuracy, and computational constraints, contributing to the strategic planning and development of sustainable and efficient district heating systems.

Key words. District Heating Systems (DHS), XGBoost model, Thermal Demand Forecasting, Renewable Energy Integration

1. Introduction

In the context of the European Union's commitment to environmental sustainability, legislative frameworks such as the Renewable Energy Directive (RED) [1] and the Energy Efficiency Directive (EED) [2] delineate ambitious objectives to transition towards a more sustainable and efficient energy system. These directives highlight the pivotal role of District Heating Systems (DHS) in achieving the EU's climate neutrality goals. By facilitating the integration of renewable energy sources and waste heat recovery, DHS serve as a cornerstone in

urban energy strategies aimed at reducing carbon footprints and enhancing energy efficiency. The flexibility and scalability of DHS allow for a more effective incorporation of diverse energy sources, underscoring the system's significance in meeting Europe's stringent 2030 and 2050 sustainability targets. The role of DHSs is particularly central in densely populated urban areas where energy demand is particularly high and land availability for the installation of renewable technologies reduced. A critical aspect of optimizing DHS operations and facilitating the integration of renewable sources, as well as waste heat and storage systems, is the development of accurate hourly load profiles. Such detailed profiling enables the simulation of DHS functioning and sizing throughout the year, ensuring that the systems can effectively meet energy demands while maximizing the use of renewable and waste heat sources. The ability to predict thermal energy demand with high precision is essential for planning investments and operational strategies of DHS, thereby ensuring their alignment with the sustainability criteria set forth by European directives.

In this context, the application of Artificial Intelligence (AI) in Energy Planning Models (EPMs) has gained prominence. Debnath and Mourshed (2018) [3] underscore the key role of forecasting in EPMs, noting the increasing reliance on AI and machine learning models to enhance prediction accuracy and system adaptability. Their comprehensive review identifies a wide array of forecasting methods, with a particular emphasis on the effectiveness of models such as Artificial Neural Networks (ANN), Support Vector Machines (SVM), and XGBoost in energy demand forecasting. Among these, XGBoost stands out for its superior performance in terms of computational efficiency and predictive accuracy, making it an invaluable tool in the context of DHS optimization.

Academic research has applied AI models to forecast thermal demand in DHSs, achieving favorable results in predicting load profiles for the subsequent 24 or 72 hours using hourly data. Gong et al. [6] compared various deep

learning and machine learning models for predicting thermal load in DHSs. These models, built using hourly data, were utilized to forecast the load profile for the following day. Results indicated that the deep neural network model performed the best followed by the support vector machine, recurrent neural network, multilayer perceptron, random forest, decision tree, and long short-term neural network. Additionally, parallel studies have applied these models to predict monthly data to obtain annual consumption values. Maljkovic et al. [7], in their study, examined regression tree models, random forest, and support vector machines applied at various levels of building feature knowledge, aiming to determine if simplifying the required model information maintained an accurate prediction level. Specifically, the random forest model demonstrated the most effective performance, even in scenarios with limited contextual knowledge of consumption data. The application of XGBoost in forecasting thermal demand within DHS has been notably explored in studies such as Runge and Saloux (2023) [4], which affirm the model's efficacy in achieving reduced forecasting errors and fast training times. Similarly, Xue et al. [5] demonstrated the applicability of XGBoost alongside SVR and DNN models, further validating the ensemble tree algorithm's capabilities in predicting thermal load profiles for subsequent hours and days with remarkable precision. Current research has therefore focused on identifying accurate forecasting models for predicting thermal demand within 24/72 h as a support tool in the operational phase. Equally important, however, is to define a planning tool that allows for the understanding of hourly variations in system thermal demand over the entire season. The ability to accurately forecast thermal demand on an hourly basis provides DHS operators with valuable information for dimensioning generation plants, selecting the most appropriate energy sources and planning maintenance activities.

By leveraging the advanced predictive capabilities of XGBoost in conjunction with real consumption data, this paper aims to address the pressing need for demand modelling in the DHS for planning purposes. The paper's objective is to delve into the efficacious methodologies for forecasting the hourly thermal demand of a DHS adopting a data-driven approach [8]. This investigation will apply the XGBoost model in two distinct manners: firstly, at the level of individual buildings, assuming knowledge of the characteristics of each building, and secondly, at an aggregate level, considering the average characteristics of the building stock served. The paper aims to compare these forecasting approaches in terms of prediction accuracy and the computational time required. Each approach's applicability is determined based on the availability of input information, the required accuracy level for specific purposes and applications. The paper is structured to provide a description of the XGBoost model, and the performance evaluation criteria used. This methodology is then applied to a case study of the Turin DHS, which comprises hourly consumption data from approximately 200 heat exchange substations over nine heating seasons. The case study serves as a real-world application to demonstrate the effectiveness of the proposed forecasting methodologies.

2. Method

This chapter delves into the application of the XGBoost machine learning algorithm (XGB) for the accurate prediction of hourly heat demand in DHS. It discusses the algorithm's parameter optimization to enhance predictive precision, the evaluation of model performance through RMSE and R^2 , and the structure of the dataset. Additionally, it outlines the multi-scale analysis methodology used to test the model's capability for generalization across various input management approaches, ensuring alignment with the broader goals of district heating management.

A. XGBoost machine learning algorithm

XGB is an advanced gradient boosting framework that improves predictive accuracy through the sequential construction of decision trees. Using the Python library, the algorithm is optimised with a careful selection of parameters, including 'n_estimators', 'learning_rate', 'max_depth'. These parameters are calibrated to balance model complexity and prevent overfitting, while ensuring sensitivity to outliers and improving predictive efficiency.

B. Performance Evaluation

Root Mean Square Error (RMSE) is utilized as the primary metric to evaluate prediction accuracy, crucial for refining XGB in energy demand forecasting. The importance of RMSE as an effective performance indicator is highlighted in the work of Wei et al. (2019) [9], emphasizing its critical role in assessing predictive models within the energy sector. Through the targeted optimization of key parameters – 'n_estimators', 'learning_rate', 'max_depth', etc. – aimed at minimizing RMSE, the model's precision is enhanced while ensuring its reliability and avoiding overfitting. This methodology adheres to established best practices in model evaluation, striving for a balance between accuracy and generalizability in thermal demand forecasting.

The performance of the trained Hourly Model XGBoost (HM) was assessed using two key metrics: the percentage relative error and the coefficient of determination (R^2). The percentage relative error provides a direct measure of the deviation of the model's predictions from the actual values, offering an intuitive indication of the accuracy of the predictions. R^2 evaluates the model's ability to capture the variability of the data, with values closer to 1 indicating a better fit of the model to the observed data.

C. Dataset structure

XGB is trained on a complex dataset that encompasses climatic, temporal, construction-related, and plant-specific variables:

- climatic: hourly external temperature, and outdoor temperature of the previous 6 hours;
- calendar: time of day, day of the week, and month are transformed into cyclic variables through sine and cosine functions, preserving the cyclic nature of energy consumption. Holidays are identified by a

boolean value. The thermal season is categorised through numerical values;

- construction-related: data such as heated volume, number of dwelling units, number of floors, construction period, and the surface-to-volume ratio (S/V) define the physical characteristics of buildings;
- plant-specific: the installed power of the thermal substation.

D. Modelling methodology and XGB application

As illustrated in the figure, the dataset was divided into two distinct subsets: one for training and one for validation (step 1). The stratification criterion was selected based on the year and identification code of each building, ensuring that both datasets maintained a varied and balanced representation of housing characteristics. This methodology ensured that the variability between buildings was consistent, allowing XGB to learn and validate itself across a several range of cases. XGB is then trained and calibrated (step 2).




Fig. 1. Model's architecture

The goal is to maximize prediction accuracy, while maintaining the model's ability to generalize to unseen data. The HM is tested on a set of buildings not involved in the training dataset, with the aim of testing the generalization capability (step 3).

1) *Multi-scale analysis*: the performance evaluation of the hourly model was conducted through three different approaches:

- Single approach (S): the inputs provided to the model are separated for each Test Building (TB). This results in an hourly heat load curve for each building;
- Aggregate approach (A): similar to the Single approach, this approach starts with separate inputs for each TB. The data obtained from the forecast are aggregated to obtain an overall hourly heat profile for all buildings;
- Weighted average approach (W): in contrast to the first two approaches, the model inputs are already aggregated, using weighted average values based on the thermal capacity of the district

heating substation installed in the buildings. This approach directly provides the overall hourly heat profile without the need for further aggregation steps.

This strategy is crucial as it allows us to assess the model's ability to harmonise individual errors, potentially improving accuracy in the aggregate. It also aligns the analysis with the operational objectives of overall district heating management, where the overall view prevails over the microscopic detail of individual buildings.

3. Result and discussion

A. Case study descriptions

XGB was trained and validated on the basis of hourly heat consumption data from 205 heat exchange substations connected to the DHS in Turin (Italy), all of which were residential buildings. These records, initially detailed at 6-minute intervals, were aggregated to generate an hourly thermal load profile.




Fig. 2. Buildings' characteristics

Information on the heated volume and the thermal capacity of the installed district heating substation is made available by the DHS operator. The construction information was defined during previous research work [10][11] and is based on elaborations of the ISTAT database of the last Italian national census [12]. There are potential inaccuracies in the information regarding the construction characteristics of buildings and heated volumes. The estimate of the heated volume is made by the district heating operator based on the gross dimensions of the building. This approach does not consider possible unheated interior spaces, such as common areas (stairs, entrances, etc.), which could significantly affect the accuracy of the estimate. Similarly, the information from the ISTAT database, being based on questionnaires filled in by citizens and not by experts, is subject to potential compilation errors. These factors introduce a level of uncertainty into the information used to predict thermal consumption, underlining the importance of considering these limitations in the analysis and application of predictive models. The ISTAT database is available nationwide and thus makes the application of the approach replicable both in Italy and in all contexts in which similar statistics are implemented. Approximately 70% of the buildings involved in the analysis have heated

volumes between 1,200 and 5,700 m³ correlating to district heating substations with a thermal power capacity of 50 to 150 kW. The characteristic construction period of the buildings is between the 1950s and 1980s, featuring a transmission loss coefficient between 0.5 and 0.7 W/m³K. The dataset spans from November 2012 through April 2021, offering a comprehensive view of consumption patterns over the years.

After establishing the optimal parameters of XGB, which will be elaborated upon in the subsequent section, we move on to the testing phase. This phase involves applying the model to a set of entirely new buildings (TB). A total of 28 buildings are included, for which weighted average characteristics based on the DH substation installed capacity have been calculated (Table I).

Table I. – TB parameters

total heating volume	m ³	158,000
total DH substation installed power	kW	4'950
transmission loss coefficient	W/m ³ K	0.604
S/V	m ⁻¹	0.338
construction period	-	1970
number of dwellings	-	38
number of floor	-	8

TB were integrated into the DHS starting from the 2018/2019 heating season. The model's hourly heat consumption forecast is extended across the three available thermal seasons. Concurrently, an analysis of the actual annual energy consumption of TB was conducted aiming to detect any consumption variations that might be attributed to the lockdown due to COVID-19. Contrary to tertiary buildings, which experienced significant fluctuations in energy consumption due to the pandemic, TB exhibited no notable changes in consumption during the lockdown period. This step is critical to ensuring that the model's performance accurately mirrors consumption patterns under standard conditions, unaffected by exceptional events.

B. XGB performances

In refining the parameters of XGB, the GridSearchCV function from the scikit-learn library was employed to conduct an exhaustive search for optimal hyperparameter combinations. This method reviewed various configurations of *n_estimators*, *max_depth*, and *learning_rate*, seeking to identify the one that provided the best balance between a low RMSE and a high R² without excessively prolonging the execution time. The model was trained and validated using a substantial dataset of 2.8 million hourly records, which was split into 60% for training and 30% for validation.

As demonstrated in the Figure 3, increasing the number of estimators enhanced the R² value until reaching a plateau; beyond 5,500 estimators, further improvements in RMSE were marginal, while execution times significantly increased. Consequently, the selected parameters were set at 5,500 estimators, with a *max_depth* of 5 and a *learning_rate* of 0.1. These parameters were affirmed as optimal through a cross-validation process, ensuring that the model remained generalizable and reliable by avoiding overfitting or underfitting.




Fig. 3. XGB performances

C. HM Performances

The test database for evaluating the performance of trained Hourly Model XGBoost on TB varied depending on the approach employed. In the Single (S) and Aggregate (A) approaches, the test database consisted of approximately 200,000 records each. In contrast, for the Weighted Average (M) approach, the number of records was reduced to around 6,400 due to the aggregation of inputs. Execution times were notably reduced for all approaches; the A and S approaches required an execution time of about 8 seconds, while the M approach, benefiting from the smaller volume of input data, achieved execution times of around 1 second. It is important to consider that the test dataset is composed of a small number of data, the application of the model to an entire building stock served by a medium to large DHS over a reference heating season would lead to higher runtimes. In this case, the application of a simplified model that reduces the time by almost 90% is the most suitable choice.




Fig. 4. approach performances

Figure 4 provides a detailed analysis of HM performance across three different input management approaches - Single (S), Aggregate (A) and Weighted (W) - in predicting the heat demand for a DHS.

The model shows excellent performance in all three approaches. Specifically, Approach S proved to be able to keep the relative percentage error within $\pm 5\%$ for 39% of the records. However, this approach tends to overestimate consumption, exceeding 100% error in 10% of cases. These errors are particularly concentrated during the

changeover periods of the heating season and during the evening hours, when buildings can stop consuming heat beforehand and the model is unable to simulate it. The variety of non-technical or climatic behaviours, such as times and dates of switching on and off of the systems that may change from period to period or from season to season, is a challenge for the model, which is less accurate to reality in a minority of cases. However, through post-processing that corrects errors by knowing the dates and times of the start and end of heating service for each building, it would be possible to improve the accuracy of the model.

With Approach A, by aggregating the hourly consumption of all buildings, errors tend to be smoothed out. This is particularly advantageous since the objective of the study is to define the hourly demand of a DHS; having optimal performance at the aggregate level is therefore more important than at the level of individual housing units. The relative percentage error remains within $\pm 5\%$ in a slightly smaller percentage of cases (32%), but the percentage of events with an error within $\pm 20\%$ increases to 68% of cases.

Approach M shows a narrower error distribution: 43% of cases fall within $\pm 5\%$ error, and 90% within $\pm 50\%$. In contrast to the other two approaches, Approach M overestimates more frequently, with less than 3% of cases showing overestimation. This type of error mainly occurs between 9pm and 11pm, when some buildings still show thermal consumption, but the model considers them to be zero.

Looking at part (b) of the graph, which represents the aggregate consumption on a daily basis, all three approaches show very good performance, with coefficients of determination (R^2) above 0.9. Of these, Approach A stands out as the best performing.

- 1) *Analysis of feature importance*: one of the significant capabilities of XGBoost is its feature importance classification, which identifies how each variable influences the final output. Analyzing the importance of variables in the XGB for predicting hourly heat demand, a marked preference for the cosine of the month over the sine is evident. This distinction can be attributed to the ability of the cosine to provide a more cohesive and direct representation of typical heating months. The cosine peaks in January and December, maintaining positive values above zero in the autumn and winter months, which correspond to the period of greatest heating demand. Furthermore, the cosine value is the same in October and April, corresponding to the start and end months of the season. This characteristic makes it particularly suitable for signaling the continuous need for heating energy during the cold season. On the contrary, the sine shows an oscillation that causes it to assume positive values between January and April and negative values between October and December. This variation implies a less intuitive representation of the heating months, as the signal changes from positive to negative precisely in the critical period for heating demand. This discontinuity could make

the sine less effective in indicating a uniform heating need.




Fig. 5. Model features importance

The cosine of the month is followed by the importance of the average hourly outdoor temperature and the power of the installed heating substation.

Building parameters appear to have less weight, which could indicate a limitation of the training dataset due to the relative homogeneity of the residential building sample considered. This suggests that the variability within the dataset may not be large enough for the model to effectively learn the impact of these characteristics on energy demand.

In contrast, variables such as the day of the week and the distinction between weekdays and holidays are of marginal importance. This is consistent with the residential context of the analysis: the buildings being primarily permanent dwellings in a non-tourist area. The constancy in residential patterns, regardless of the type of day, shows that the thermal behavior of buildings is moderately influenced by these factors in an urban context not linked to seasonal fluctuations in tourism.

4. Conclusion

In this study, the XGBoost machine learning algorithm was used to predict the hourly heat demand within district heating systems (DHSs). The algorithm was trained and validated with hourly heat consumption data from 205 residential substations of the Turin DHS (Italy) over nine heating seasons. The hourly model was tested on a subset of 28 buildings for three heating seasons. The primary objective was to accurately simulate the hourly load profile for effective DHS management, which is crucial for the integration of renewable energy sources, waste heat and thermal storage optimisation. The types of input needed by the model to predict demand are climatic (outdoor temperatures), calendar (month, day, hour), construction-related and plant-specific of the building. The information on construction characteristics is derived from analyses of ISTAT databases whose data were collected nationwide through census questionnaires. This makes the application of the model replicable throughout the country and in countries where census data are similarly available.

Three distinct approaches were used to evaluate the performance of the model. In the Single (S) approach the construction and plant characteristics of each of the 28 test buildings were provided as input to the model. The Approach S produced high accuracy, with approximately 39% of the records falling within an error range of $\pm 5\%$. However, high over-estimates occurred in 10% of the records, particularly during seasonal transitions and evening hours, when heating is used with variably logic. In the Aggregate Approach (A) the inputs to the model coincide with those used in Approach S. Aggregating the data from all buildings improved the distribution of errors and resulted in more accurate daily heat consumption predictions. This approach aligns well with the objective of managing the overall hourly demand of a DHS, thus being the most effective strategy demonstrated in this study. Finally, with the Weighted Average (M) Approach, the required input data coincide with the average construction and plant characteristics of the building stock considered. This method also gave promising results, with a narrower error range, but showing an understated percentage of underestimation of the hourly demand. Approach A therefore appears to be the one with the best performance. Approach M proves its usefulness when detailed information on the building stock is less available. Approach M also appears to be the most appropriate choice in the case of forecasting data on an extensive building stock of a medium to large DHS where model run times are reduced by almost 90% compared to approach A

XGB made it possible to weight the importance of the inputs provided to the model. Specifically, the model gave higher priority to the cosine of the month, the outdoor temperature, and the thermal capacity of the substations. The lesser emphasis on building characteristics, and on the building's transmission loss coefficient, is a limitation of the model for estimating the impact of energy efficiency measures in reducing thermal demand.

The study focuses exclusively on residential buildings, which tend to have more predictable energy consumption patterns than tertiary structures. In order to generate a comprehensive hourly heat demand profile for a DHS, it is essential to incorporate different building categories. Future research aims to expand the application of the model to include different building types, ultimately creating a heat load profile that reflects the entire DHS.

This research successfully demonstrated the applicability of XGBoost to predict hourly thermal energy consumption at various levels of detail on the building stock. The study offers a significant hourly thermal demand forecasting tool for energy planning. The hourly forecast is crucial for planners and operators of DHSs to design generation plants and select the most appropriate energy sources in order to plan investments and decarbonization strategies. The tool can also be used in operations to plan maintenance activities.

References

- [1] Directive (EU) 2023/2413 of The European Parliament and of the Council (2023), "Amending Directive (EU) 2018/2001, Regulation (EU) 2018/1999 and Directive 98/70/EC as regards the promotion of energy from

renewable sources, and repealing Council Directive (EU) 2015/652", 2023

- [2] Directive (EU) 2023/1791 of The European Parliament and of The Council (2023) on energy efficiency and amending Regulation (EU) 2023/955
- [3] K. B. Debnath, M. Mourshed, "Forecasting methods in energy planning models", *Renewable and Sustainable Energy Reviews* (2018), Vol. 88, pp. 297-325
- [4] Jason Runge, Etienne Saloux, "A comparison of prediction and forecasting artificial intelligence models to estimate the future energy demand in a district heating system", *Energy* 269 (2023) 126661
- [5] Puning Xue, Yi Jiang, Zhigang Zhou, Xin Chen, Xiumu Fang, Jing Liu, "Multi-step ahead forecasting of heat load in district heating systems using machine learning algorithms", *Energy* (2019) Vol.188
- [6] Mingju Gong, Haojie Zhou, Qilin Wang, Sheng Wang & Peng Yang, "District heating systems load forecasting: a deep neural networks model based on similar day approach", *Advances in Building Energy Research* (2020), 14:3, 372-388
- [7] Danica Maljkovic, Bojana Dalbelo Basic, "Determination of influential parameters for heat consumption in district heating systems using machine learning", *Energy* (2020), vol. 201 117585
- [8] Kadir Amasyali, Nora M. El-Gohary A., "A review of data-driven building energy consumption prediction studies", *Renewable and Sustainable Energy Reviews* (2018) Vol.81 pp.1192
- [9] N. Wei, C. Li, X. Peng, F. Zeng, X. Lu "Conventional models and artificial intelligence-based models for energy consumption forecasting: A review", *Journal of Petroleum Science and Engineering* (2019), vol. 181 106187
- [10] G. Cerino Abidin, "Modelli per la pianificazione locale di filiere energetiche da biomassa forestale", *Doctoral Thesis*, 2015
- [11] Regione Piemonte, "Piano Regionale di Qualità dell'Aria (PRQA) - Allegato C - Analisi dei consumi energetici e riduzioni emissive ottenibili", 2019
- [12] ISTAT. [Online]. Available: <http://dawinci.istat.it/>.

Nomenclature

DHS	District Heating System
HM	trained Hourly Model XGBoost
R ²	coefficient of determination
RMSE	Root Mean Square Error
S/V	loss surfaces to heating volume ratio
XGB	XGBoost machine learning algorithm
A	Aggregate Approach
S	Single Approach
W	Weighted Average Approach

Acknowledged

This paper and related research have been conducted during and with the support of the Italian inter-university PhD course in sustainable development and climate change (link: www.phd-sdc.it)