



energies



Article

Optimizing Space Heating in Buildings: A Deep Learning Approach for Energy Efficiency

Fernando Almeida, Mauro Castelli, Nadine Corte-Real and Luca Manzoni



<https://doi.org/10.3390/en18102471>

Article

Optimizing Space Heating in Buildings: A Deep Learning Approach for Energy Efficiency

Fernando Almeida ^{1,*} , Mauro Castelli ¹ , Nadine Corte-Real ¹ and Luca Manzoni ² 

¹ NOVA Information Management School (NOVA IMS), Universidade NOVA de Lisboa, 2829-516 Caparica, Portugal; mcastelli@novaims.unl.pt (M.C.); nreal@novaims.unl.pt (N.C.-R.)

² Dipartimento di Matematica, Informatica e Geoscienze, Università degli Studi di Trieste, 34127 Trieste, Italy; lmanzoni@units.it

* Correspondence: m20200957@novaims.unl.pt

Abstract: Building energy management is crucial in reducing energy consumption and maintaining occupant comfort, especially in heating systems. However, achieving optimal space heating efficiency while maintaining consistent comfort presents significant challenges. Traditional methods often fail to balance energy consumption with thermal comfort, especially across multiple zones in buildings with varying operational demands. This study investigates the role of deep learning models in optimizing space heating while maintaining thermal comfort across multiple building zones. It aims to enhance heating efficiency by developing predictive models for building temperature and heating consumption, evaluating the effectiveness of different deep learning architectures, and analyzing the impact of model-driven heating optimization on energy savings and occupant comfort. To address this challenge, this study employs Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Transformer models to forecast area temperatures and predict space heating consumption. The proposed methodology leverages historical building temperature data, weather station measurements such as atmospheric pressure, wind speed, wind direction, relative humidity, and solar radiation, along with other weather parameters, to develop accurate and reliable predictions. A two-stage deep learning process is utilized: first, temperature predictions are generated for different building zones, and second, these predictions are used to estimate global heating consumption. This study also employs grid search and cross-validation to optimize the model configurations and custom loss functions to ensure energy efficiency and occupant comfort. Results demonstrate that the Long Short-Term Memory and Transformer models outperform the Gated Recurrent Unit regarding heating reduction, with a 20.95% and 20.69% decrease, respectively, compared to actual consumption. This study contributes significantly to energy management by providing a deep learning-driven framework that enhances energy efficiency while maintaining thermal comfort across different building areas, thereby supporting sustainable and intelligent building operations.

Keywords: optimization; heating; comfort; thermal; machine learning



Academic Editor: Vincenzo Costanzo

Received: 20 March 2025

Revised: 25 April 2025

Accepted: 5 May 2025

Published: 12 May 2025

Citation: Almeida, F.; Castelli, M.; Corte-Real, N.; Manzoni, L. Optimizing Space Heating in Buildings: A Deep Learning Approach for Energy Efficiency. *Energies* **2025**, *18*, 2471. <https://doi.org/10.3390/en18102471>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Building energy consumption represents a significant fraction of global energy use, especially in heating, ventilation, and air conditioning (HVAC). According to the International Energy Agency (IEA), more than 40 percent of the world's total energy is consumed by buildings [1], and nearly all of that is used for space heating in colder climates. As energy

efficiency in building environments becomes increasingly important, significant research has been devoted to optimizing energy consumption while retaining occupant comfort.

Space heating consumption (SHC) optimization in this broader domain is significant [2]. SHC systems can save vast amounts of energy wastage and greenhouse gas emissions and provide sustainability to the built environment [3]. New approaches to predict and optimize SHC have become possible with recent advances in data-driven technologies and the increasing availability of real-time data from intelligent buildings [4]. Building management systems (BMS) time series data coupled with external factors (i.e., weather) and sensor data are used to dynamically control heating loads and improve energy efficiency in these approaches. For example, studies have shown that incorporating real-time weather data, such as temperature, humidity, and wind speed, can improve temperature regulation and lead to a reduction in energy consumption by adjusting heating systems to external conditions more efficiently [5,6]. Additionally, sensor data, including occupancy patterns and indoor temperature readings, allow for adaptive control strategies that maintain comfort while minimizing energy waste. The optimal temperature for occupant comfort typically ranges between 21 °C and 23 °C in winter [7], but this can vary depending on external weather conditions such as ambient temperature and humidity. In colder climates, slightly higher indoor temperatures may be necessary to maintain comfort, whereas in milder climates, lower settings may be sufficient.

A recent interest has been found in integrating machine learning (ML) techniques including deep learning (DL) into the optimization of SHC [8]. According to a global report from the Global Energy & CO₂ Status [5], AI-based models could cut building energy consumption by 10 to 20 percent through efficiency gains and automation of control systems. They have been extensively used for time series forecasting and applied to predicting building temperature dynamics and energy demand [9–11]. The models can learn complicated temporal relations using historical data and make accurate temperature control and SHC forecasting predictions.

However, optimizing SHC while maintaining occupant comfort is a difficult task. The challenge is to balance energy savings and indoor comfort, keeping environments at comfortable temperature ranges since occupancy patterns, weather conditions, and the thermal characteristics of different building zones are inherently variable [12]. Fast temperature changes make the occupant uncomfortable, demanding smooth temperature transitions [13]. These conflicting priorities are managed through real-time decision making with large amounts of data to leverage comfort and energy efficiency.

These challenges are well addressed by machine learning, including deep learning models, as these models can then learn patterns from historical temperature and consumption data and make informed predictions on energy consumption, allowing the adjustment of heating systems dynamically concerning energy efficiency and comfort [14,15]. Additionally, hyperparameter tuning and optimization methods like grid search cross-validation ensure that the models are set up for the best performance.

This study aims to develop and compare DL models for predicting SHC in buildings and optimizing space heating systems. We mainly explore Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Transformer models to forecast temperature in different building zones and predict SHC. In this two-stage approach, temperatures are predicted based on SHC estimation, and a model is selected to minimize energy consumption at an acceptable level of occupant comfort. This study aims to provide input to more sustainable building management practices through real-time data and advanced predictive techniques. The contributions of this study are the following:

- i. This study develops a robust temperature prediction model using LSTM, GRU, and Transformer architectures, specifically tailored to individual building zones. This

- approach ensures accurate thermal management, enhancing occupant comfort while minimizing energy consumption.
- ii. A two-stage machine learning framework is proposed, where temperature predictions are first generated for different building zones and then used to estimate global space heating consumption (SHC). By leveraging these predictions, this study optimizes energy use across multiple zones, ensuring efficient heating distribution.
 - iii. Custom loss functions are introduced to regulate temperature variations while minimizing energy consumption. These loss functions enforce thermal comfort constraints and promote energy efficiency, making them crucial for balancing occupant comfort with sustainable building management.

2. Related Work

In recent years, ML techniques including DL have intensely improved the capability to predict and optimize energy consumption in buildings, especially space heating [13]. Table 1 compares various predictive models applied to forecast temperatures and energy usage within building environments. These models leverage time series data, building characteristics, and external factors like weather conditions to improve accuracy and efficiency [16]. Prior studies have also explored methods to balance energy efficiency with occupant comfort, introducing custom loss functions and optimization strategies to achieve both goals in intelligent energy systems.

2.1. Energy Consumption Forecasting in Buildings

Energy consumption forecasting in buildings has become an essential area of research, as it aids in optimizing HVAC operations, improving energy efficiency, and ensuring occupant comfort. Various predictive models are being explored to accurately forecast energy usage based on time series data, building characteristics, weather conditions, and other influencing factors. These techniques range from machine learning models to more advanced hybrid systems.

One of the prominent approaches in this domain is the use of Long Short-Term Memory (LSTM) networks for predicting power consumption. Mahjoub et al. [17] explored the effectiveness of LSTM in forecasting short-term power consumption, demonstrating that it outperformed other models like GRU and Drop-GRU in terms of accuracy and precision. This model helps in managing power consumption by preventing load peaks. Similarly, in the domain of HVAC systems, Ghahramani et al. [18] employed a systematic analysis of HVAC control policies using simulations, including building-level annual control and zone-level control, to forecast energy savings and improve thermal comfort, finding up to 50% energy savings depending on the climate zone.

In a different approach, Chen et al. [19] applied data-driven models using machine learning algorithms like Random Forest and Support Vector Machines (SVM) to optimize energy consumption in commercial buildings. These techniques were used to ensure that the buildings maintained thermal comfort while achieving significant cost savings throughout their lifecycle. Their results suggest that predictive control strategies based on these models help improve the efficiency of HVAC systems while ensuring energy efficiency.

Lastly, in the context of smart buildings, Almeida et al. [20] proposed an adaptive strategy for optimizing HVAC systems' setpoints to enhance energy efficiency during the heating season. The approach, which used smart thermostat data from European households, estimated energy savings of up to 5.2%. Despite being computationally lightweight and suitable for real-world implementation, it faces challenges in handling data drift and adjusting models without additional fine-tuning.

2.2. Thermal Comfort and HVAC System Optimization

The optimization of HVAC systems, while maintaining thermal comfort, is another crucial research focus. Many studies aim to enhance HVAC system operations to balance energy consumption and thermal comfort, leading to the development of advanced control strategies and personalized comfort profiles. These innovations improve occupant satisfaction while reducing energy wastage.

Talami et al. [21] focused on developing a temperature setpoint optimizer that adjusts dynamically according to weather conditions and occupancy. By integrating external factors, their method showed energy savings compared to using fixed setpoints. The study analyzed HVAC energy consumption by considering varying occupancy rates and outdoor temperatures. It quantified energy reductions by comparing HVAC usage at different occupancy levels against a 100% occupancy baseline. Dynamic setpoints were adjusted between 5 °C and 32 °C to balance heating, free-running, and cooling modes. However, the model focused on air temperature and did not account for mean radiant temperature, air velocity, or relative humidity, which are essential for thermal comfort. The study's limitations included the lack of variability in occupancy rates, which might hinder its application across different real-world scenarios. Similarly, in a more generalized study, Ghahramani et al. [22] proposed a knowledge-based system that integrates personalized comfort preferences with energy consumption patterns. By utilizing online learning and optimizing zone temperature setpoints, the system tailored comfort to individual occupants while minimizing energy usage. The authors incorporated thermal comfort parameters such as air temperature, mean radiant temperature, air velocity, and relative humidity. Personalized comfort preferences are modeled at the zone level, with temperature setpoints optimized for energy efficiency while maintaining comfort and indoor air quality. The main challenge lies in scalability across different building types and occupant preferences.

On the other hand, Yadav and Kowli [16] introduced a simulation framework to model HVAC behavior and energy consumption, addressing limitations in existing simulators. Their approach simulated temperature distribution and power savings in homes, outperforming traditional zone comfort optimization techniques. The framework incorporated thermal comfort metrics, including air temperature, mean radiant temperature, air velocity, and relative humidity, to optimize AC setpoint control based on local comfort rather than zonal comfort. While the results were promising, the framework's reliance on limited data sources restricted the generalizability of its findings. Lastly, Li et al. [23] explored a demand response (DR) strategy for air-conditioning systems, combining thermal comfort models with outdoor temperature predictions and time-of-use pricing to optimize indoor temperatures dynamically. Their model incorporated the Predicted Mean Vote (PMV) indicator, a widely used thermal comfort metric based on heat balance equations, to assess human thermal comfort. The PMV index accounted for key environmental parameters, including air temperature, mean radiant temperature, air velocity, and relative humidity, ensuring a comprehensive evaluation of indoor comfort levels. The strategy adjusted comfort thresholds during peak and off-peak electricity pricing periods to balance energy efficiency and occupant comfort. Their model achieved notable energy savings, though its applicability across varying building types and climates remains a limitation.

2.3. Prediction Models for Smart Buildings and District Heating

The integration of machine learning and predictive models for controlling smart building systems and district heating is gaining traction. These models aim to predict energy demand and optimize system operations for greater energy efficiency and comfort.

In this area, Cui et al. [24] proposed a heat load prediction model for district heating systems, combining Whale Optimization Algorithm (WOA) with LSTM, attention mech-

anism (ATT), and convolutional neural networks (CNN). This hybrid model addressed challenges in predicting heat load under fluctuating conditions, though its reliance on a single dataset and short observation periods limited its effectiveness. In a similar vein, Huang et al. [25] developed a two-stage data-driven framework to predict indoor temperatures and control heating systems in smart buildings. Their approach combined a univariate AR model for ambient conditions with a multivariate XGBoost model for temperature predictions. While it proved effective in managing energy use, its assumption of stationary ambient conditions reduced its adaptability in dynamic environments.

Meanwhile, Yang et al. [26] explored an AE-GWO-GRU-based method for predicting heat load in district heating systems, using autoencoder-based data augmentation and grey wolf optimization for parameter tuning. This approach addressed issues such as heat load fluctuations and transmission lags but was limited by its reliance on a small dataset, which was augmented for better performance. Lastly, Qaisar et al. [27] introduced OPTnet, a Transformer-based model for building occupancy prediction, which uses multi-sensor data, including HVAC operations. Although it significantly improved HVAC control and energy optimization, its dataset was limited to HVAC operational periods, reducing the scope for broader application.

Table 1. Comparison of existing techniques and methods.

Ref.	Year	Focus	Dataset	Technique	Result	Limitation
Ghahramani et al. [22]	2014	Personalized HVAC control for energy efficiency	The dataset collected from an occupant	Knowledge-based optimization	Average daily airflow reduction of 57.6 m ³ /h (12.08%)	Potential limitations in scalability and generalizability
Ghahramani et al. [18]	2018	Analyzing HVAC control policies	DOE references small office buildings in three U.S. climate zones	EnergyPlus simulations	Building-level annual control achieves 27.76% to 50.91% energy savings	Reliance on simulation data and simplified thermal comfort modeling
Aryal and Becerik-Gerber [28]	2018	Comfort-driven zone-level setpoints	RP884 database	Simulation and quantification	25% increase in occupant satisfaction, 2.1% energy savings	Difficulty in meeting ASHRAE requirements for occupant satisfaction
Huang et al. [25]	2020	Indoor temperature prediction and heating system control	Data from a smart building sensor network	AR, XGBoost	Predicted warm-up time: 19 min, Target temperature: 22.4 °C	Assumes stationary ambient conditions for predictive accuracy
Chen et al. [19]	2021	Data-driven thermal comfort optimization	ASHRAE Global Thermal Comfort Database II	Support Vector Machine (SVM), Random Forest	Total cost savings exceed CNY 1.5 million	Scalability to different building types or climates.
Yadav and Kowli [16]	2022	Simulation framework for HVAC optimization	Thermal behavior and energy consumption data collected from various sensors	Regression-based spatial thermal mapping	More energy savings compared to traditional zonal comfort optimization	Uncertain reliability of model performance

Table 1. Cont.

Ref.	Year	Focus	Dataset	Technique	Result	Limitation
Mahjoub et al. [17]	2022	Short-term power consumption forecasting	Power consumption data from Péronne city, France	LSTM, GRU, and Drop-GRU	Drop-GRU produced better accuracy and prediction speed, with low RMSE and MAE, and high R (near +1)	Learning time depends on the forecasting method used
Talami et al. [21]	2023	HVAC energy consumption optimization	U.S. DOE reference building energy models	HVAC zone temperature setpoint optimizer	Additional energy savings of 2–10% achieved	Static occupancy rates and specific problem formulation may limit generalizability to different scenarios
Almeida et al. [20]	2023	Adaptable strategy for HVAC setpoints	Smart thermostat data from European households	Predictive indoor temperature models	Energy savings up to 5.2%	Feasibility on resource-constrained smart thermostat devices
Li et al. [23]	2023	Demand response control strategy	Historical temperature and humidity data of the office building	Heat balance equations, RBFneuralnetwork	5.92% reduction in electricity consumption and 6.81% decrease in operational costs during DR periods	Applicability to different building types and climates may vary
Qaisar et al. [27]	2023	Building occupancy prediction	Operational data from HVAC system	Transformer network (OPTnet)	Superior performance in accuracy and mean squared error compared to other method	Dataset only includes HVAC operational hours
Cui et al. [24]	2024	Heat load prediction	1100 h dataset from a heat exchanger station	CNN, LSTM	$R^2 = 0.9962$, MSE = 0.0001, MAE = 0.0082	Limited to a single dataset and short observation period
Yang et al. [26]	2024	Heat load prediction in district heating systems	Data from DHS in a multifunctional region	Autoencoder, grey wolf optimization, Gated Recurrent Unit	RMSE: 47.90, MAPE: 2.17%	Insufficient data, requiring augmentation for improved model performance
Jing et al. [29]	2024	Energy management optimization in buildings	Real data from a building in Nanjing	Deep policy gradient decision making	Effective balance of energy efficiency and user comfort	Optimized only for a specific time frame

3. Dataset

The European Central Bank (ECB) Headquarters in Frankfurt is a state-of-the-art building equipped with multiple sensor networks as part of an advanced Building Management System, enabling granular analysis.

Data are collected from the district heating network that supplies the building, using measurements from devices located in the heat exchange units responsible for transferring heat from the district system to the building. Additionally, data from the building's weather station (BWS) offer a valuable foundation for analyzing the relationship between environmental parameters, area temperatures (AT), and heating demand. By analyzing trends from the BWS, we can identify key patterns and seasonal variations in space heating.

3.1. Heating Consumption Data

The district heating system operates as the baseload supplier, which increases the building foundation's thermal reliability and diversification. These data are collected by a smart meter which is located in the premises' technical areas. This device tracks energy consumption, with measurements taken at 15 min intervals. Although the heating consumption data are recorded at 15 min intervals, real-time data are also collected from Building Management Systems (BMSs) and sensors, ensuring inputs for more sustainable building management practices. The data monitor is installed in the exchange technical area between the provision heating network and the building's heating system. This data monitor is an integral part of the overall data acquisition process.

One of the parameters that can be used to forecast the energy demand needed to heat a building is the Heating Degree Day (HDD). The number of degrees below an established reference temperature proves a day's mean temperature necessary to use heating in buildings [30]. In this case, the 15 °C threshold has been selected based on expert knowledge, ensuring accuracy and relevance in measuring the building heating requirements. In this study, Heating Degree Hours (HDH) are employed because it disaggregates and yields a more accurate assessment of heating requirements, especially when working with hourly temperature data compared to HDD. The formula for calculating HDH is as follows:

$$\begin{cases} \frac{15 - T_{\text{out}}(h)}{24} & \text{if } T_{\text{out}}(h) < 15 \\ 0 & \text{if } T_{\text{out}}(h) \geq 15 \end{cases}$$

where T_{out} is the hourly outside temperature, and HDH represents the hourly degree hours.

3.2. Weather Variables from Building Weather Station (BWS)

The weather variables dataset consists of a broad range of meteorological parameters for evaluating atmospheric conditions, including temperature, humidity, wind speed, and many others. These factors indicate weather pattern dynamics that govern the operations of buildings and their energy use. Obtained directly from a weather station built on the top of the ECB skyscraper, this dataset provides accurate information that applies to the immediate surroundings of the building. Positioning the weather station at such an elevation maximizes gathering appropriate data that echo the building's atmosphere.

The integration with the Building Automation System (BAS) significantly improves data efficiency and accessibility. The seamless connection between the weather station and BAS ensures smooth data transmission and storage, making real-time weather monitoring readily available to building managers and operators. This integration enables informed decision making in control rooms, leveraging up-to-date weather information for better building management. Additionally, the BAS allows for the retrospective analysis of historical weather data, enabling the identification of trends and patterns.

The inclusion of ATs is a key feature of the dataset, as they play a crucial role in understanding building thermal dynamics. The temperatures extracted from the BAS provide insights into how heating demands vary across different zones of the building. Assessing the thermal behavior of each area requires careful consideration of these areas' role in the business activity. When combined with weather data and SHC, they form the foundation for accurately modeling the relationship between comfort, environmental conditions, and energy usage, enabling more precise energy management and optimization.

3.3. Hourly Temperature Trends in Building Zones

The analysis of hourly temperature trends across various building zones reveals distinct patterns between weekdays and weekends, highlighting the influence of occupancy and operational demands on heating systems as shown in Figure 1. Temperatures remain

higher almost daily during the week, particularly between the hours of working morning through evening. This implies that heating systems operate more on weekdays where they are more likely to be needed to maintain thermal comfort for building occupants.




Figure 1. Average hourly air temperature across different zones on weekdays and weekends.

There is a clear hourly variation in the sense that temperatures decrease slowly during the early morning hours (0–6 a.m.) and then increase during the day. Between 12 p.m. and 6 p.m., a gap between weekday and weekend temperatures can be seen, and this is the peak heating time. The lower occupancy and operational demand on weekends are confirmed further by reduced heating operations during non-working days.

Temperature trends in different building zones differ, which implies different heating control strategies. For example, building zones GC_0_5 and HN_0_13 are more distinct across the weekday and weekend, indicating that the heating adjustment takes place in different ways on different floors. In addition, more heating seems to be needed in zones characterized with higher sun exposure, such as the south face, and zones in shade show somewhat more heating. This variation underscores the importance of tailored heating strategies to optimize energy efficiency and occupant comfort in different building areas. Zones like GC_0_5 and HN_0_13, with lower sun exposure or higher occupancy during weekdays, show the need for more heating, while areas such as south-facing zones or those with higher sunlight exposure may require less heating.

3.4. Space Heating Consumption and Air Temperature

Figure 2 presents the monthly scaled SHC from 2018 to early 2024. It reveals a substantial seasonal variation in SHC, with the highest values occurring consistently during the colder months—Q1 (January, February, March: winter to early spring) and Q4 (October, November, December: fall to early winter).

3.5. Variables and Aggregation

The variables are categorized into three main groups: building weather station variables, time-related variables, and AT. Below is detailed information on each group and the aggregation methods applied to some variables.




Figure 2. Relationship between SHC and external air temperature in winter of each year—reflecting the increased demand for heating. In contrast, SHC significantly drops during the warmer months—Q2 (April, May, June: spring to early summer) and Q3 (July, August, September: summer to early fall)—suggesting reduced heating needs. The trend repeats yearly, with slightly higher peaks in some years, possibly indicating colder winters or increased heating requirements. The data show a stable seasonal consumption pattern, with a notable peak again in early 2024, suggesting similar heating demands.

3.5.1. Area Temperature Variables

Specific building zones are used to predict AT, and the following (Table 2) variables are aggregated from different building-specific areas. Thus, the aggregated variables are averages of respective zone temperatures to obtain a clear picture of the fluctuations in the corresponding areas of the building.

Table 2. Area Temperature Variables and Aggregation.

Aggregated Variable	Component Variables	Aggregation Description
CA_0_5	'CA0_E', 'CA0_W', 'CA1', 'CA2_E', 'CA2_W', 'CA3', 'CA4_A1', 'CA4_A2', 'CA5'	Common Areas on floors 0–5
NT_0_13	'NT0_1_NO', 'NT0_1_NW', 'NT10_13_NO', 'NT10_13_NW', 'NT2_5_NO', 'NT2_5_NW', 'NT6_9_NO', 'NT6_9_NW'	North Tower on floors 0–13
ST_0_13	'ST0_1_SO', 'ST0_1_SW', 'ST10_13_SO', 'ST10_13_SW', 'ST2_5_SO', 'ST2_5_SW', 'ST6_9_SO', 'ST6_9_SW'	South Tower on floors 0–13
NT_14_25	'NT22_25_NO', 'NT22_25_NW', 'NT18_21_NO', 'NT18_21_NW', 'NT14_17_NO', 'NT14_17_NW'	North Tower on floors 14–25
ST_14_25	'ST14_17_SO', 'ST14_17_SW', 'ST18_21_SO', 'ST18_21_SW', 'ST22_25_SO', 'ST22_25_SW'	South Tower on floors 14–25
NT_26_37	'NT26_29_NO', 'NT26_29_NW', 'NT30_33_NO', 'NT30_33_NW', 'NT34_37_NO', 'NT34_37_NW'	North Tower on floors 26–37
ST_26_37	'ST26_29_SO', 'ST26_29_SW', 'ST30_33_SO', 'ST30_33_SW', 'ST34_37_SO', 'ST34_37_SW'	South Tower on floors 26–37

3.5.2. Building Weather Station Variables

The building weather station (BWS) variables provide crucial environmental data influencing AT predictions and SHC estimations. These variables include the following:

- BWS_air_pressure: Atmospheric pressure at the building location.
- BWS_air_temperature: External air temperature recorded by the BWS.
- BWS_wind_speed: Speed of the wind measured by the BWS.
- BWS_wind_direction: Direction from which the wind is blowing.
- BWS_relative_humidity: Humidity level in the atmosphere.
- BWS_global_radiation: Amount of solar radiation received.

3.5.3. Space Heating Consumption Variables

In the case of SHC prediction, the target variable used is the total SHC of the building. This variable is computed with the help of the predicted AT, which is one of the main inputs in the model. Other related features are also incorporated into the analysis to improve the accuracy of the SHC estimates, apart from AT forecasts. Such features include environmental conditions, for instance, wind speed, air pressure, and time features, such as the hour of the day and the season.

3.5.4. Time-Related Variables

Time-related variables are integral to modeling, accounting for seasonal variations and daily patterns affecting temperature and heating needs. The following (Table 3) variables are considered:

Table 3. Time variables and their description.

Time Variable	Description
year	Year of the observation
month	Month of the observation
day	Day of the month
hour	Hour of the day
day_of_week	Day of the week (0 = Monday, 6 = Sunday)
season_fall	Indicator for fall season
season_spring	Indicator for spring season
season_summer	Indicator for summer season
season_winter	Indicator for winter season
HDH15	Heating Degree Hours (base 15 °C)
weekend	Indicator for weekend days (1 if weekend, 0 otherwise)

The variables listed above will be utilized across the three research goals. The AT variables will support Goal 1, while the SHC variable will be used explicitly for Goal 2. The time-related variables will provide essential context for temperature predictions and SHC modeling. Together, these variables enable a comprehensive analysis to optimize energy consumption and maintain occupant comfort in building environments.

4. Proposed Methodology

The methodology of this study involves a two-stage machine learning framework for energy consumption prediction, starting with preprocessing steps that extract date and time components to capture temporal patterns, such as seasonal trends and daily

variations, in energy usage. Missing values are handled primarily through imputation, with linear regression used as a backup method when necessary. Feature engineering includes the introduction of key features for heating demand to improve model accuracy. The modeling phase develops a robust temperature prediction model using LSTM, GRU, and Transformer architectures, specifically tailored to individual building zones, ensuring accurate thermal management while minimizing energy consumption. Although LSTM, GRU, and Transformer models are widely used, they are particularly well suited for time series forecasting, which is critical for this study's goal of predicting temperature and heating consumption patterns. Performance is evaluated using metrics such as R-squared, Mean Squared Error (MSE), and Mean Absolute Error (MAE), with thresholds set based on model-specific requirements. The novelty of this work lies not only in the application of these architectures to space heating optimization but also in the integration of custom loss functions to balance energy efficiency and occupant comfort. The optimization of space heating consumption and comfort is achieved through two models: one predicting area temperatures, with smoothness enforced through a penalty, and another predicting heating demand, which leverages temperature predictions to estimate global SHC and optimize energy use across multiple zones. The two-stage predictive modeling approach for both temperature forecasting and heating consumption prediction adds further value, distinguishing this study from traditional applications of these models in the energy domain. Custom loss functions are introduced to regulate temperature variations while minimizing energy consumption, enforcing thermal comfort constraints and promoting energy efficiency in sustainable building management.

4.1. Preprocessing

The data preprocessing phase is critical in preparing the dataset for effective modeling and analysis. This process incorporates different data cleaning, transformation, and preparation methods to make the raw data in the correct format for modeling. These include conversion of the extracted date and time into their paramount components, recognition of seasons in the data stream, mechanism handling of missing data, and feature engineering of quantitative and categorical characteristics that synthetically dissect the applied dataset. In addition, the data are classified into more generalized forms and normalized to the same scale, and the data are organized into sequences for time series. Last but not least, the dataset is split into train, validation, and test sets to make appropriate model assessments possible.

4.1.1. Extracting Date and Time Components

The preprocessing begins by extracting essential date and time components from the 'Datetime' column. The code identifies and separates the following elements: year, month, day, hour, and day of the week. This breakdown transforms the timestamp into individual components that can effectively be utilized as features in the predictive models.

4.1.2. Determining Seasons

A function is defined to enrich further the dataset, which assigns a specific season to each data point based on the month. The seasons are categorized as follows: Winter encompasses December, January, and February; Spring includes March, April, and May; Summer covers June, July, and August; and Fall consists of September, October, and November. One-hot encoding is applied to create separate binary columns for each season, such as 'season_winter' and 'season_spring', enabling the models to account for seasonal variations in energy consumption.

4.1.3. Handling Missing Values

The preprocessing process also addresses the issue of missing values within the dataset. The code identifies variables that may have missing entries, including 'SHC', 'BWS_air_pressure', and others. For each variable with missing data, the code fills these gaps using the mean values of groups defined by 'year', 'month', 'hour', and 'BWS_air_temperature'. Suppose any missing values persist after this initial imputation. In that case, a linear regression model is employed to predict and fill them based on the 'BWS_air_temperature' variable, ensuring the dataset is as complete as possible for analysis.

4.1.4. Feature Engineering

One of the most straightforward and compulsory techniques to enhance the effectiveness of the predictive analysis of the model involves feature engineering. To describe the heating demand when the outside temperature is less than or equal to 15 °C, a new variable 'HDH15' is introduced, and more information about the models is given. The weekends binary column added indicates if the data point is on the weekend or not. This feature is valuable for recording variations in heating requirements within the working and weekend days.

4.1.5. Grouping Building Areas

To manage the complexity of the dataset, individual building ATs are grouped into seven distinct categories by averaging related columns. For instance, Group 1, referred to as 'CA_0_5', combines temperatures from various areas such as 'CA0_E', 'CA1', and others. Similar aggregations are applied to form Groups 2 through 7. Following this grouping, the original columns used to create the groups are discarded. This step allows for a focus on broader temperature patterns rather than individual ATs, facilitating more effective modeling.

4.1.6. Data Scaling

When applied, feature scaling prepares the dataset for modeling. The feature columns, the weather data, and time components go through the MinMaxScaler, which scales all features in the dataset to a range of 0 and 1. Likewise, the target variables, such as the grouped temperatures, are scaled for consistency across the dataset. In addition, the scalers used for both features and targets are equally saved with joblib, which can also be used to scale during the next prediction.

4.1.7. Data Splitting

Finally, the dataset, consisting of 61,320 rows (representing 24 h × 365 days × 7 years), is divided into three sets for model training and evaluation. The training set, constituting 60% of the data, is utilized to train the models. The validation set, comprising 20%, is employed to fine-tune model parameters and mitigate the risk of overfitting, using a 5-fold cross-validation approach. The remaining 20% forms the test set, which assesses the final models' performance.

4.2. Area Temperature and Space Heating Consumption Prediction in Building Areas

This methodology has two primary objectives. The first objective is to predict the temperatures of specific categories of building areas. The second objective is to predict the SHC using LSTM, GRU, and Transformer algorithms. Based on the history of the building temperature, data from the BWS, and parameters that define time, the approach identifies the best models for temperature and SHC estimation, as depicted in Figure 3. To ensure that the models have been used at their best, grid search cross-validation has been used for the selection of hyperparameters in order to yield the best performance out of them.




Figure 3. Proposed methodology for forecasting ATs and SHC.

4.2.1. LSTM Model

LSTM is a type of recurrent neural network (RNN) that is especially suited for sequence data and overcoming the deficiencies of the traditional RNN to preserve long-term dependencies [31]. In contrast, other RNNs using simple feedback mechanisms are prone to get stuck in vanishing gradient problems; LSTMs have a specific memory cell design that allows the network to remember or forget information over relatively long intervals. Each LSTM cell includes three key gates—input, forget, and output gates—that allow control over which data should be allowed to enter or leave the cell state [32]. This gating mechanism allows LSTMs to retain essential previous information and discard irrelevant details, which makes the application of LSTMs ideal for applications that incorporate long-term contexts, such as language comprehension, time series forecasting, intricate system simulation, and complex system modeling.

In this study, the LSTM model is developed using the Keras Sequential API that enables the creation of an architecture suitable to efficiently capture temporal structures, especially in time series data [33]. In its basic form, the model utilizes an LSTM layer activation and functions to develop the hierarchical interconnections of the inputs. To address the risk of overfitting, a dropout layer is incorporated, which randomly deactivates a portion of the neurons during training. Additionally, a dense output layer is included to align with the number of target variables, allowing the model to predict multiple temperature groups. The model is compiled with the Adam optimizer and utilizes the mean squared error loss function to assess its performance. Training is conducted on the designated training set, followed by validation on a separate validation set to ensure robustness. After training, the model's predictive capability is evaluated on the test set, and predictions are generated. R-squared scores are then calculated for each target variable, providing valuable insight into the model's accuracy and effectiveness in temperature prediction.

4.2.2. GRU Model

GRU network is a more compact LSTM network designed for sequence modeling to be less computationally expansive while still efficient [10]. Like LSTMs, GRUs were developed to overcome the problem of vanishing gradients inherent in traditional RNNs and restrict the sequence information output [34]. GRUs achieve this through two primary gates: the update gate and the reset gate. The update gate controls how much of the previous time steps should be passed on to the next layer and, hence, how much information should be kept in the current step, while the reset gate controls how much of the previous hidden state should be forgotten in the current step. Its structure permits GRUs to learn long-term dependencies of a sequential input without burdening the algorithm with as many parameters as LSTMs, making for much quicker training. The GRU model's training, evaluation, and prediction are similar to the LSTM model in that they provide a proper comparison between the two architectures.

4.2.3. Transformer Model

Transformers are a deep learning architecture designed to process a sequence better than the RNN since it replaces recurrence with self-attention mechanisms [35]. Unlike RNNs, Transformers generalize information simultaneously and thus are very effective. They can capture dependencies from the quadratic time complexity of sequences, making them very effective in processing long sequences. The most crucial concept of Transformer architecture is self-attention, which allows the model to decide on the importance of a particular input in the output [35]. This mechanism is organized in several "heads" across attention layers and works with various relations inside the data, leading to more complex representations.

This paper develops a custom Transformer model to effectively capture long-term dependencies within the data, enhancing the model's ability to understand complex relationships over extended periods [36]. The model's architecture consists of multiple Transformer encoder blocks; the attention mechanism is present to concentrate on the most critical parts of the input sequence. Furthermore, employing layer normalization and multi-head attention helps stabilize the model's training and smooth the extraction of complex relationships within the data. After the encoder blocks, the feed-forward network, global average pooling layer, and dense layers are used to generate the output. When it comes to the Transformer model, the training and testing protocols are the same as those used for LSTM and GRU, thus presenting a clear comparison between them all.

4.2.4. Evaluation and Model Saving

R-squared metrics, along with mean squared error and mean absolute error, are utilized to assess the predictive accuracy of temperatures for specific groups of building areas and SHC. Following this evaluation, the trained models, labeled as `lstm_model1`, `gru_model1`, and `transformer_model1`, are saved using the `model save()` method for Keras models. This functionality ensures future access to the models without retraining, facilitating subsequent analyses and implementations.

4.3. Optimizing Space Heating Consumption While Maintaining Occupant Comfort

The third objective of this study is to reduce the building's SHC while maintaining occupants' comfort levels. This is achieved through a two-stage machine learning process: the first is a model to generate temperature estimates regarding the different building zones, and the second is to use those temperature predictions to estimate global SHC. The methodology uses GRU, LSTM, Transformer blocks, and custom loss functions to match desired energy efficiency and comfort, as shown in Figure 4.




Figure 4. Methodology of optimizing SHC while maintaining occupant comfort.

4.3.1. Model 1: Predicting Area Temperatures

In the first stage, temperature is forecasted in the specific areas of the building using historical data from the BWS and time-related variables. The model employs deep learning architecture such as GRU, LSTM, or Transformer blocks to capture temporal dependencies.

A custom loss function is implemented to regulate temperature predictions in line with occupant comfort. During work hours (7 a.m. to 6 p.m.), the model applies a high penalty for predicted temperatures falling below the recommended range of 21–23 °C, as outlined in EN 16798-1 for winter thermal comfort [7]. Deviations from this threshold incur significant penalties to ensure the temperature remains within the comfort zone. In contrast, during non-work hours, the model allows the temperature to drop to 18 °C, promoting energy savings. Although penalties for deviations during this period are less than work hours, they do achieve the comfort and energy efficiency balance. When the building is typically unoccupied on weekends, the model imposes penalties for temperatures exceeding 18 °C during those times to capitalize on energy conservation opportunities.

Additional constraints in the loss function prevent excessive heating by applying strong penalties for temperatures exceeding 23 °C, in accordance with the upper limit of the recommended winter comfort range in EN 16798-1, which results in unnecessary energy consumption. The model also discourages rapid temperature changes—predicted

temperature shifts exceeding 0.5 °C between consecutive hours incur penalties, promoting smoother, more efficient heating.

- Input and Output Handling:

The model inputs sequential environmental and time-related features, including historical temperature, wind speed, and time of day, normalized between 0 and 1 with specific cutoffs for work and non-work hours. The output consists of the predicted ATs, passed to the second model for SHC prediction. The training process for Model 1 utilizes the custom loss function, `custom_loss_model1`, as previously described. To prevent overfitting, validation is conducted using the `val_step_model1` function. Additionally, smoothness constraints within the loss function ensure gradual transitions in temperature predictions, promoting comfort and energy efficiency.

4.3.2. Model 2: Predicting Space Heating Consumption

Once Model 1 makes the AT predictions, they serve as additional inputs for Model 2, which forecasts the overall SHC required to maintain the predicted temperatures. Model 2 uses original input features (e.g., outdoor temperature, pressure, and wind speed) and the temperature predictions from Model 1. This is achieved by concatenating the outputs of Model 1 with other input features.

A custom loss function is implemented for Model 2 to minimize energy waste. This loss function imposes a heavier penalty for overestimating heating needs, as such overestimations result in unnecessary energy consumption. The model accurately predicts the exact amount of heating needed to maintain occupant comfort without overshooting by placing greater emphasis on penalizing over-predictions than under-predictions.

The input is a concatenation of predicted temperatures from Model 1 and various environmental variables. Combining these datasets constitutes a comprehensive basis for the model to process and analyze. Model 2's output is the predicted values of SHC. Real-world energy consumption data are then used to compare with these predicted values in order to effectively evaluate the model and its performance in predicting actual consumption data.

4.3.3. Custom Loss Function Design

The success of both models hinges on their tailored loss functions, which strategically balance occupant comfort and energy efficiency. For Model 1 (AT Prediction), the loss function emphasizes maintaining comfort thresholds during work hours (21 °C–22 °C) while allowing for energy-saving measures during non-work hours and weekends by relaxing temperature constraints. Additional penalties are applied to avoid excessive heating (>22 °C) and to ensure smooth transitions in temperature predictions, promoting both comfort and operational efficiency. The custom loss function operates on multiple layers, ensuring that temperature predictions remain close to the lower limit to minimize energy consumption. At the same time, it enhances consistency by reducing volatility, leading to a more stable temperature profile that improves overall comfort. This two-phase design balances energy efficiency with stability, preventing abrupt fluctuations that could disrupt occupant comfort.

For Model 2 (SHC Prediction), the loss function prioritizes accuracy in heating predictions, with a heavier penalty on over-predictions to prevent unnecessary energy consumption. Under-predictions are penalized less severely to allow for efficient heating within acceptable comfort ranges. By aligning Model 1's stable temperature predictions with Model 2's optimized SHC forecasts, the custom loss functions ensure that energy efficiency is achieved without compromising thermal comfort. Together, these custom loss functions align the predictive goals of both models—ensuring the temperature predictions from

Model 1 lead to optimized SHC predictions in Model 2, achieving a cohesive strategy for energy-efficient building management.

4.3.4. Mathematical Formulation of Loss Function

The custom loss functions for both models are designed to balance occupant comfort and energy efficiency.

For Model 1, the loss function $\mathcal{L}_1$ combines the temperature prediction error with penalties for excessive heating and smoothness. It is designed as follows:

$$\mathcal{L}_1 = \alpha \sum_t (|T_{\text{pred}}(t) - T_{\text{actual}}(t)| + \beta \sum_t (\max(0, T_{\text{pred}}(t) - 22)) + \gamma \sum_t ((T_{\text{pred}}(t) - T_{\text{pred}}(t-1))^2) \quad (1)$$

where $T_{\text{pred}}(t)$ is the predicted temperature at time t , $T_{\text{actual}}(t)$ is the actual measured temperature at time t , α, β, γ are hyperparameters that control the weights of each term. The first term minimizes the temperature prediction error. The second term penalizes temperatures above 22 °C, which is the upper comfort limit. The third term penalizes the volatility in temperature changes, ensuring smooth transitions and reducing abrupt fluctuations.

For Model 2, the loss function $\mathcal{L}_2$ penalizes over-predictions more severely than under-predictions to prioritize energy savings. The equation is defined as follows:

$$\mathcal{L}_2 = \delta \sum_t (\max(0, \hat{Q}_{\text{pred}}(t) - \hat{Q}_{\text{actual}}(t))^2) + \epsilon \sum_t (\max(0, \hat{Q}_{\text{actual}}(t) - \hat{Q}_{\text{pred}}(t))^2) \quad (2)$$

where $\hat{Q}_{\text{pred}}(t)$ is the predicted space heating consumption at time t , $\hat{Q}_{\text{actual}}(t)$ is the actual measured heating consumption at time t , δ and ϵ are hyperparameters, where $\delta > \epsilon$ to penalize over-predictions more heavily than under-predictions.

The final combined loss function $\mathcal{L}_{\text{final}}$ integrates the objectives of both models (temperature and heating consumption), ensuring that the temperature predictions from Model 1 guide the heating consumption predictions from Model 2:

$$\mathcal{L}_{\text{final}} = \mathcal{L}_1 + \lambda \mathcal{L}_2 \quad (3)$$

where λ is a scaling factor that balances the contributions from both loss functions, ensuring that both temperature prediction and heating consumption optimization are aligned.

4.3.5. Model Evaluation

The models were evaluated through a comprehensive analysis of their performance metrics on unseen data. Each model was assessed based on specific criteria, such as accuracy in predicting temperature deviations from target values during defined periods, adherence to established upper-temperature limits, and the smoothness of temperature transitions over consecutive hours. Additionally, average deviations from target temperatures during working, non-working, and weekend hours were calculated to evaluate how well each model maintained comfort levels. The effectiveness of the models was further examined by comparing their predicted values against actual observations, focusing on the overall reduction in percentage difference between actual and predicted outcomes.

5. Results and Discussion

This section presents a comprehensive analysis of the results obtained from the implemented machine learning models for temperature and SHC predictions. The results demonstrate these predictive capabilities' contribution to energy usage optimization in the building while maintaining comfort standards across dynamic operational conditions.

5.1. Area Temperature and SHC Forecasting

The results of the temperature prediction model for various heating zones demonstrate notable differences in performance, as measured by the R-squared values as shown in Table 4. Across the heating zones, the Transformer model consistently achieves the highest R-squared values, particularly in the CA_0_5 zone, where it reaches 0.945718, indicating its superior capability in capturing the underlying patterns in the temperature data. The LSTM and GRU models also perform well, with R-squared values close to the Transformer, especially in the NT_0_13 and NT_14_25 zones, where the LSTM model achieves R-squared values of 0.935583 and 0.948459, respectively. This suggests that while LSTM and GRU models are compelling, the Transformer model may leverage its attention mechanisms to better understand the temporal dependencies and interactions in the data.

Table 4. Models evaluation in area temperature forecasting.

Heating Zone	LSTM R-Squared	GRU R-Squared	Transformer R-Squared
CA_0_5	0.92	0.92	0.95
NT_0_13	0.94	0.94	0.94
ST_0_13	0.92	0.92	0.91
NT_14_25	0.95	0.95	0.95
ST_14_25	0.91	0.91	0.91
NT_26_37	0.95	0.95	0.95
ST_26_37	0.92	0.91	0.91

In the ST_0_13 zone, all three models show relatively similar performance, with R-squared values hovering around 0.92, indicating consistent temperature predictions regardless of the model used. However, the least favorable performance is observed in the ST_26_37 zone, where both the LSTM and GRU models yield slightly better R-squared values compared to the Transformer model. This suggests that the complexities of this specific heating zone's data may not align as well with the Transformer's architecture.

The results of the SHC prediction models demonstrate strong performance across the LSTM, GRU, and Transformer architectures, with R-squared values indicating high accuracy. The GRU achieves the highest R-squared value of 0.9522, closely followed by the LSTM at 0.9519, highlighting their effectiveness in capturing complex relationships within the dataset. The GRU's slight advantage may stem from its efficient management of temporal dependencies through its gating mechanism. In contrast, the Transformer model shows a respectable R-squared value of 0.9389 but falls short of the other two models, suggesting that its capabilities may not be fully utilized for this specific dataset, as shown in Figure 5. Overall, these findings emphasize the suitability of the GRU and LSTM models for SHC prediction, which will significantly enhance energy management strategies and optimize energy usage in buildings.

The Transformer model shows the lowest R-squared value, suggesting it may not be the optimal choice for energy consumption prediction despite its ability to handle sequential data. This could be due to its requirement for large datasets to capture long-range dependencies, which was not fully possible with our smaller dataset.




Figure 5. Model evaluation in SHC forecasting.

5.2. Optimizing SHC for Occupant Comfort Using Predictive Modeling

5.2.1. Experimental Setup

The models were evaluated with comfort metrics tracked on unseen data, consisting of February and March 2024 temperatures. Table 5 shows the comfort metrics used in this paper's evaluation.

Table 5. Comfort metrics with their description.

Comfort Metric	Description
Upper Limit Violation Percentage	The percentage of time temperatures in each area exceeded 22 °C, a critical threshold for energy efficiency.
Deviation from 21 °C during Work Hours	The average deviation from the ideal temperature of 21 °C during work hours (6 a.m. to 6 p.m.) is crucial for maintaining occupant comfort.
Deviation from 18 °C during Non-Work Hours	Average deviation from 18 °C during non-working hours (7 p.m. to 5 a.m.), representing energy-saving opportunities.
Deviation from 18 °C during Weekends	Deviation from 18 °C during weekends when the building is unoccupied is needed to assess energy-saving potential.
Smoothness Violation Percentage	The percentage of ATs that changed by more than 0.5 °C from one hour to the next is essential for avoiding discomfort and unnecessary energy consumption.

The analysis compared the models' performances using the above metrics to the actual building data without a model, allowing for a detailed evaluation of the comfort profiles achieved by each predictive model. Table 6 presents the comfort metrics for the unseen data when no predictive model was applied.

Table 6. Comfort metric of unseen data.

Area	Upper Limit Violation %	Deviation from 21 °C During Work Hours	Deviation from 18 °C During Non-Work Hours	Deviation from 18 °C During Weekends	Smoothness Violation %
CA_0_5	0.00	0.48	3.33	3.26	0.00
NT_0_13	3.76	0.43	2.99	2.83	0.00
ST_0_13	23.12	0.75	3.41	3.31	0.74
NT_14_25	7.39	0.56	2.95	2.77	0.07
ST_14_25	27.22	0.82	3.50	3.38	0.94
NT_26_37	8.94	0.62	2.90	2.73	0.40
ST_26_37	26.81	0.95	3.37	3.26	3.16
Average	13.89	0.66	3.21	3.08	0.76

5.2.2. Performance Evaluation

The LSTM model's results show a significant reduction compared to the actual values of SHC, with a 20.95% reduction. The Upper Limit Violation was successfully avoided over all areas by analyzing the comfort metrics in Table 7, where thresholds were not exceeded above 22 °C, a critical point for energy efficiency. The model was able to keep the temperatures close to the 21 °C ideal for occupant comfort in 0.347 °C average deviation from 21 °C during 8:00 to 18:00 (work hours).

Table 7. Performance evaluation of LSTM SHC optimization.

Area Zones	Upper Limit Violation %	Deviation from 21 °C During Work Hours	Deviation from 18 °C During Non-Work Hours	Deviation from 18 °C During Weekends	Smoothness Violation %
CA_0_5	0.00	0.14	2.07	2.00	0.76
NT_0_13	0.00	0.53	1.02	0.91	6.25
ST_0_13	0.00	0.30	1.29	1.16	9.38
NT_14_25	0.00	0.46	0.19	0.02	14.31
ST_14_25	0.00	0.39	1.22	1.07	11.81
NT_26_37	0.00	0.30	0.35	0.17	14.72
ST_26_37	0.00	0.32	0.91	0.75	12.71
Average	0.00	0.35	1.01	0.87	9.99

The deviation from 18 °C averaged 1.006 °C during non-working hours, showing that the model maintained a reasonable energy-saving vs. comfort balance during periods of lower occupancy. The deviation from 18 °C over weekends was also lower by 0.869 °C, indicating that the model could optimize energy use when the building was empty.

The Smoothness Violation showed that around 10% of temperatures changed over 0.5 °C within one hour. While a slight loss of comfort sensation could occur, the model performs reasonably well overall in keeping temperature transitions stable, given the significant reduction in energy consumption.

The GRU model moderately improved SHC, achieving a 2.33% decrease from actual SHC. In comfort metrics, the Upper Limit Violation stayed consistently at 0%, meaning the predicted temperatures never exceeded the 22 °C threshold in any zone and maintained energy efficiency, as shown in Table 8.

Table 8. Performance evaluation of GRU in SHC optimization.

Area Zones	Upper Limit Violation %	Deviation from 21 °C During Work Hours	Deviation from 18 °C During Non-Work Hours	Deviation from 18 °C During Weekends	Smoothness Violation %
CA_0_5	0.00	0.98	2.00	2.00	0.07
NT_0_13	0.00	2.05	0.91	0.91	0.69
ST_0_13	0.00	1.79	1.16	1.16	0.83
NT_14_25	0.00	2.91	0.02	0.02	1.67
ST_14_25	0.00	1.86	1.07	1.07	1.46
NT_26_37	0.00	2.71	0.17	0.17	2.29
ST_26_37	0.00	2.12	0.75	0.75	2.36
Average	0.00	2.06	0.87	0.87	1.34

However, there were differences from the ideal temperatures (particularly during work hours), with an average deviation of 2.06 °C from the ideal 21 °C. The GRU model's performance is higher than the LSTM model's, suggesting it may have had difficulty maintaining optimal comfort temperature during work hours.

The average deviation of 18 °C from non-work hours was 0.87 °C, with variation but still within an acceptable range for saving energy. The deviation from 18 °C during weekends was about 0.87 °C, which indicated that the model worked well with unoccupied periods.

The Smoothness Violation metric had a much better result (1.34%) than the LSTM model. This implies that the resulting temperature shifts in the GRU model exceeded 0.5 °C in fewer occasions of cases, ensuring comfort sensation and decreasing energy usage.

The Transformer model exhibited high energy efficiency, with a saving percentage (20.69%) similar to the LSTM model. The Transformer model's performance in the Upper Limit Violation metric showed that 0% of zones violated their temperature, which helped accomplish the energy-saving goals.

As we can verify in Table 9, the average deviation from the ideal 21 °C of 0.40 °C did represent a significant improvement over the GRU and all LSTM models during work hours. This implies that the Transformer model could have a more consistent and comfortable indoor environment during working hours and manage occupant comfort.

Table 9. Performance evaluation of Transformer in SHC optimization.

Area Zones	Upper Limit Violation %	Deviation from 21 °C During Work Hours	Deviation from 18 °C During Non-Work Hours	Deviation from 18 °C During Weekends	Smoothness Violation %
CA_0_5	0.00	0.19	2.07	2.00	5.76
NT_0_13	0.00	0.44	1.02	0.91	9.44
ST_0_13	0.00	0.39	1.29	1.16	10.14
NT_14_25	0.00	0.38	0.19	0.02	11.53
ST_14_25	0.00	0.56	1.23	1.08	11.39
NT_26_37	0.00	0.34	0.35	0.17	11.88
ST_26_37	0.00	0.47	0.91	0.75	10.97
Average	0.00	0.40	1.01	0.87	10.16

The average deviation from 18 °C for non-work hours was 1.01 °C, slightly higher than other models but still within an acceptable range for energy efficiency. Likewise, the deviation was 0.87 °C, which is comparable to the non-work hour performance, and reasonable temperature control is achieved during non-used hours.

The percentage of Smoothness Violations was considerable, 10.16%, which meant that the Transformer model had more abrupt temperature changes than the rest of the models.

It is clear from the above experimental analysis results that the LSTM and Transformer models have similar results on both energy savings and comfort metrics. Both perform better than the GRU in terms of SHC while preserving occupant comfort. These results show the effectiveness of energy savings balanced with comfort using the LSTM and Transformer models with a 20.95% reduction in heating consumption. Additionally, it recorded the lowest average deviations from target temperatures over work and non-work hours and the fewest Upper Limit Violations and Smoothness Violations so that the temperature changes were steady and comfortable. In comparison, the Transformer model achieved strong performance with a 20.69% reduction in consumption while transitioning between the temperatures. However, Smoothness Violations are more indicative of less stable temperature transitions. While the GRU model achieved the lowest energy savings (2.33%), comfort metrics were consistently performed.

Although the proposed models are trained on historical data, their design supports extension to real-time applications. The modular two-stage architecture allows for incremental updates using new data streams, enabling periodic retraining or online learning strategies. While the dataset used in this study is limited in size, it offers fine-grained, high-frequency measurements from a real-world multi-zone building, capturing diverse spatial and temporal heating dynamics.

5.3. Assessment of Indoor Thermal Comfort

Thermal comfort is a critical factor in evaluating indoor environmental quality, particularly in buildings that implement demand response (DR) strategies for heating and cooling optimization. To ensure a comprehensive long-term thermal comfort analysis, we incorporate the three primary environmental categories outlined in EN 16798-1:2019 [7]. These categories classify indoor environments based on their suitability for different occupant groups, defining operative temperature ranges for winter and summer conditions. Given that our analysis covers the first three months of the year (January, February, and March)—corresponding to winter conditions—the target comfort range for our study is 21–23 °C, as specified for Category I (high comfort).

To assess the indoor environmental quality of the studied building, we computed the total area average temperature, along with specific averages for working days (Monday to Friday including the 24 h of the day) and working hours (Monday to Friday, 7 a.m.–6 p.m.). These calculations help determine whether the observed temperatures align with the EN 16798-1 winter comfort range. The results are summarized in Table 10.

Table 10. Computed average temperatures and comparison with EN 16798-1 winter comfort range.

Category	Computed Average Temperature (°C)	Target Comfort Range (Winter, °C)	Compliance with EN 16798-1 (Category I: 21–23 °C)
Total Area Average	21.02	21–23	Within Range
Total Area Average Working Days	21.15	21–23	Within Range
Total Area Average Working Hours	21.31	21–23	Within Range

6. Limitation

This study contributes to the field by demonstrating the effectiveness of deep learning models in predicting heating demand using readily available building and environmental data, offering a scalable approach for energy optimization in institutional buildings.

However, this study has certain limitations. Direct occupancy measurements were not available due to the absence of sensors; instead, temporal variables such as hour of day, day of week, and season were used as proxies, given their typical correlation with occupancy patterns in institutional settings. While these features helped capture occupancy-driven heating variations, the lack of real-time occupancy data may limit the models' adaptability to dynamic usage.

Additionally, this study does not specifically address the stricter thermal comfort and air quality standards required for IEQ Category I spaces, which are designed for sensitive groups such as children or the elderly. The assumption that indoor air, radiant, and operative temperatures are equal ($t_a = t_r = t_o$) was made for simplicity, particularly suited to winter conditions where differences are minimal. However, this may not apply in other seasons.

7. Conclusions

This study uses LSTM, GRU, and Transformer deep learning models to predict the ATs and SHC while maintaining occupant comfort. Models reduced SHC by up to 20.95% through a two-stage process based on LSTM and 20.69% through a two-stage process based on Transformer. This showcases that these models could optimize the energy consumption in complex building environments from increased SHC. The historical temperature data and time-specific parameters have been used to forecast precisely and act as good energy managers in different space heating zones. These results underscore the importance of leveraging advanced machine-learning techniques to improve energy efficiency and sustainability in modern buildings.

Future research can explore integrating advanced deep learning models like reinforcement learning to enhance prediction accuracy and real-time optimization. Furthermore, the framework can be extended further to introduce renewable energy sources; real-time feedback of occupancy patterns can be expected to add further enhancements to real-time energy management. The use of hybrid models combining multiple machine-learning techniques may also be explored to improve model reliability and interpretability. In addition, multi-agent reinforcement learning (MARL) can be considered for dynamic management of multiple HVAC zones. Finally, the methodology can be applied to larger-scale, multi-building environments to gain further insight into their robustness and scalability for a wide variety of operational conditions.

Author Contributions: Conceptualization, F.A., M.C., L.M. and N.C.-R.; methodology, F.A.; software, F.A.; validation, M.C., L.M. and N.C.-R.; formal analysis, F.A.; investigation, F.A.; resources, F.A.; data curation, F.A.; writing—original draft preparation, F.A.; writing—review and editing, M.C. and N.C.-R.; visualization, F.A.; supervision, M.C., L.M. and N.C.-R.; project administration, F.A., M.C. and N.C.-R.; funding acquisition, M.C. and N.C.-R. All authors have read and agreed to the published version of the manuscript.

Funding: This work was funded by the European Union under NextGenerationEU. Ecosystem INEST Interconnected Nord-Est Innovation; PNRR Mission 4, Component 2, Investment 1.5, Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or The European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This work was also supported by national funds through FCT (Fundação para a Ciência e Tecnologia), under the

project- UIDB/04152/2020 (doi:10.54499/UIDB/04152/2020)-Centro de Investigação em Gestão de Informação (MagIC)/NOVA IMS.

Data Availability Statement: The datasets used and/or analyzed during the current study are available from the corresponding author on reasonable request.

Conflicts of Interest: The authors declare no conflict of interest.

Nomenclature

Term	Description
AT	Area Temperature
BWS	Building Weather Station
CA	Common Area
LSTM	Long Short-Term Memory model
GRU	Gated Recurrent Unit model
MSE	Mean Squared Error
MAE	Mean Absolute Error
NT	North Tower
R-squared (R^2)	Coefficient of determination, a statistical measure of the goodness of fit
SHC	Space Heating Consumption
ST	South Tower

References

1. Ezzati, P.; Mohammadi, A.; Reza, M.; Fard, A. Effect of insulation on energy saving of buildings with closed cell flexible elastomeric foam (FEF) insulation. In Proceedings of the 5th International Conference on Management Optimization and Development of Energy Infrastructures. 2022.
2. Yao, R.; Steemers, K. A method of formulating energy load profile for domestic buildings in the UK. *Energy Build.* **2005**, *37*, 663–671. [CrossRef]
3. Li, X.; Yao, R. A machine-learning-based approach to predict residential annual space heating and cooling loads considering occupant behaviour. *Energy* **2020**, *212*, 118676. [CrossRef]
4. Lopes, M.N.; Lamberts, R. Development of a metamodel to predict cooling energy consumption of HVAC systems in office buildings in different climates. *Sustainability* **2018**, *10*, 4718. [CrossRef]
5. Liu, H.; Liu, Y.; Guo, X.; Wu, H.; Wang, H.; Liu, Y. An energy consumption prediction method for HVAC systems using energy storage based on time series shifting and deep learning. *Energy Build.* **2023**, *298*, 113508. [CrossRef]
6. Liu, Z.; Zhang, X.; Sun, Y.; Zhou, Y. Advanced controls on energy reliability, flexibility and occupant-centric control for smart and energy-efficient buildings. *Energy Build.* **2023**, *297*, 113436. [CrossRef]
7. Boerstra, A.C. New health & comfort promoting CEN standard. *REHVA J.* **2019**, *4*, 2. Available online: <https://www.rehva.eu/rehva-journal/chapter/new-health-comfort-promoting-cen-standard> (accessed on 21 December 2024).
8. Nuthakki, S.; Kulkarni, C.S.; Kathiriya, S.; Nuthakki, Y. Artificial Intelligence Applications in Natural Gas Industry: A Literature Review. *Int. J. Eng. Adv. Technol.* **2024**, *13*, 64–70. [CrossRef]
9. Nishimwe, A.M.R.; Reiter, S. Using artificial intelligence models and degree-days method to estimate the heat consumption evolution of a building stock until 2050: A case study in a temperate climate of the Northern part of Europe. *Clean. Responsible Consum.* **2022**, *5*, 100069. [CrossRef]
10. Sadeghian Broujeny, R.; Ben Ayed, S.; Matalah, M. Energy Consumption Forecasting in a University Office by Artificial Intelligence Techniques: An Analysis of the Exogenous Data Effect on the Modeling. *Energies* **2023**, *16*, 4065. [CrossRef]
11. Ravi, A.; Chaturvedi, V.; Shafique, M. ViT4Mal: Lightweight Vision Transformer for Malware Detection on Edge Devices. *ACM Trans. Embed. Comput. Syst.* **2023**, *22*, 1–26. [CrossRef]
12. Khalel, A.; El-Saban, M. Automatic Pixelwise Object Labeling for Aerial Imagery Using Stacked U-Nets. *arXiv* **2018**, arXiv:1803.04953.
13. Wei, P.; Xia, S.; Jiang, X. Energy saving recommendations and user location modeling in commercial buildings. In Proceedings of the UMAP 2018—Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization, Singapore, 8–11 July 2018; pp. 3–11. [CrossRef]
14. Zhao, H.X.; Magoulès, F. A review on the prediction of building energy consumption. *Renew. Sustain. Energy Rev.* **2012**, *16*, 3586–3592. [CrossRef]

15. Moradzadeh, A.; Mohammadi-Ivatloo, B.; Abapour, M.; Anvari-Moghaddam, A.; Roy, S.S. Heating and Cooling Loads Forecasting for Residential Buildings Based on Hybrid Machine Learning Applications: A Comprehensive Review and Comparative Analysis. *IEEE Access* **2022**, *10*, 2196–2215. [\[CrossRef\]](#)
16. Yadav, Y.; Kowli, A. Data-driven simulation for energy and local comfort optimization: Case study of a laboratory. *J. Build. Eng.* **2022**, *54*, 104612. [\[CrossRef\]](#)
17. Mahjoub, S.; Chrifi-Alaoui, L.; Marhic, B.; Delahoche, L. Predicting Energy Consumption Using LSTM, Multi-Layer GRU and Drop-GRU Neural Networks. *Sensors* **2022**, *22*, 4062. [\[CrossRef\]](#)
18. Ghahramani, A.; Dutta, K.; Becerik-Gerber, B. Energy trade off analysis of optimized daily temperature setpoints. *J. Build. Eng.* **2018**, *19*, 584–591. [\[CrossRef\]](#)
19. Chen, Z.; Tao, Z.; Chang, A. A data-driven approach to optimize building energy performance and thermal comfort using machine learning models. In Proceedings of the 2021 1st International Conference on Control and Intelligent Robotics, Guangzhou, China, 18–20 June 2021; pp. 464–469. [\[CrossRef\]](#)
20. Almeida, R.; Georgieva, P.; Martins, N. Heating setpoint recommendation strategy for thermal comfort and energy consumption optimization. *Energy Build.* **2023**, *296*, 113406. [\[CrossRef\]](#)
21. Talami, R.; Dawoodjee, I.; Ghahramani, A. Demystifying energy savings from dynamic temperature setpoints under weather and occupancy variability. *Energy Built Environ.* **2023**, *5*, 878–888. [\[CrossRef\]](#)
22. Ghahramani, A.; Jazizadeh, F.; Becerik-Gerber, B. A knowledge based approach for selecting energy-aware and comfort-driven HVAC temperature set points. *Energy Build.* **2014**, *85*, 536–548. [\[CrossRef\]](#)
23. Li, Z.; Meng, Q.; Wei, Y.; Zhang, L.; Sun, Z.; Lei, Y.; Yang, L.; Yan, X. Dynamic room temperature setpoints of air-conditioning demand response based on heat balance equations with thermal comfort model as constraint: On-site experiment and simulation. *J. Build. Eng.* **2023**, *65*, 105798. [\[CrossRef\]](#)
24. Cui, X.; Zhu, J.; Jia, L.; Wang, J.; Wu, Y. A novel heat load prediction model of district heating system based on hybrid whale optimization algorithm (WOA) and CNN-LSTM with attention mechanism. *Energy* **2024**, *312*, 133536. [\[CrossRef\]](#)
25. Miles, H. A Sequential Modelling Approach for Indoor Temperature Prediction and Heating Control in Smart Buildings. *arXiv* **2020**, arXiv:2009.09847.
26. Yang, Y.; Yan, J. Applied sciences A Heat Load Prediction Method for District Heating Systems Based on the AE-GWO-GRU Model. *Appl. Sci.* **2024**, *14*, 5446. [\[CrossRef\]](#)
27. Qaisar, I.; Sun, K.; Zhao, Q.; Xing, T.; Yan, H. Multi-Sensor-Based Occupancy Prediction in a Multi-Zone Office Building with Transformer. *Buildings* **2023**, *13*, 2002. [\[CrossRef\]](#)
28. Aryal, A.; Becerik-Gerber, B. Energy consequences of Comfort-driven temperature setpoints in office buildings. *Energy Build.* **2018**, *177*, 33–46. [\[CrossRef\]](#)
29. Jing, Z.; Chen, X.; Bu, L.; Xu, W.; Chen, J.; Yu, K.; Shen, J.; Jing, Z.; Chen, X.; Bu, L.; et al. Intelligent optimization for building energy management considering indoor heat transfer Intelligent optimization for building energy management considering indoor heat transfer. *Innov. Energy* **2024**, *1*, 1–10. [\[CrossRef\]](#)
30. Meng, Q.; Xi, Y.; Zhang, X.; Mourshed, M.; Hui, Y. Evaluating multiple parameters dependency of base temperature for heating degree-days in building energy prediction. *Build. Simul.* **2021**, *14*, 969–985. [\[CrossRef\]](#)
31. Sahu, A.K.; Kumar Gupta, P.; Dohare, A.K.; Kumar Singh, A.; Mishra, A.; Rao, A.; Jha, S. Stock Market Prediction Using Machine Learning. In Proceedings of the 2023 International Conference on Computational Intelligence, Communication Technology and Networking (CICTN), Ghaziabad, India, 20–21 April 2023; pp. 512–515. [\[CrossRef\]](#)
32. Gers, F.A.; Schmidhuber, J.; Cummins, F. Learning to forget: Continual prediction with LSTM. *Neural Comput.* **2000**, *12*, 2451–2471. [\[CrossRef\]](#)
33. Chung, H.; Shin, K.S. Genetic algorithm-optimized long short-term memory network for stock market prediction. *Sustainability* **2018**, *10*, 3765. [\[CrossRef\]](#)
34. Van Houdt, G.; Mosquera, C.; Nápoles, G. A review on the long short-term memory model. *Artif. Intell. Rev.* **2020**, *53*, 5929–5955. [\[CrossRef\]](#)
35. Beltagy, I.; Peters, M.E.; Cohan, A. Longformer: The Long-Document Transformer. *arXiv* **2020**, arXiv:2004.05150.
36. Patterson, D.; Gonzalez, J.; Le, Q.; Liang, C.; Munguia, L.-M.; Rothchild, D.; So, D.; Texier, M.; Dean, J. Carbon Emissions and Large Neural Network Training. *arXiv* **2021**, arXiv:2104.10350.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.