Intelligent Environment Monitoring and Control

Author
Muhammad Faizan Faiz
Regn Number
362940

Supervisor
Dr. Sara Ali

ROBOTICS AND INTELLIGENT MACHINE ENGINEERING
SCHOOL OF MECHANICAL & MANUFACTURING ENGINEERING
NATIONAL UNIVERSITY OF SCIENCES AND TECHNOLOGY
ISLAMABAD

May 2023

Intelligent Environment Monitoring and Control

Author

Muhammad Faizan Faiz

Regn Number

362940

A thesis submitted in partial fulfillment of the requirements for the degree of

MS Robotics and Intelligent Machine Engineering

Thesis Supervisor:

Dr. Sara Ali

Thesis Supervisor's Signature: ____

ROBOTICS AND INTELLIGENT MACHINE ENGINEERING

SCHOOL OF MECHANICAL & MANUFACTURING ENGINEERING

NATIONAL UNIVERSITY OF SCIENCES AND TECHNOLOGY,

ISLAMABAD

May 2023

Declaration

I certify that this research work titled "*Intelligent Environment Monitoring and Control*" is my own work. The work has not been presented elsewhere for assessment. The material that has been used from other sources it has been properly acknowledged / referred.

Signature of Student

2021-NUST-Ms-RIME-362940

Plagiarism Certificate (Turnitin Report)

This thesis has been checked for Plagiarism. Turnitin report endorsed by Supervisor is attached.

Signature of Student

Registration Number

362940

Signature of Supervisor

ORIGINALITY REPORT

7%	7%	1%	7%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Higher Education Commission Pakistan Student Paper	4%
2	Submitted to Middle East College of Information Technology Student Paper	1%
3	www.journaltocs.ac.uk Internet Source	1%
4	repository.londonmet.ac.uk Internet Source	1%

Exclude quotes	On	Exclude matches	< 1%
Exclude bibliography	On		

Copyright Statement

Acknowledgements

I am thankful to my Creator Allah Subhana-Watala to have guided me throughout this work at every step and for every new thought which You setup in my mind to improve it. Indeed I could have done nothing without Your priceless help and guidance. Whosoever helped me throughout the course of my thesis, whether my parents or any other individual was Your will, so indeed none be worthy of praise but You.

I am profusely thankful to my beloved parents who raised me when I was not capable of walking and continued to support me throughout every department of my life.

I would also like to express special thanks to my supervisor Dr. Sara Ali for his help throughout my thesis and for courses which she has taught me. I can safely say that I haven't learned any other engineering subject in such depth than the ones which she has taught.

I would also like to pay special thanks to Dr. Muhmmad Sajid for his tremendous support and cooperation. Each time I got stuck in something, he came up with the solution. Without his help I wouldn’t have been able to complete my thesis. I appreciate his patience and guidance throughout the whole thesis. I would also like to acknowledge AIMS Lab, SINES that provided me with much conducive research environment and facilities.

I would also like to thank Dr. Kashif Javed and Dr. Yasir Ayaz for being on my thesis guidance and evaluation committee and express my special thanks to Saeed Iqbal, Ali Javaid and Muhamad Faizan Qadir for their help in setting up the experimental setup.

Finally, I would like to express my gratitude to all the individuals who have rendered valuable assistance to my study.

Dedicated to my exceptional parents and adored siblings whose tremendous support and cooperation led me to this wonderful accomplishment.

Abstract

Heating, Ventilation, and Air Conditioning (HVAC) systems play a vital role in building energy management by controlling the indoor environment and ensuring the occupant's comfort. These systems are responsible for regulating the temperature and air quality inside buildings, thereby creating a comfortable and healthy indoor environment for occupants. However, the energy consumption of HVACs contributes significantly towards overall energy usage of a building and carbon footprint creating a challenge for building energy management. To address this challenge, this research proposes the development of a predictive model for HVAC temperature forecasting using Machine Learning (ML) algorithms to optimize energy efficiency while maintaining thermal comfort in buildings. The study focuses on comparing the performance of Transformer Neural Networks and CNN-LSTM, a seq2seq model combining Convolutional Neural Networks (CNN) and Long-Short Term Memory (LSTM) on multiple forecasting horizons. Both models are validated using data obtained from multiple devices which are deployed in a room verified by feedback survey forms filled by occupants. The transformer model outperformed, achieving an R2 score of 0.936 at a 1 minute forecasting horizon, surpassing the performance of CNN-LSTM model at all tested forecasting horizons. The transformer model yielded significant energy savings thereby reducing energy consumption by almost 50 percent compared to the non-AI conventional methods, particularly at forecasting horizons of 1 minute and 60 minutes, while the occupant survey also favored a 60-minute forecasting horizon indicating that the proposed model can effectively balance energy efficiency with occupant comfort. The performance of transformer model particularly with a 60-minute forecasting horizon underscores its potential to optimize energy efficiency while ensuring thermal comfort in building energy management systems.

Key Words: *Heating Ventilation and Air Conditioning (HVAC), Transformer-based model, Energy optimization, Indoor environment*

Table of Contents

Declaration i

Plagiarism Certificate (Turnitin Report) ii

Copyright Statement iv

Acknowledgements v

Abstract vii

Table of Contents viii

List of Figures x

List of Tables xii

CHAPTER 1: INTRODUCTION 1

1.1 Background, Scope and Motivation 1

1.1.1 Machine Learning and Time Series Forecasting 4

1.1.2 HVACs and Data Driven Control 6

CHAPTER 2: LITERATURE REVIEW 8

CHAPTER 3: METHODOLOGY 16

3.1 Data Acquisition 16

3.1.1 ASHRAE standard 55 and Survey Data 16

3.1.2 Air Quality 17

3.1.3 Power Quality 18

3.1.4 Indoor Ambiance 20

3.1.5 Meteorological Data 20

3.2 Data Pre-Processing: 23

3.3 Forecasting Models 24

3.3.1 CNN-LSTM 25

3.3.2 Transformer Neural Networks 28

3.4 Evaluation Metrics 32

CHAPTER 4: MATERIALS AND METHODS 34

4.1 Subjects 34

4.2 Environmental Setup 34

4.3 Experiment Design 36

CHAPTER 5: RESULTS AND DISCUSSION 42

5.1 Model Evaluation 42

5.2 Energy Optimization 47

5.3 Thermal Comfort 52

CHAPTER 6: CONCLUSION 55

REFERENCES 56

List of Figures

Figure 1: A schematic representation of the proposed methodology illustrating the data acquisition process into cloud storage utilized by a deep learning model that generates temperature predictions to control HVAC. 15

Figure 2 Air Quality Analyzer 18

Figure 3 Power Quality Analyser 19

Figure 4: Schematic of a CNN- LSTM architecture illustrating the input data is first fed into several convolution layers. The output from these layers is then passed on to LSTM cells, and finally to a dense layer that generates predictions. 26

Figure 5: Architecture diagram of transformer neural network representing the flow of input sequences through multiple layers. 29

Figure 6: Trend of indoor and outdoor temperatures 35

Figure 7: Layout of the experimentation room 35

Figure 8 IR-Controller 39

Figure 9: Flow diagram of forecasting temperature on HVAC 40

Figure 10: Actual vs Predicted temperatures using CNN-LSTM with temperature on y-axis and Total no. of predictions in terms of days are represented on x-axis. 45

Figure 11: Actual vs Predicted Temperatures using Transformers with Temperature on y-axis and Total no. of predictions in terms of days are represented on x-axis. 47

Figure 12 Bar plot representing the net power consumption over 3 time periods in a day. 49

Figure 13 Bar Plot representing daily Average Power Consumption for all models. 50

Figure 14 Trend of Power Consumption among all models 51

Figure 15 Nested donuts representing the thermal comfort of the occupants for models on different forecasting horizons at three multiple time periods of the day where outermost donut represents morning, middle donut represents afternoon and innermost donut repres 53

List of Tables

Table 1 Details of sensors used in the experiments 22

Table 2 List of Hyperparameters.. 27

Table 3 Comparison of performance between CNN-LSTM and Transformers at multiple forecasting horizons 43

CHAPTER 1: INTRODUCTION

1.1 Background, Scope and Motivation

Heating, Ventilation, and Air Conditioning (HVAC) systems have become an integral part of modern building design and operation. Their significance lies in their ability to maintain optimal indoor air quality, thermal comfort, and energy efficiency. As environmental sustainability and occupant well-being take precedence in our society, understanding and optimizing HVAC systems have gained increased importance. The primary objective of HVAC systems is to create and sustain a comfortable indoor environment by regulating temperature, humidity, and air quality. By effectively controlling these factors, HVAC systems ensure that occupants can work, relax, or engage in activities in a comfortable setting, regardless of external weather conditions. Furthermore, these systems play a crucial role in removing indoor pollutants, allergens, and excessive humidity, thus promoting healthier indoor spaces and reducing the risk of respiratory problems. The prevalence of HVAC systems can be observed worldwide, as they are implemented in various types of buildings, including residential, commercial, and industrial facilities.

However, the extensive usage of these systems comes at a cost. HVAC systems consume a significant amount of energy, primarily derived from fossil fuels, contributing to greenhouse gas emissions and exacerbating environmental concerns on a global scale. According to a study conducted by Longo et al. [1], it has been revealed that HVAC systems account for a substantial portion of global energy consumption. In fact, their findings indicate that a staggering 38% of all energy consumed can be attributed to HVAC systems. This energy usage is distributed across various sectors, with residential sources accounting for 32% and tertiary industries contributing as much as 47%. In residential settings, heating, cooling, and ventilation requirements directly impact energy demand. Residential buildings alone contribute to approximately one-third of the total energy consumed by HVAC systems. These systems also play a vital role in tertiary industries, encompassing commercial buildings, office spaces, educational institutions, and other non-residential facilities. These sectors account for a significant share of energy consumed by HVAC systems. Also, the Earth's climate system has experienced a steady increase in the global mean

Footnote may be given with Font size 10 of Times New Roman.

surface temperature, with a staggering rise of 1.1^{o}C since the advent of the industrial revolution [2]. This substantial temperature increase has had far-reaching consequences, including heatwaves, extreme weather events, and altered climatic patterns across the globe. As a direct response to these rising temperatures, the demand for HVAC systems has intensified, as they provide essential cooling and temperature regulation mechanisms in various settings. From residential buildings to commercial establishments, educational institutions, healthcare facilities, and industrial complexes, HVAC systems have become an integral part of maintaining comfortable and habitable indoor environments. They are indispensable in ensuring the well-being, productivity, and safety of occupants in the face of increasingly challenging climatic conditions.

The unprecedented rise in global mean surface temperature has had a cascading effect on various sectors, necessitating the widespread adoption of HVAC systems to counter the adverse impacts of heat stress. Heatwaves, characterized by prolonged periods of intense heat, pose severe health risks, particularly to vulnerable populations such as the elderly, children, and individuals with pre-existing medical conditions. HVAC systems, equipped with efficient cooling mechanisms, play a crucial role in mitigating these risks by maintaining suitable indoor temperatures and reducing heat-related illnesses and fatalities. Moreover, the effects of climate change extend beyond human health and comfort. In sectors such as agriculture, manufacturing, and data centres, temperature regulation is paramount for preserving product quality, ensuring operational efficiency, and safeguarding critical equipment from heat-related damage. HVAC systems are relied upon to provide optimal temperature and humidity conditions, thereby supporting essential processes, enhancing productivity, and minimizing economic losses. Given the undeniable influence of climate change and the consequent rise in global temperatures, the increased adoption of HVAC systems is a pragmatic response to the need for climate adaptation. Apart from the impact of man-made climate change, it is crucial to acknowledge the additional environmental concerns associated with HVAC systems. These systems not only contribute to the emission of greenhouse gases like carbon dioxide but also release dangerous particulate matter and other pollutants into the air. Furthermore, the use of refrigerants in HVAC systems, which are often toxic and harmful, can have detrimental effects on the ozone layer, exacerbating ecological imbalances [3]. The cumulative impact of these emissions poses significant challenges to the environment and necessitates a comprehensive approach to mitigate their adverse effects.

Footnote may be given with Font size 10 of Times New Roman.

Considering that individuals spend approximately 90% of their time indoors, it becomes increasingly imperative to prioritize clean air and suitable indoor temperatures. People aspire to breathe air of good quality and to reside in spaces that maintain a comfortable thermal environment. These systems act as a vital defence against the adverse effects of heatwaves and extreme temperatures, enabling individuals, communities, and industries to adapt and thrive in an ever-changing climate landscape.

As societies increasingly prioritize sustainability and energy efficiency, addressing the energy consumption and environmental impact of HVAC systems has become an urgent imperative. To mitigate the environmental impact and enhance energy efficiency, various strategies are being employed in the design, operation, and maintenance of HVAC systems. One of the key approaches is the utilization of advanced technologies and equipment. High-efficiency heating and cooling equipment, such as heat pumps and geothermal systems, can significantly reduce energy consumption compared to conventional systems. Additionally, the integration of renewable energy sources, such as solar panels or wind turbines, into HVAC systems can further reduce reliance on fossil fuels and decrease greenhouse gas emissions. Furthermore, optimizing the design and layout of HVAC systems plays a crucial role in improving energy efficiency. Proper insulation, efficient ductwork, and zoning techniques that enable different temperature settings in various areas of a building can contribute to substantial energy savings. The installation of smart controls and programmable thermostats allows for precise temperature scheduling and adaptive operation, optimizing energy usage based on occupancy patterns and environmental conditions. Building occupants also play a significant role in achieving energy efficiency in HVAC systems. Educating occupants about energy-saving practices, such as adjusting thermostat settings, maintaining proper ventilation, and minimizing air leaks, can lead to substantial energy conservation. Additionally, occupants' behavior and awareness regarding energy usage can be influenced through energy management programs and incentives.

As HVAC systems continue to evolve, new technologies and approaches are being explored to enhance their performance. The development of smart HVAC systems, which utilize advanced sensors, data analytics, and machine learning algorithms, enables real-time monitoring and control of indoor environmental conditions. These systems can adapt to dynamic parameters,

Footnote may be given with Font size 10 of Times New Roman.

such as occupancy levels, outdoor weather conditions, and indoor air quality, to optimize energy consumption and occupant comfort. The development of such smart HVAC systems becomes essential that predicts the temperature to optimize energy and thermal comfort. Accurate predictions of indoor temperature enable the optimization of HVAC system efficiency while simultaneously ensuring the comfort and well-being of occupants [4]. By deploying forecasting models, it becomes possible to judiciously manage energy resources, thereby fostering sustainability and minimizing environmental impacts. The accurate estimation of indoor temperature serves as a fundamental aspect of creating comfortable indoor environments. It allows for a fine-tuning of HVAC systems' operation, ensuring that the desired temperature ranges are consistently achieved. This not only enhances occupant comfort but also optimizes energy consumption by avoiding excessive heating or cooling. The efficient utilization of HVAC systems not only reduces energy waste but also contributes to cost savings and promotes environmental stewardship. Moreover, precise temperature forecasting models enable proactive energy management strategies, as they provide valuable insights into the expected heating and cooling demands of a space. By anticipating temperature fluctuations and occupant preferences, building operators can implement pre-emptive measures to optimize energy consumption. This proactive approach ensures that HVAC systems operate in a targeted and energy-efficient manner, reducing unnecessary energy usage and minimizing the associated environmental impacts.

1.1.1 Machine Learning and Time Series Forecasting

Machine learning, a subfield of artificial intelligence, has emerged as a powerful tool in various domains, including finance, healthcare, marketing, and supply chain management. One particularly important application of machine learning is in time series forecasting, where it has revolutionized the field of predictive analytics. By leveraging advanced algorithms and statistical models, machine learning techniques enable accurate predictions and insights into the behavior and trends of time-dependent data.

Time series data refers to a sequence of observations collected over time, typically at regular intervals. It encompasses a wide range of real-world phenomena, such as stock prices, weather patterns, energy consumption, website traffic, and many others. Analyzing and forecasting time series data is crucial for decision-making, resource allocation, risk management, and optimization in various industries. Traditional statistical methods for time series forecasting often rely on assumptions of linearity, stationarity, and specific data distribution characteristics. However, these assumptions may not hold true in complex real-world scenarios where time series data often exhibit nonlinear patterns, trends, seasonality, and irregularities. Machine learning techniques, on the other hand, excel in handling such complexities and offer more flexible and robust models. Machine learning algorithms employ data-driven approaches to automatically learn patterns, relationships, and dependencies from historical time series data. They can capture nonlinearity, handle missing values, adapt to changing dynamics, and incorporate multiple variables into the forecasting process. By analyzing large volumes of data and detecting intricate patterns, machine learning models can generate accurate predictions and valuable insights that were previously unattainable.

One of the key advantages of machine learning in time series forecasting is its ability to handle high-dimensional and unstructured data. With the exponential growth of data sources and the advent of the Internet of Things (IoT), organizations are faced with vast amounts of time series data from diverse sources. Machine learning techniques, such as deep learning, recurrent neural networks, and ensemble models, can effectively extract relevant features and capture complex temporal dependencies from such data, enabling improved forecasting accuracy. Furthermore, machine learning algorithms can adapt and self-adjust over time, making them well-suited for dynamic and evolving time series data. They can automatically update their models based on new observations and continuously improve their forecasting performance. This adaptability is particularly valuable in domains where time series data exhibit non-stationary behavior, seasonality shifts, or sudden changes due to external events or market dynamics.

Time series forecasting also plays a vital role in HVAC temperature control, enabling accurate predictions of future temperature patterns and facilitating efficient heating and cooling operations. In HVAC systems, time series data consists of temperature measurements taken at regular intervals, reflecting the dynamic nature of indoor and outdoor environments. By analyzing historical temperature data, time series forecasting models can identify underlying trends, seasonal patterns, and cyclic variations, enabling HVAC systems to anticipate temperature changes and

Footnote may be given with Font size 10 of Times New Roman.

adjust heating and cooling settings accordingly. This proactive approach improves energy efficiency by ensuring that HVAC systems operate optimally and avoid unnecessary energy consumption.

The integration of machine learning with time series forecasting has resulted in significant advancements in predictive analytics. These techniques have been applied to diverse fields, such as weather forecasting, financial markets, energy demand forecasting, sales forecasting, anomaly detection, and predictive maintenance. The accurate predictions and insights derived from machine learning models enable organizations to make informed decisions, optimize resource allocation, mitigate risks, and gain a competitive edge in a rapidly changing environment.

1.1.2 HVACs and Data Driven Control

As concerns about energy consumption, environmental impact, and occupant well-being grow, there is an increasing need for advanced control strategies that optimize HVAC system performance. The emergence of data-driven control approaches, enabled by advancements in sensor technology, data analytics, and machine learning, has revolutionized the way HVAC systems are managed. By leveraging data and analytics, data-driven control techniques offer the potential to enhance energy efficiency, occupant comfort, and system reliability in HVAC operations.

Traditional control strategies for HVAC systems have often relied on fixed setpoints and pre-determined control algorithms. However, these approaches may not adapt well to dynamic conditions, changing occupant requirements, or variations in external factors such as weather patterns. Data-driven control, on the other hand, leverages real-time sensor data and historical information to optimize control actions and dynamically adjust HVAC system parameters. The foundation of data-driven control in HVAC systems lies in the availability of sensor data that captures relevant environmental and operational information. Sensors placed throughout the building monitor variables such as temperature, humidity, occupancy, and air quality. This data is then processed and analyzed using advanced analytics techniques, including machine learning algorithms, to extract patterns, identify correlations, and make informed decisions about system operation.

Footnote may be given with Font size 10 of Times New Roman.

Data-driven control offers several key advantages for HVAC systems. Firstly, it enables personalized comfort control, as it can learn occupant preferences and adjust settings accordingly. By analyzing historical data on individual occupant behavior, the system can optimize temperature, airflow, and ventilation rates in different zones of the building, tailoring the HVAC operation to specific needs and enhancing occupant comfort. Secondly, data-driven control enables real-time optimization of energy consumption. By continuously monitoring and analyzing sensor data, the system can identify energy-saving opportunities, adapt HVAC operation based on occupancy patterns and external conditions, and dynamically adjust setpoints and control strategies. This optimization can lead to significant energy savings without compromising comfort, resulting in lower utility costs and reduced environmental impact.

However, data-driven control in HVAC systems also presents challenges. Issues such as data quality, privacy concerns, model interpretability, and computational requirements must be carefully addressed. Additionally, the implementation of data-driven control requires expertise in data analytics, machine learning, and HVAC system knowledge to ensure effective integration and operation.

Footnote may be given with Font size 10 of Times New Roman.

CHAPTER 2: LITERATURE REVIEW

According to B. Dong et al [5] and A. Foucquier et al [6], currently Physical knowledge-based, data-driven, and hybrid models are the three types of models used most frequently to predict energy consumption and temperature.

The first type, physical knowledge-based models also called white-box models, plays a crucial role in the development of prediction models by employing mathematical representations of the fundamental principles of physics [7]. These models rely on a deep understanding of the underlying physical processes governing the behavior of the system being modeled. By incorporating equations derived from thermodynamics, heat transfer, and fluid dynamics, white-box models aim to provide accurate predictions based on well-established scientific principles. However, it is important to note that the utilization of white-box models presents certain drawbacks. One significant limitation is the requirement for extensive knowledge and understanding of the specific building or area being modeled. To develop an effective white-box model, researchers and practitioners must possess comprehensive information about the structural characteristics, materials, and system configurations of the building. This necessitates detailed data collection and analysis, including factors such as building layouts, insulation properties, ventilation systems, and occupancy patterns. Another challenge associated with white-box models is the presence of assumptions within the mathematical equations used. These assumptions are based on simplifications and idealizations of the real-world system behavior. While these assumptions are necessary to formulate tractable mathematical models, they may not always perfectly align with the actual behavior observed in practice. Variations in operating conditions, occupant behavior, or unexpected system interactions can introduce discrepancies between the model predictions and real-world outcomes [8]. Overall, these models incorporate a deep understanding of the physical properties and behaviors of HVAC systems and the surrounding environment. By leveraging fundamental scientific principles, physical knowledge-based models aim to accurately simulate the energy consumption and temperature dynamics of a building or space. They typically require detailed inputs such as building specifications, system configurations, and environmental conditions to generate predictions. These models are valuable

Footnote may be given with Font size 10 of Times New Roman.

for their ability to provide insights into the underlying physical processes governing energy consumption and temperature regulation.

Machine Learning (ML) and other data-driven black-box techniques have gained significant attention and demonstrated superior performance compared to traditional methods in various domains. These approaches leverage the power of advanced algorithms to analyze large datasets and extract complex patterns and relationships. In the context of thermal dynamics modeling, black-box techniques based on data have shown promising capabilities in simulating and predicting temperature dynamics without the need to explicitly identify zone-specific characteristics, such as heat capacity and size [10]. The use of machine learning (ML) and data-driven black-box techniques in thermal dynamics modeling offers increased flexibility and adaptability across various building types and environments. These models can handle diverse settings and variations in thermal dynamics without requiring manual adjustments for specific zone characteristics. This scalability and generalizability are particularly advantageous in scenarios involving multiple zones or complex buildings with varying thermal characteristics. However, it is important to acknowledge that black-box models come with limitations. While they excel in capturing complex relationships, their lack of interpretability can hinder understanding of the underlying physical processes. Unlike white-box models that provide explicit equations and assumptions, black-box models are treated as "black boxes" with non-interpretable internal mechanisms. This lack of interpretability makes it challenging to gain insights into the factors driving predictions or identify specific physical phenomena at play.

Gray-box models represent a hybrid approach that seeks to bridge the gap between white-box models and the complexities of actual building conditions. These models combine the transparency and interpretability of white-box techniques with the incorporation of actual data, moving away from relying solely on assumptions. By integrating real-world data, gray-box models aim to provide more accurate predictions of thermal dynamics and energy consumption in buildings [9]. However, gray-box models do come with their own set of limitations. One significant drawback is the requirement for substantial prior information about the building being modeled. This includes detailed knowledge of the building's physical characteristics, such as its layout, construction materials, and HVAC system specifications. Without comprehensive data, the

Footnote may be given with Font size 10 of Times New Roman.

accuracy of the gray-box model may be compromised. Additionally, gray-box models often rely on mathematical presumptions and simplifications to achieve the desired accuracy. These assumptions are necessary to strike a balance between model complexity and computational efficiency. However, these simplifications may introduce certain degrees of uncertainty and limitations to the model's predictive capabilities. Therefore, careful consideration and validation of the assumptions made within the gray-box model are necessary to ensure reliable predictions.

Indoor temperature forecasting has garnered significant attention in recent years, prompting extensive research into various machine learning techniques that aim to deliver precise predictions. Among the approaches explored, support vector machines (SVM) [11], [12], random forests [13], extreme gradient boosting (XGBoost) [14], [15], and artificial neural networks (ANNs) [16] have emerged as prominent contenders. These techniques offer unique advantages and have been successfully employed in achieving accurate indoor temperature forecasts. SVMs leverage statistical learning theory to find optimal decision boundaries, random forests employ ensemble learning with decision trees, XGBoost employs gradient boosting algorithms, and ANNs mimic the biological neural network structure to model complex relationships. With ongoing advancements, the field of indoor temperature forecasting continues to witness improvements in prediction accuracy and application potential. An artificial neural network (ANN) is a nonlinear statistical technique widely used for various machine learning tasks. It comprises three fundamental components: an input layer, a hidden layer, and an output layer. While ANNs have proven effective in modeling complex relationships, they encounter challenges when applied to time series data, particularly in capturing temporal dependencies between input and output data. To address these challenges, time correlation and recurrent neural networks (RNNs) have emerged as valuable tools [17]. LSTM (Long Short-Term Memory) is a variant of the recurrent neural network (RNN) algorithm that has gained significant popularity in the field of time-series forecasting. One of the main challenges faced by traditional RNNs is the issue of exploding or vanishing gradients, which can hinder the network's ability to learn long-term dependencies. LSTM addresses this problem by introducing a specialized memory cell that allows the network to selectively retain and propagate information over longer sequences. Researchers, such as Xu et al. [18], have recognized the potential of LSTM for indoor temperature prediction and have conducted studies comparing it with other models. In their work, they focused on predicting indoor

Footnote may be given with Font size 10 of Times New Roman.

temperature with a forecast horizon of 5 minutes and compared the performance of LSTM against Support Vector Machines (SVM) and Back Propagation Neural Networks (BPNN). The results of the study demonstrated that LSTM outperformed the other models in terms of predictive accuracy. Specifically, the evaluation metric used, R-squared (R^2), indicated that LSTM achieved an R^2 value of 0.79. This implies that approximately 79% of the variability in the indoor temperature could be explained by the LSTM model. By comparison, SVM and BPNN exhibited lower R^2 values, suggesting lower performance in capturing the intricate relationships within the data.

In a study conducted by Mtibaa et al. [19], the performance of two models, NNARX (Neural Network-based Autoregressive Model with Exogenous Inputs) and LSTM, was compared in the context of a 6-hour prediction model. The evaluation metrics used were MAPE (Mean Absolute Percentage Error) and MAE (Mean Absolute Error). However, it is worth noting that the execution time for NNARX was considerably slower compared to LSTM, which is an important practical consideration in real-time applications. In contrast, LSTM demonstrated faster execution time, completing a 4-hour prediction in less than 15 seconds. On the other hand, NNARX took more than 5 minutes to achieve the same prediction. Despite the difference in execution time, both models yielded similar results in terms of prediction accuracy. Additionally, the study also examined the standard deviation of error for both models. The standard deviation of error for NNARX was reported to be 0.45 °C slightly larger than the LSTM. This indicates the variability or spread of errors around the predicted values. While LSTM performed better in terms of accuracy and execution time, the slightly larger standard deviation of error for NNARX suggests that it may have a slightly higher level of inconsistency in its predictions compared to LSTM. In another study conducted by Fang et al. [20], a seq2seq LSTM model was proposed for predicting the indoor temperature of different building zones. The model consisted of three architectures: LSTM-Dense, LSTM-LSTM, and LSTM-Dense-LSTM. The key concept behind this research was the utilization of an encoder-decoder framework in the seq2seq model. The encoder component captured important features from the input sequence and passed them to the decoder, which was responsible for generating forecasts based on the provided information. Among the three architectures examined, LSTM-Dense demonstrated the best performance with an RMSE (Root Mean Square Error) value of 0.458. The RMSE is a measure of the difference between predicted values and actual values, with lower values indicating better accuracy. Another seq2seq model utilized

Footnote may be given with Font size 10 of Times New Roman.

convolutional neural networks (CNN) for feature extraction on the input data, followed using a long short-term memory (LSTM) model to predict the interior temperature fluctuation for the next 120 minutes. This approach aimed to leverage the strengths of CNN in extracting meaningful features from the input sequence and LSTM's ability to capture temporal dependencies. By employing CNN as a feature extraction mechanism, the model could effectively identify relevant patterns and relationships within the input data. This process allowed the model to capture important spatial information and extract higher-level features that contribute to accurate temperature predictions. After the feature extraction phase, the extracted features were fed into an LSTM model. By incorporating an LSTM layer, the model could effectively model the temporal dynamics and dependencies in the interior temperature fluctuations. The results of this seq2seq model were highly promising, with a high R-squared value of above 0.9. The R-squared value indicates the proportion of the variance in the dependent variable (interior temperature) that is explained by the model. A value above 0.9 suggests a strong correlation between the predicted temperature values and the actual temperature values. Jiang et al. [21] utilized an attention-LSTM model that was employed to predict the indoor temperature. The attention mechanism is a mechanism that allows the model to focus on relevant parts of the input sequence while making predictions. The attention-LSTM model was compared to two other recurrent neural network (RNN) variants, namely the Gated Recurrent Units (GRU) and the traditional LSTM. The results of the study indicated that the attention-LSTM model outperformed both the GRU and LSTM models when the forecasting horizon was 30 minutes or greater. The superiority of the attention-LSTM model can be attributed to the attention mechanism, which enables the model to selectively attend to specific parts of the input sequence that are deemed more important for making accurate predictions. This mechanism allows the model to assign higher weights or attention to relevant temporal features, thereby capturing crucial information and improving prediction performance. The comparison results suggest that for longer forecasting horizons, where capturing intricate temporal dependencies becomes more critical, the attention-LSTM model excels in capturing and leveraging the relevant information effectively. The ability of the attention-LSTM model to dynamically focus on important features within the input sequence leads to improved forecasting accuracy compared to the GRU and LSTM models.

However, LSTM-based models do have certain limitations that should be considered. One of the challenges associated with LSTM networks is their ability to effectively capture long-term dependencies in data, particularly when dealing with long sequences. LSTM networks are designed to address the vanishing gradient problem and can capture short-term dependencies effectively, but they may struggle to capture complex long-term dependencies. The vanishing gradient problem refers to the issue where the gradients diminish exponentially over time, making it difficult for the network to propagate useful information across many time steps. This limitation becomes more pronounced as the length of the input sequence increases. Consequently, when training LSTM-based models on long sequences, there is a risk of losing important information and compromising the model's performance. Furthermore, LSTM networks process input sequences sequentially, one timestep at a time. This sequential processing can result in slower training times compared to models that can leverage parallel processing. Each timestep in an input sequence must be processed and propagated through the network before proceeding to the next timestep, which can be computationally demanding and time-consuming.

When comparing RNN-based techniques to transformer-based approaches, it has been observed that transformer models tend to achieve lower prediction errors. This is primarily attributed to the transformer's capability to recognize long-term latent patterns within the entirety of the input data [22]. Unlike RNN-based models, transformers employ self-attention mechanisms to process input sequences. Self-attention allows transformers to consider all timesteps simultaneously, rather than relying on sequential processing. By considering the entire sequence at once, transformers can capture dependencies and patterns across long sequences more effectively. The self-attention mechanism in transformers enables the model to assign different weights or importance to different parts of the input sequence, depending on their relevance to the prediction task. This allows the model to attend to relevant information and disregard irrelevant or redundant information, leading to improved performance in capturing long-term dependencies. The ability of transformers to process input sequences in parallel rather than in a sequential manner is a significant advantage. By considering all timesteps at once, transformers can consider the entire history of the sequence, facilitating the recognition of complex patterns and dependencies that may span across long periods of time.

Footnote may be given with Font size 10 of Times New Roman.

This study aims to investigate the effectiveness of transformers in predicting energy usage data obtained from multiple devices and surveys. While the use of transformer models for time-series modeling in indoor temperature forecasting has received limited attention in the literature, especially for HVAC temperature forecasting, this research explores their potential. Four different forecasting horizons (1, 15, 30, and 60 minutes) were considered, and the models were trained using collected data. Real-time control of the HVAC system was conducted using a specially designed infrared sensor based remote to evaluate the accuracy of the models' forecasts. Energy usage data was collected and compared across different forecasting horizons. Surveys were also conducted to gather information on occupants' thermal comfort. This work is particularly significant as it contributes to power optimization and enhances occupant comfort. The trained model can be applied to optimize HVAC systems in similar setups with sensors and data recording devices, indicating broader applicability. The study's findings provide insights into the viability of using transformers for energy forecasting and their potential role in creating energy-efficient buildings. A schematic representation of the proposed method is given in Figure 1.

Footnote may be given with Font size 10 of Times New Roman.

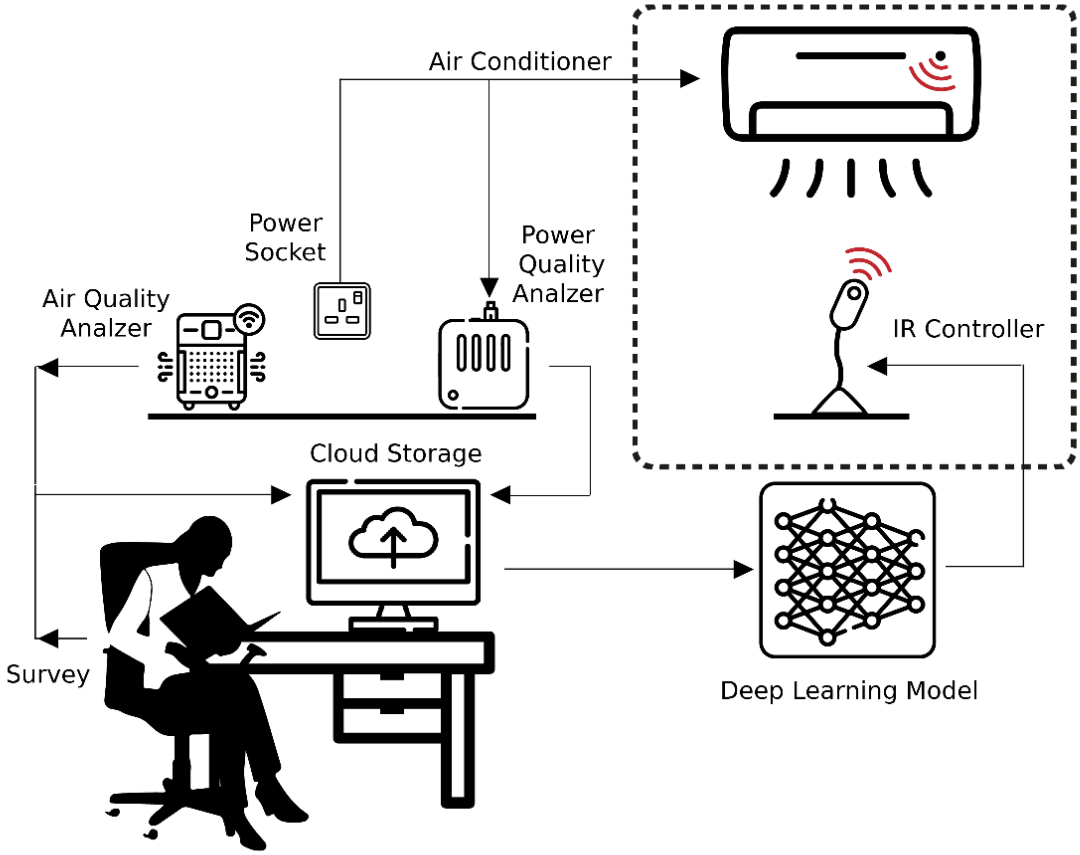

Figure 1: A schematic representation of the proposed methodology illustrating the data acquisition process into cloud storage utilized by a deep learning model that generates temperature predictions to control HVAC.

Footnote may be given with Font size 10 of Times New Roman.

CHAPTER 3: METHODOLOGY

3.1 Data Acquisition

The dataset used in this research encompasses a wide range of information obtained from various sources, including an air quality analyzer, power quality analyzer, and surveys. The data collection process involved sampling the variables at a rate of 1 minute, ensuring a comprehensive representation of the dynamic nature of the indoor environment. To organize the collected data effectively, it was divided into five distinct classes based on the types of information being captured.

3.1.1 ASHRAE standard 55 and Survey Data

The ASHRAE Standard 55, developed by the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE), provides guidelines and recommendations for thermal environmental conditions in buildings [23]. It specifically focuses on factors such as temperature, humidity, air movement, and occupant activity levels to ensure thermal comfort and well-being.

ASHRAE Standard 55 provides guidance on how to assess and consider human factors that influence thermal comfort. This includes factors such as body dimensions, clothing insulation, metabolic rates, and personal preferences. By incorporating these factors into the analysis, the standard helps to establish a more comprehensive understanding of how occupants perceive and experience thermal comfort in different environments.

The class of anthropometric measurements in the dataset is derived from surveys conducted based on the ASHRAE Standard 55. In order to address any uncertainty in the data, the study conducted by Kim et al. [24] proposed the utilization of a minimum of 60 surveys per person to achieve stable predictions of thermal preference. Therefore, a total of 84 surveys were collected per participant to ensure robustness and mitigate potential uncertainties in the data. The survey encompassed various personal parameters including metabolic rate, clothing insulation, Body Mass Index (BMI), and body temperature. The metabolic rate, which indicates the rate at

Footnote may be given with Font size 10 of Times New Roman.

which an individual generates heat, was estimated based on the activity level and clothing type of the occupants, following the guidelines provided by the ASHRAE standard. Clothing insulation, which determines the thermal resistance provided by the clothing worn, was also assessed in the survey. The occupants were asked to provide information about the type and thickness of their clothing, which allowed for the estimation of clothing insulation values. Body Mass Index (BMI), a measure of body composition based on height and weight, was recorded to consider the potential influence of body size on thermal comfort. To measure body temperature, a non-invasive laser thermometer was used. This tool enables surface temperature measurements and is commonly employed for assessing body temperature without physical contact. In addition to collecting objective data, the survey also included the subjective thermal perception of the occupants. The participants were asked to rate their thermal sensation on a 5-point scale, in accordance with the categories defined by the ASHRAE standard 55. The scale ranged from -2 to 2, representing the categories of cool, slightly cool, neutral, slightly warm, and warm, respectively.

These measurements are essential for understanding the individual variations in thermal comfort preferences and responses to indoor temperature conditions. By incorporating the anthropometric data based on the ASHRAE Standard 55, the predictive model can account for the diverse range of occupant profiles and their impact on thermal comfort. This allows for a more personalized and adaptive approach to HVAC temperature forecasting, optimizing energy efficiency while maintaining occupant comfort.

3.1.2 Air Quality

The second class of data, air quality, focused on assessing the quality of indoor air and its potential impact on occupant well-being. To gather this information, an air quality analyser was utilized, capable of measuring multiple parameters related to air pollution and evaluating the overall air quality within the building.

The air quality analyser collected data on 15 different parameters that are known to contribute to air pollution and can have significant implications for indoor environments. These

Footnote may be given with Font size 10 of Times New Roman.

parameters included carbon dioxide (CO_2) levels, particulate matter (PM) concentrations, Volatile Organic Compounds (VOCs), and ozone (O_3) levels, among others.

Figure 2 Air Quality Analyzer

3.1.3 Power Quality

The third class of data, power quality, aimed to monitor and analyse various electrical parameters associated with the HVAC system, including voltage, current, and power consumption.

Footnote may be given with Font size 10 of Times New Roman.

A power quality analyser was employed to measure and record these parameters, providing valuable insights into the energy consumption patterns of the HVAC system and its overall impact on the building's energy usage.

Figure 3 Power Quality Analyser

The power quality analyser was specifically configured to monitor the energy consumption parameters of the HVAC system installed in the room. It collected data on a total of 32 parameters, which included voltage levels, current values, active power, and reactive power.

Footnote may be given with Font size 10 of Times New Roman.

3.1.4 Indoor Ambiance

The fourth class of data, indoor environment ambiance, focused on capturing various variables that contribute to the overall indoor environment within the monitored space. These variables included lighting levels, indoor temperature, humidity, pressure, and occupancy status. Understanding and monitoring these factors are crucial as they significantly influence occupant comfort, well-being, and overall satisfaction within a building.

Lighting levels were measured to assess the intensity and quality of illumination within the space. Proper lighting is essential for creating a comfortable and productive indoor environment, and it plays a vital role in occupant visual comfort and task performance. Indoor temperature and humidity are key parameters that directly affect thermal comfort. Maintaining an optimal temperature and humidity level is important to ensure that occupants feel neither too hot nor too cold and to prevent issues such as condensation or excessive dryness. Pressure measurements were also recorded to monitor any variations in air pressure within the indoor environment. Changes in pressure can impact ventilation and air circulation, which in turn affects indoor air quality and comfort. Occupancy status was monitored to determine the number of people present in the space at different times. Occupancy data is valuable for analysing patterns, understanding occupant behaviour, and evaluating the impact of occupancy on indoor conditions and energy consumption.

3.1.5 Meteorological Data

The meteorological data used in the study played a crucial role in understanding the outdoor conditions and their impact on the indoor environment. These data were obtained from NOAA's (National Oceanic and Atmospheric Administration) weather database [25], a reputable source for comprehensive and reliable meteorological information.

The dataset comprised 13 parameters that encompassed a range of meteorological factors. One of the key parameters was the outdoor temperature, which provided insights into the ambient temperature outside the building. Understanding the outdoor temperature is vital as it directly influences the indoor temperature and the heating or cooling requirements of the HVAC system.

The NOAA (National Oceanic and Atmospheric Administration) weather database is a comprehensive and reliable source of meteorological data. It is maintained by the U.S. government

Footnote may be given with Font size 10 of Times New Roman.

agency responsible for monitoring and forecasting weather and climate conditions. The database collects data from various weather stations across the United States and other regions, ensuring wide coverage and accuracy. NOAA's weather database contains a vast amount of information related to weather conditions, including temperature, humidity, precipitation, wind speed, atmospheric pressure, UV index, and more. These data points are collected at regular intervals from different locations, providing a comprehensive picture of weather patterns and variations.

The sensors used in the devices along with their typology and operational range, to record data are provided in Table 1.

Footnote may be given with Font size 10 of Times New Roman.

Table 1 Details of sensors used in the experiments

Sensor Name	Typology	Operational Range
High Precision Laser Dust Concentration Sensor	Type: Optical sensor • Working principle: Measures the scattered light intensity from dust particles using a laser beam • Applications: Environmental monitoring, indoor air quality control, industrial process control	Particle measurement range 0.3 to 1.0; 1.0 to 2.5; 2.5 to 10 (Micron) Operating temperature range -10 ~ + 60 (Celsius) Operating humidity range 0 ~ 99% Storage temperature range -40 ~ + 80 (Celsius)
Low Concentration Ozone Gas Detection Sensor	Type: Chemical sensor • Working principle: Measures the change in electrical conductivity of a sensing material exposed to ozone gas • Applications: Air purification systems, gas detectors, sterilization equipment	10～1000ppm Ozone
High Concentration Ozone Gas Detection Sensor	Type: Electrochemical sensor • Working principle: Measures the change in electrical current generated by a reaction between ozone and a sensing electrode • Applications: Ozone generators, industrial process control, sterilization equipment	10～1000ppm Ozone
CO2/ TVOC Sensor	Type: Gas sensor • Working principle: Measures the change in electrical conductivity of a sensing material exposed to CO2 and/or volatile organic compounds (TVOCs) • Applications: Indoor air quality monitoring, HVAC systems, smart homes	CO2 = 400ppm to 8192ppm TVOC = 0ppb to 1187ppb
Dust Sensor	Type: Optical sensor • Working principle: Measures the light scattering or transmission	Output Voltage range (Min 3.4 Volts) Operating temperature range -10 to +65 (Celsius)

Footnote may be given with Font size 10 of Times New Roman.

	caused by dust particles using a LED or laser beam • Applications: Environmental monitoring, indoor air quality control, industrial process control	Output Voltage at no dust (0 to 1.5 Volts)
Air Quality Co, NH_3, NO Gas Sensor	Type: Electrochemical sensor • Working principle: Measures the change in electrical current generated by a reaction between the target gas and a sensing electrode • Applications: Environmental monitoring, gas detectors, industrial process control	Carbon monoxide CO 1 – 1000ppm Nitrogen dioxide NO_2 0.05 – 10ppm Ethanol C_2H_5OH 10 – 500ppm Hydrogen H_2 1 – 1000ppm Ammonia NH_3 1 – 500ppm Methane CH_4 >1000ppm Propane C_3H_8 >1000ppm Iso-butane C_4H_{10} >1000ppm
Current Transformer	Type: Current sensor • Working principle: Measures the magnetic field produced by an electric current passing through a wire using a magnetic core and a coil • Applications: Power monitoring, energy management, industrial control systems	Instantaneous Energy for Phase 1, 2 & 3: 0 to 10s kJ Power Factor: 0 to 1 Frequency of Energy Line: 50/60 Hz Current Amperes: 0 to 1000

3.2 Data Pre-Processing:

Data pre-processing plays a crucial role in the overall performance of a machine learning model. Several steps were taken to ensure the data was appropriately prepared before training and evaluating the predictive models.

Initially, the data from all sensors and devices were combined into a single dataset. This integration involved aligning the timestamps of the data points to create a unified timeline. By concatenating the data, a comprehensive table was formed with 40,320 rows representing the recorded data instances and 78 columns representing the various features or parameters.

To ensure consistent scaling and to eliminate any potential bias caused by varying measurement units or scales, the dataset was normalized to a range of [0, 1]. Normalization allows all input variables to be on the same scale, preventing certain features from dominating the learning

Footnote may be given with Font size 10 of Times New Roman.

process due to their larger magnitudes. In order to capture the temporal patterns and incorporate historical information into the predictive models, the input variables were transformed into 60-minute lag sequences. This decision was based on previous research and preliminary experiments, which indicated that a 60-minute lag sequence offered a balance between predictive accuracy and the risk of overfitting. By including past observations as input, the models could leverage temporal dependencies and capture patterns in the data that contribute to future temperature forecasting.

To further enhance the training process and introduce additional variability, a sliding window technique was applied. This technique involved dividing the time-series data into smaller sequences or windows. This division not only facilitates data augmentation but also enables the models to learn from shorter sequences within the overall time-series. In this study, sliding window sizes of 1, 15, 30, and 60 were chosen to create different forecast horizons.

After the pre-processing steps, the dataset was split into training, validation, and testing sets. The split was performed in a ratio of 60:40, with 60% of the data allocated for training and the remaining 40% for validation and testing. This division allows for model training on a substantial portion of the data while retaining separate sets for evaluating model performance and generalization. The specific choices made regarding the training set size were based on experimentation and analysis. It was observed that varying the training set size to 50% and 70% led to decreased predictive accuracy, with the R2 score dropping below the desired threshold. Thus, a training set size of 60% was determined to strike the right balance between data availability for training and the model's ability to capture the underlying patterns effectively.

3.3 Forecasting Models

Time-series forecasting is a widely studied problem, and researchers have developed various models to address it effectively. In the context of this research, a transformer model is utilized for time-series forecasting, which was originally introduced for Natural Language Processing (NLP) tasks with promising results [26]. Transformers have demonstrated their ability to capture long-term dependencies in sequences, making them well-suited for time-series analysis. To assess the performance of the proposed transformer model, it is compared with a recent model introduced by Elmaz et al. [10]. The comparison model combines Convolutional Neural Network

Footnote may be given with Font size 10 of Times New Roman.

(CNN) and Long Short-Term Memory (LSTM) network, commonly referred to as CNN-LSTM. The CNN-LSTM architecture is specifically designed for time-series forecasting, as it can effectively identify both short- and long-term dependencies in the data.

3.3.1 CNN-LSTM

Convolutional Neural Network (CNN) coupled with a Long Short-Term Memory (LSTM) network, known as CNN-LSTM, is a powerful model architecture that excels in time-series forecasting tasks. This architecture effectively captures both short and long-term dependencies in the data, making it highly suitable for addressing complex time-series forecasting problems.

In the CNN-LSTM model, a CNN is employed as a feature extractor to extract relevant features from the input data. CNNs are well-known for their ability to automatically learn and identify important spatial features in various types of data, including images, signals, and time-series data [28]. The CNN component of the model uses 1D convolutions to capture important patterns and features from the input time-series data. The extracted features from the CNN are then fed into the LSTM component of the model. LSTM, which stands for Long Short-Term Memory, is a type of recurrent neural network (RNN) that is specifically designed to address the challenges of modeling long-term dependencies in sequential data [29]. LSTMs overcome the issues of gradient disappearance and explosion that are commonly encountered in traditional RNNs. They achieve this by incorporating a memory cell that can retain information over extended time steps, enabling the network to selectively retain and discard information as needed [30]. This memory mechanism allows LSTMs to capture and remember long-term dependencies in the time-series data, contributing to their effectiveness in time-series forecasting tasks.

By combining the strengths of CNNs in feature extraction and LSTMs in modeling temporal dependencies, the CNN-LSTM architecture offers a comprehensive solution for time-series forecasting. It leverages the spatial and temporal information in the input data to make

Footnote may be given with Font size 10 of Times New Roman.

accurate predictions and capture complex patterns present in the time-series data

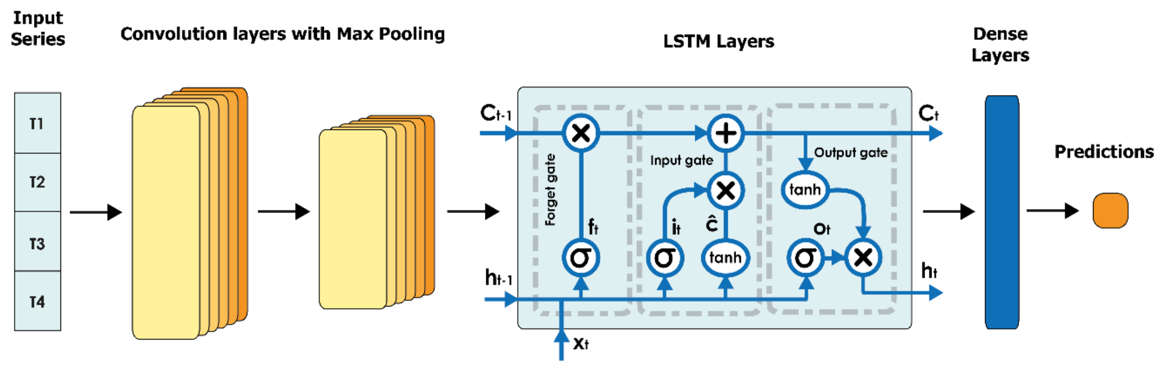

Figure 4: Schematic of a CNN- LSTM architecture illustrating the input data is first fed into several convolution layers. The output from these layers is then passed on to LSTM cells, and finally to a dense layer that generates predictions.

Elmaz et al. [10] introduced a CNN-LSTM architecture for time-series forecasting, as depicted in Figure 4. The architecture consists of several interconnected layers designed to extract features from the input series and make accurate prediction.

The input series is initially fed into the CNN component of the model. The CNN consists of two stacked layers, with 64 and 128 kernels respectively. These kernels have kernel sizes of 32 and 16, indicating the receptive field size over the input data. The use of multiple layers with different kernel sizes allows the model to capture features of varying scales and complexities. A stride of 1 is employed to scan the input series, extracting relevant features. Following the convolutional layers, max pooling is applied to reduce the dimensionality of the features and retain the most salient information.

The output from the CNN component is then passed on to the LSTM layer. The LSTM layer consists of a single layer with 60 units, which are responsible for capturing and modeling the temporal dependencies within the input series. The LSTM layer utilizes its memory cell and gates to selectively remember and forget information, effectively capturing long-term dependencies in the data.

Finally, the flattened and fully connected layer, connected to the LSTM layer, contributes to the prediction process. The flattened layer transforms the output from the LSTM layer into a

Footnote may be given with Font size 10 of Times New Roman.

one-dimensional vector, which is then fed into the fully connected layer. This layer performs computations on the extracted features, further refining them for prediction purposes.

During the training process, various hyperparameters are tuned to optimize the performance of the CNN-LSTM architecture. These hyperparameters include the number of layers and units in the CNN and LSTM components, the kernel sizes, and the stride in the CNN, among others. The specific values chosen for these hyperparameters are detailed in Table 2 of the paper.

Table 2 List of Hyperparameters

Learning Rate	0.001
L2 Regularization Rate	0.1
Dropout	0.25
Batch Size	32

The batch size of 32 was selected as a common practice in deep learning to strike a balance between convergence speed and memory requirements [31]. This choice ensures that a sufficient amount of data is processed in each iteration while considering the limitations of the hardware's memory capacity. To prevent overfitting, a regularization rate of 0.1 was applied. Regularization is a technique that adds a penalty term to the loss function, discouraging the model from relying too heavily on any particular set of features. By regularizing the model, it becomes more robust and less prone to overfitting the training data. A dropout rate of 0.25 was also incorporated in the architecture. Dropout is a regularization technique where randomly selected neurons are ignored during the training process. This helps in preventing overfitting by forcing the model to learn more robust and generalized representations [31].

The learning rate was set to 0.001. The learning rate determines the step size taken during the optimization process. A carefully chosen learning rate ensures a balance between the convergence speed and the risk of overfitting or underfitting the data. A smaller learning rate allows for more precise updates but may lead to slower convergence, while a larger learning rate may cause the model to overshoot the optimal solution.

Footnote may be given with Font size 10 of Times New Roman.

3.3.2 Transformer Neural Networks

That's correct. The Transformer model is a type of sequence-to-sequence model, commonly used in natural language processing tasks such as machine translation, text summarization, and language generation. It excels in handling sequential data by employing an attention mechanism that allows it to capture dependencies across the entire input sequence. The Transformer model consists of an encoder and a decoder. The encoder processes the input sequence and generates a representation that encapsulates the input's contextual information. The decoder then uses this representation to generate the output sequence, one element at a time. This sequence-to-sequence architecture makes the Transformer model versatile and applicable to various tasks that involve encoding and decoding sequential data, not limited to language-related tasks.

The architecture of the transformer model, as depicted in Figure 5, is employed for HVAC temperature forecasting. The input sequence is first passed through a position encoding layer, which aids the transformer in capturing temporal patterns and dependencies within the data [32]. The positional encoder enhances the model's ability to understand the sequential nature of the input sequence.

The output from the positional encoding layer is then fed into multiple transformer blocks, with our architecture incorporating four such blocks. Each transformer block consists of eight transformer heads, with each head having a size of 256. These transformer heads are utilized in multi-head attention layers, allowing the model to attend to different parts of the input sequence simultaneously and capture relevant information effectively.

Footnote may be given with Font size 10 of Times New Roman.

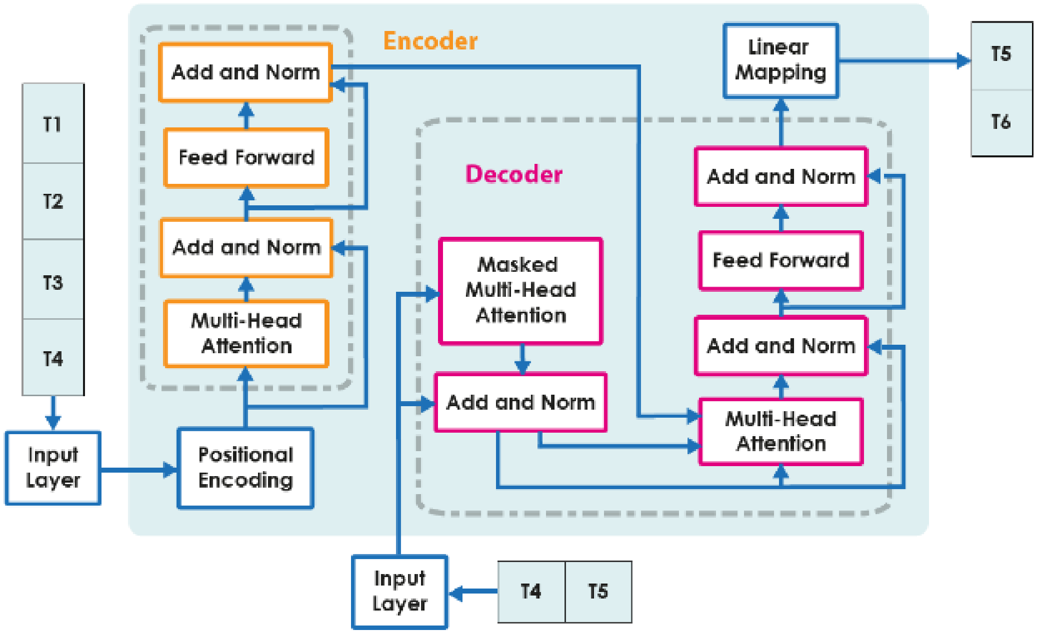

Figure 5: Architecture diagram of transformer neural network representing the flow of input sequences through multiple layers.

Furthermore, the architecture incorporates two 1D convolutional layers in the feed-forward layer. These convolutional layers have kernel sizes of 1 and filter sizes of 8 and 74, respectively, corresponding to the number of features present in the dataset. The convolutional layers serve to further process and extract meaningful representations from the input data. After passing through the multiple transformer blocks and the convolutional layers, the output data is directed to a dense layer with 256 units. This dense layer performs additional computations and generates the final output, which represents the predicted HVAC temperature.

The transformer architecture's utilization of self-attention mechanisms is a key feature that sets it apart. Self-attention allows the network to assign relative importance to different elements within the input sequence when making predictions [33]. This dynamic weighting mechanism enables the transformer to effectively process input sequences of varying lengths. Unlike traditional Recurrent Neural Networks (RNNs), which have a fixed-length context, the transformer can capture long-range dependencies in the data, making it suitable for tasks that involve extensive temporal or sequential information [34].

Footnote may be given with Font size 10 of Times New Roman.

Another notable aspect of the transformer is its employment of multi-head attention. With multi-head attention, the model can simultaneously process information from multiple representation subspaces and different positions within the input sequence [33]. By considering different subspaces, the transformer can capture diverse patterns and relationships within the data, enhancing its ability to understand complex and multi-faceted dependencies. This multi-head attention mechanism allows the model to efficiently handle both local and global information, resulting in more comprehensive and accurate predictions. The self-attention mechanism in the transformer can be described mathematically as follows:

$$\textbf{Attention (Q, K, V) = softmax} \left(\frac{QK^T}{d\_k}\right)\textbf{V} \qquad 1$$

In Equation 1, the parameter d_k represents the dimensionality of the key representations in the self-attention mechanism. The matrices Q, K, and V are learned during the training process and are responsible for mapping the input elements to their respective query, key, and value representations. These representations play a crucial role in the self-attention process.

During self-attention, the model calculates the similarity between each query and its corresponding key representation. This similarity score determines the importance or weight assigned to each value representation. The key representations capture information about the relationships and dependencies within the input sequence, while the query representations indicate the elements of interest that require attention. The value representations contain the actual information or features associated with each input element. By computing the dot product between the query and key representations and scaling the result by the square root of the dimensionality d_k, the self-attention mechanism produces attention weights. These weights reflect the relevance of each value representation to the corresponding query. The value representations are then multiplied by their respective attention weights and summed to obtain a weighted sum, representing the aggregated information or context relevant to the query.

The self-attention mechanism allows the model to selectively attend to different parts of the input sequence, focusing on relevant elements and capturing important dependencies. By dynamically weighing the value representations based on their similarity to the query, the model can effectively capture the relevant information needed for accurate predictions. This process of attending and aggregating information is performed across all positions in the input sequence,

Footnote may be given with Font size 10 of Times New Roman.

enabling the transformer to capture global dependencies and effectively model long-term relationships in the data [33].

The flexibility and adaptability of the transformer architecture are particularly advantageous when dealing with time-series data. Its self-attention mechanisms enable it to effectively model long-term dependencies, while the multi-head attention mechanism allows it to capture diverse patterns and relationships within the data. These capabilities make the transformer well-suited for tasks involving temporal sequences, such as time-series forecasting. By leveraging these characteristics, the transformer architecture has demonstrated remarkable success in various domains, including natural language processing and now extending to time-series forecasting. The pseudo code for complete architecture of transformer is presented below:

Algorithm 2 Pseudocode of Transformers Neural Networks

```
FUNCTION transformer_model (inputs, head_size, num_heads, dropout, ff_dim):
    x ← LayerNormalization (inputs)
    x ← MultiHeadAttention (key_dim=head_size, num_heads=num_heads, dropout=dropout) (x,x)
    x ← Dropout (dropout)(x)
    res ← x + inputs
    x ← LayerNormalization (res)
    x ← Conv1D (filters=ff_dim, kernel_size=1, activation="relu")(x)
    x ← Dropout (dropout)(x)
    x ← Conv1D (filters=inputs.shape[-1], kernel_size=1)(x)
    RETURN x + res

FUNCTION build (look_back, n_features, num_transformer_blocks, head_size, num_heads, dropout,
ff_dim, mlp_units, mlp_dropout, horizon):
    inputs ← Input (shape = (look_back, n_features))
    x ← inputs
    for i in range(num_transformer_blocks):
        x ← transformer_model(x, head_size, num_heads, dropout, ff_dim)
    x ← GlobalAveragePooling1D(x)
    for dim in mlp_units:
        x ← Dense (dim, activation="relu")(x)
        x ← Dropout (mlp_dropout)(x)
    outputs = Dense(x)
    RETURN Model (inputs, outputs)
```


Footnote may be given with Font size 10 of Times New Roman.

3.4 Evaluation Metrics

To assess the efficiency of the implemented temperature forecasting models, three commonly utilized error metrics, namely Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-Squared (R^2) are used.

The Mean Absolute Error (MAE) is a commonly used metric for evaluating the performance of regression models, including those used for temperature forecasting. It quantifies the average magnitude of the errors between the actual and predicted temperature values. MAE is calculated by taking the absolute difference between the actual temperature and the predicted temperature for each data point, and then averaging these absolute differences over the entire dataset. By considering the absolute differences, MAE captures the magnitude of the errors without considering their direction, making it a reliable metric for assessing the overall accuracy of the predictions. One advantage of using MAE is that it treats all errors equally, without giving additional emphasis to extreme errors [35]. This characteristic is particularly useful in situations where extreme errors may occur sporadically but do not significantly impact the overall model performance. MAE provides a straightforward measure of the average prediction error, allowing for a direct comparison of different models or techniques. This makes it a reliable metric for comparing the performance of different models as shown in equation 2.

$$\boldsymbol{MAE = \frac{1}{n}\sum_{i=1}^{n}|y_{predicted} - y_{actual}|} \qquad \textbf{2}$$

where n denotes the number of data points in the dataset, $y_{predicted}$ represents the predicted values and y_{actual} represents the actual values.

The Root Mean Squared Error (RMSE) is a commonly used metric for assessing the performance of regression models, including those used for temperature forecasting. It measures the square root of the average of the squared differences between the actual and predicted temperature values. RMSE is calculated by taking the squared difference between the actual temperature and the predicted temperature for each data point, summing these squared differences, dividing by the total number of data points, and then taking the square root of the result. By squaring the differences, RMSE gives additional weight to larger errors, including outliers, and provides a measure of the overall dispersion or spread of the prediction errors [36]. One advantage

Footnote may be given with Font size 10 of Times New Roman.

of using RMSE is its sensitivity to extreme values or outliers in the temperature predictions. Since the error term is squared in the calculation, RMSE amplifies the impact of these extreme values, allowing for their identification and evaluation. This can be particularly useful in scenarios where outliers have significant consequences or need to be carefully considered in the assessment of model performance. RMSE is determined using the following formula as shown in equation 3:

$$\boldsymbol{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_{predicted} - y_{actual})^2} \qquad 3$$

The coefficient of determination (R^2) is a statistical metric used to assess the quality of a regression model, including those used for temperature forecasting. It measures the proportion of the variance in the target variable (actual temperature values) that can be explained by the model's predictions. R^2 is calculated by comparing the sum of squares of the differences between the actual temperature values and the mean of the actual values with the sum of squares of the differences between the predicted temperature values and the mean of the actual values This is calculated using the following formula as shown in equation 4:

$$\boldsymbol{R^2} = 1 - \frac{\sum_{i=1}^{n}(y_a - y_p)^2}{\sum_{i=1}^{n}(y_a - \overline{y_a})^2} \qquad 4$$

R^2 is a useful metric for evaluating the goodness of fit between the actual and predicted temperature values [37]. A higher R^2 score indicates a better fit between the model's predictions and the actual temperature values. It signifies that a larger proportion of the variance in the target variable can be accounted for by the model. In other words, a high R^2 score suggests that the model's predictions closely match the actual values, indicating a good level of accuracy. R^2 is a valuable metric because it provides a standardized way to compare the performance of different models. Since R^2 ranges from 0 to 1, with 1 indicating a perfect fit, it allows for easy interpretation and comparison of models. Models with higher R^2 scores are generally considered to have better predictive performance, as they explain a larger proportion of the variance in the target variable.

These error metrics provided a comprehensive evaluation of the models' performance and helped in determining the best-performing model for temperature forecasting.

CHAPTER 4: MATERIALS AND METHODS

4.1 Subjects

The objective of the study was to investigate thermal comfort levels in relation to HVAC temperature forecasting. The research involved conducting experiments in a controlled room environment with 4-6 participants. The participants selected for the study were between the ages of 24 to 30 years and had a weight range of 55 to 80 kg.

Before the study began, the researchers obtained informed consent from each participant, ensuring that they were aware of the purpose and procedures of the study. Confidentiality of the participants' data was maintained throughout the research process, ensuring that their personal information and experimental results were kept private and protected. By conducting experiments with human participants, the study aimed to gather data and insights regarding thermal comfort in indoor environments. The participants' responses and feedback, along with the temperature forecasting models, were used to assess and analyse the relationship between HVAC temperature and thermal comfort.

4.2 Environmental Setup

The experiments in this study were conducted in a specifically selected room that provided a controlled indoor environment, allowing to have precise control over the temperature while closely monitoring other parameters. To accurately measure and analyse the various parameters, data collection devices were strategically placed in different locations within the room, as depicted in Figure 7. For instance, the power quality analyser (labelled as "3" in the figure) was connected to the power source of the HVAC system, allowing for the measurement of electrical parameters such as voltage, current, and power consumption. Similarly, the air quality analyser (labelled as "5" in the figure) was placed on the workstation, enabling the measurement of various air pollution parameters including carbon dioxide, particulate matter, volatile organic compounds (VOCs), and ozone. This placement ensured that the air quality measurements were representative of the immediate vicinity where the occupants were located.

Footnote may be given with Font size 10 of Times New Roman.

The experiments were specifically conducted during the winter season when the weather conditions were characterized by low temperatures and low humidity levels. This choice of season allowed the researchers to evaluate the performance of the HVAC system in maintaining thermal comfort under challenging environmental conditions. The trend of temperature during the testing phase is shown in Figure **6**6.

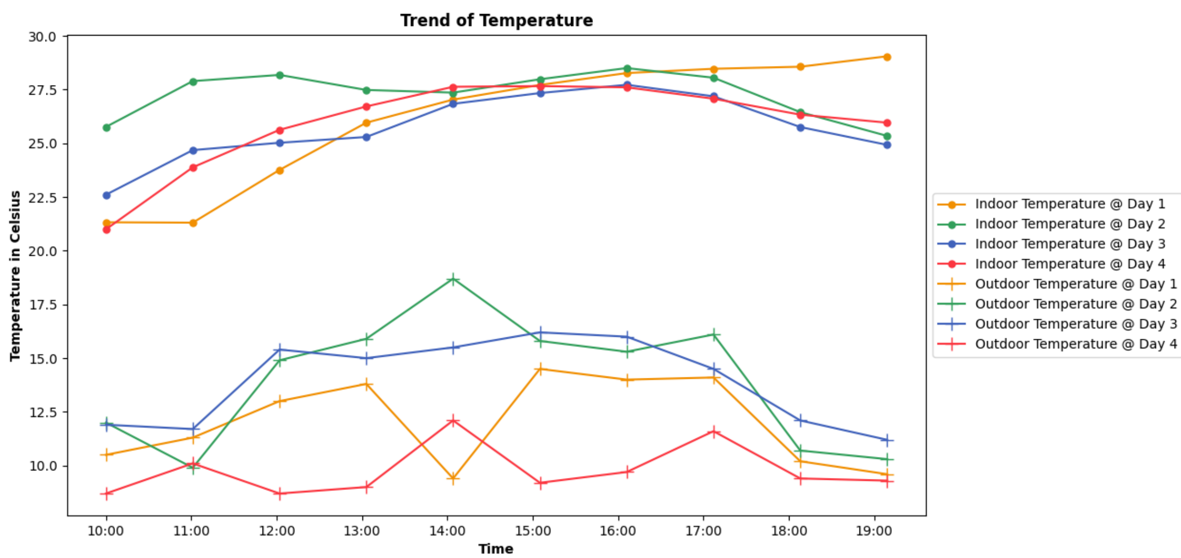

Figure 6: Trend of indoor and outdoor temperatures

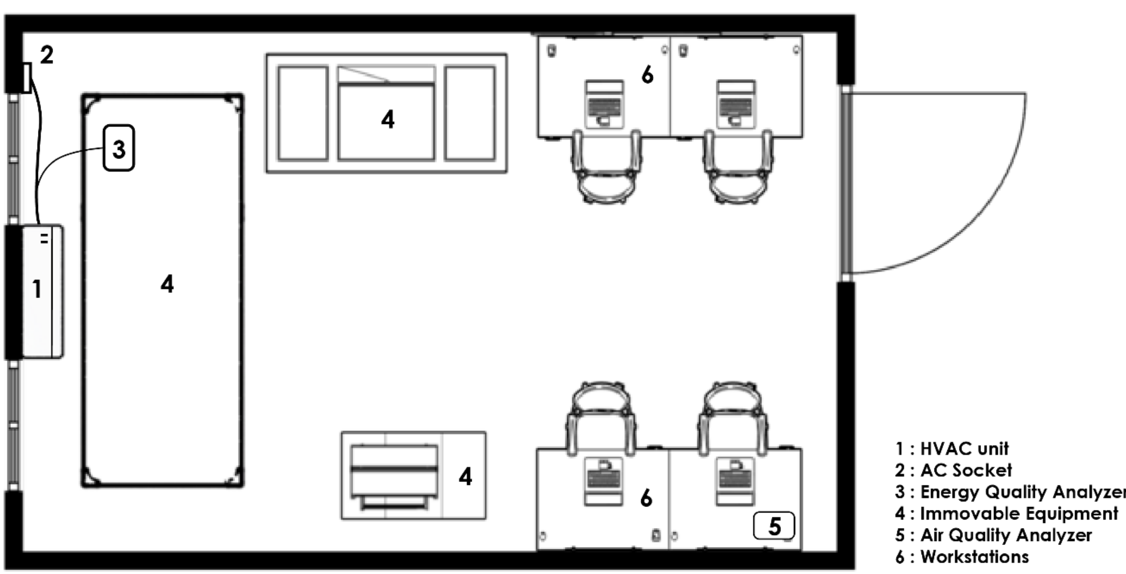

Figure 7: Layout of the experimentation room

Footnote may be given with Font size 10 of Times New Roman.

4.3 Experiment Design

In this experiment, a comprehensive dataset was created by collecting data from various sources, including multiple devices and surveys. The goal was to gather a wide range of information related to thermal comfort and HVAC temperature forecasting. To assess thermal comfort, efficient surveys were designed based on the ASHRAE standard 55, which provides guidelines for indoor thermal environmental conditions and human occupancy comfort. The surveys covered personal parameters such as metabolic rate, clothing insulation, body mass index (BMI), and body temperature. The ASHRAE standard also provided a method to estimate the metabolic rate and clothing insulation based on the occupants' activity level and clothing type. To measure body temperature, a laser thermometer was used, which is a non-invasive tool commonly employed to measure surface temperatures. Additionally, the thermal perception of the occupants was recorded using a scale based on the ASHRAE standard 55. This 5-point scale ranged from -2 to 2, corresponding to the categories cool, slightly cool, neutral, slightly warm, and warm. Meteorological data was obtained from NOAA's weather database. This data consisted of 13 parameters, including outdoor temperature, UV index, humidity, and other relevant meteorological factors. Air quality data was recorded using an air quality analyser, which measured multiple parameters contributing to air pollution. This helped evaluate the quality of the indoor air and its potential effects on occupant well-being. The analyser measured 15 parameters, including carbon dioxide, particulate matter, volatile organic compounds (VOCs), and ozone. Power quality monitoring was carried out using a power quality analyser, which measured various electrical parameters related to the HVAC system, such as voltage, current, active power, and reactive power. The indoor environment ambiance was assessed by considering variables such as lighting levels, indoor temperature, humidity, pressure, and occupancy status.

Once the data was collected from all these sources, the initial phase of data pre-processing involved combining the data from all sensors and devices into a single dataset. This was achieved by aligning the timestamps of the data points and concatenating them into a single table. The columns of the dataset represented the features/parameters, while the rows represented the recorded data points. To ensure that all input variables were on the same scale, the dataset was normalized to the range of [0, 1]. To enable the prediction of HVAC temperature, the input

Footnote may be given with Font size 10 of Times New Roman.

variables were converted into 60-minute lag sequences based on literature and preliminary experiments. This lag sequence approach allowed for the inclusion of past observations in the prediction, which is a common technique used in time-series forecasting. Additionally, a sliding window technique was employed to divide the time-series into smaller sequences, serving as a method of data augmentation. The forecasting sequence sizes were selected as 1, 15, 30, and 60, representing the desired forecasting horizons. Shorter forecasting horizons such as 1 or 15 minutes provide real-time information about temperature changes in the HVAC system, while longer forecasting horizons such as 30 or 60 minutes offer a more comprehensive view of temperature trends over an extended period.

The data was then split into training, validation, and testing sets in a ratio of 60:40. This allowed the models to be trained and evaluated using a significant portion of the data while reserving a portion for independent testing. The training set size was chosen to be 60% based on experimentation, which indicated that varying the training set size between 50% and 70% did not significantly affect model performance.

The transformer neural network architecture, as shown in Figure 5, was utilized for temperature forecasting. The input sequence was passed through a position encoding layer, which helped the transformer model learn temporal patterns and dependencies in the data. The output of the positional encoder was then directed to multiple transformer blocks, with the architecture incorporating four such blocks. Each transformer block consisted of eight transformer heads, each with a size of 256, which were utilized in multi-head attention layers. Following the attention mechanism, two 1D convolution layers were implemented in the feed-forward layer, with kernel sizes of 1 and filter sizes of 8 and 74 (corresponding to the count of features present in the dataset), respectively. The output data from the multiple transformer blocks was sent to a dense layer with 256 units, which generated the final output.

In addition to the transformer model, the CNN-LSTM architecture proposed by Elmaz et al. [10] was also employed for temperature forecasting. This architecture combines a Convolutional Neural Network (CNN) and a Long Short-Term Memory (LSTM) network, making it particularly suitable for capturing both short and long-term dependencies in the data. The CNN component was used to extract features from the input data, while the LSTM component processed the data over time. This combination allowed the model to effectively handle both spatial and temporal information in the input data.

During the training phase, various hyperparameters were set to optimize the performance of the models. The batch size was chosen as 32, following the common practice of balancing convergence speed and memory requirements. A regularization rate of 0.1 and dropout of 0.25 were applied to prevent overfitting. A learning rate of 0.001 was selected to balance the training process and prevent the models from overfitting or underfitting.

To evaluate the performance of the models, three key metrics were used: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R2). MAE is a useful metric as it directly measures the average magnitude of the error, without giving additional emphasis to extreme errors. RMSE, on the other hand, takes into account the squared error terms, thereby giving additional weight to extreme values. R2 measures the proportion of the variance in the target variable that can be accounted for by the model's predictions. A high R2 score indicates a close agreement between the model's predictions and the actual values.

The trained models were then tested in real-time on actual HVAC systems. The test data, including the input variables and the corresponding actual temperature values, was provided to the selected models to generate predictions. These predictions were then fed into a microcontroller-based Infrared (IR) module, which controlled the HVAC system based on the predicted temperatures. To assess the effectiveness of the AI models in managing the HVAC system and predicting energy usage, the energy consumption data was logged using a power quality analyser over a period of four days. The data obtained through the AI models was compared with the energy consumption data obtained through a traditional, non-AI method of HVAC control. This comparison allowed for the evaluation of the AI models' ability to optimize energy consumption and maintain thermal comfort.

Footnote may be given with Font size 10 of Times New Roman.

Figure 8 IR-Controller

Furthermore, a survey was conducted to obtain feedback from the occupants regarding their thermal comfort levels and to assess how well the HVAC system, controlled by the proposed architecture, was able to maintain a comfortable indoor environment. The survey included questions about occupant satisfaction, perceived thermal comfort, and any discomfort experienced during the experiment. This feedback was used to further validate the performance of the transformer models in managing the HVAC system and predicting energy usage data. Figure 9 illustrates the complete flow of the experimental procedure.

Footnote may be given with Font size 10 of Times New Roman.

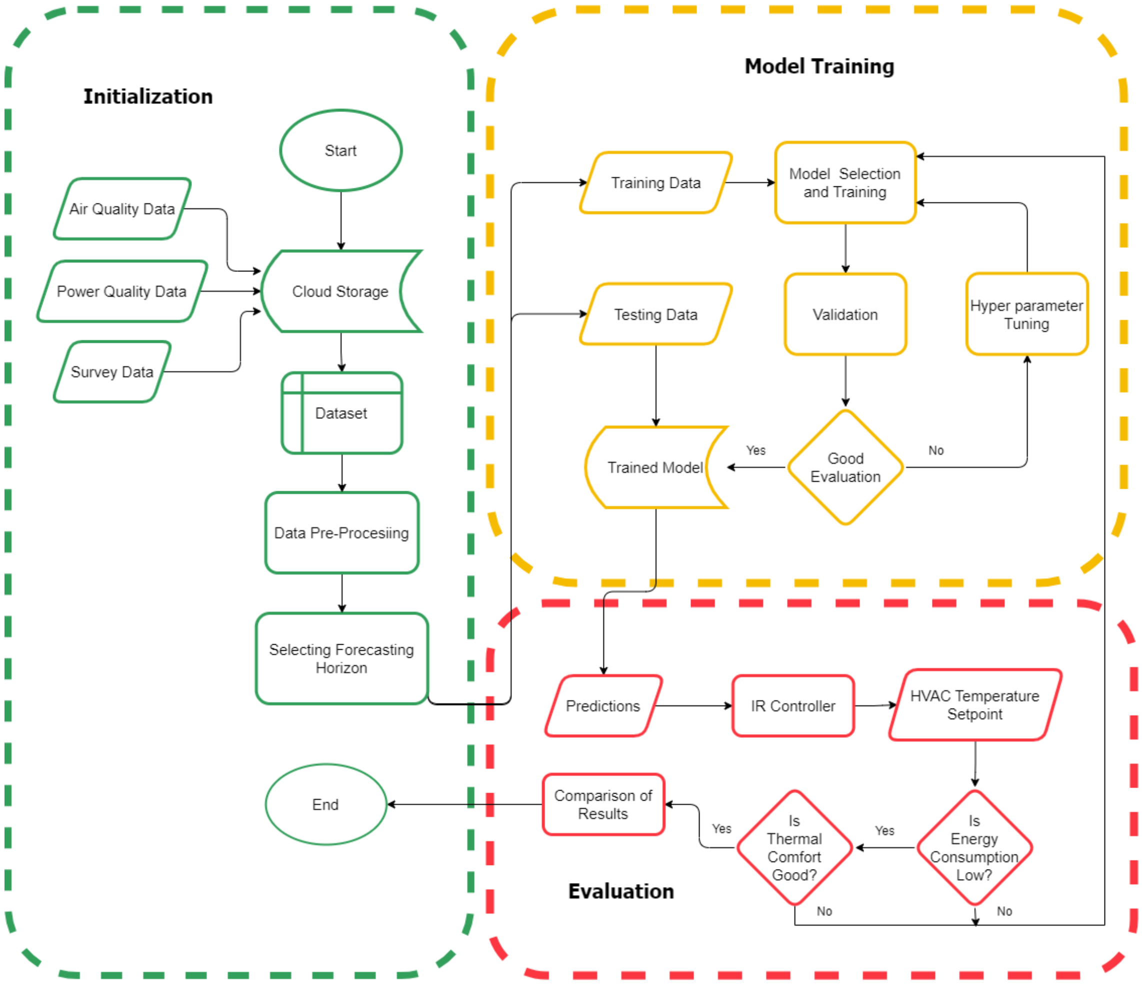

Figure 9: Flow diagram of forecasting temperature on HVAC

Footnote may be given with Font size 10 of Times New Roman.

Algorithm 2 Experimental Design

```
data_sources = ["devices", "surveys"]
data = collect_data(data_sources)
clean_data = preprocess_data(data)
forecasting_horizons = [1, 15, 30, 60]
best_model = None
best_score = None

For horizon in forecasting_horizons:
    training_set, validation_set, testing_set = split_dataset(clean_data)
    model = train_model (training_set, horizon)
    mae, rmse, R2 = evaluate_model (model, testing_set)
    if best_score is None or R2 < best_score:
        best_score = R2
        best_model = model
microcontroller_module = setup_microcontroller()

For horizon in forecasting_horizons:
    test_data = run_test(best_model, horizon)
    energy_consumption = log_energy_consumption(test_data)
    feed_predictions(microcontroller_module, test_data)

traditional_energy_data = get_energy_data("non-AI")
ai_energy_data = get_energy_data("AI")
compare_energy_consumption(traditional_energy_data, ai_energy_data)

survey_data = conduct_survey()
validate_model_performance(best_model, survey_data)
```


Footnote may be given with Font size 10 of Times New Roman.

CHAPTER 5: RESULTS AND DISCUSSION

The present study evaluates the effectiveness of a deep learning model in predicting HVAC setting to optimize energy efficiency while maintaining thermal comfort in buildings by utilizing dataset obtained from sensors. The machine learning algorithms were trained for this purpose using various time step predictions on the collected dataset for assessing the accuracy of the model's predictions. The energy consumption data collected from the resulting prediction outcomes were compared with the actual energy consumption to assess the usefulness of machine learning in energy optimization. The outcomes of the comparison were then used to determine the extent to which the deep learning model is effective in predicting HVAC temperature and optimizing energy consumption in HVAC systems supported by survey data used to assess the occupants' thermal comfort

5.1 Model Evaluation

The evaluation of the deep learning models, CNN-LSTM and transformers involved the use of different evaluation metrics to assess their performance in forecasting HVAC temperatures. The results of the study indicated that both models were effective in predicting HVAC temperatures with varying forecasting horizons.

Table 3 presents the Mean Absolute Error (MAE) results obtained from the evaluation of the models. The MAE metric measures the average absolute difference between the actual and predicted temperature values. In this study, the MAE values were compared across the same forecasting horizons to determine the accuracy of the models. The results showed that the transformers model outperformed the CNN-LSTM model in terms of MAE values. The transformers model exhibited a lower range of MAE values, ranging from 0.02 to 0.06 across the different forecasting horizons. On the other hand, the CNN-LSTM model had a higher range of MAE values, ranging from 0.06 at a horizon of 1 minute to 0.10 at a horizon of 60 minutes.

Footnote may be given with Font size 10 of Times New Roman.

Table 3 Comparison of performance between CNN-LSTM and Transformers at multiple forecasting horizons

Evaluation Metrics	Forecasting Horizon	CNN-LSTM	Transformers
MAE	1 minute	0.06	0.02
	15 minutes	0.08	0.03
	30 minutes	0.09	0.06
	60 minutes	0.10	0.05
RMSE	1 minute	0.06	0.03
	15 minutes	0.09	0.04
	30 minutes	0.09	0.07
	60 minutes	0.10	0.05
R^2	1 minute	0.736	0.936
	15 minutes	0.541	0.873
	30 minutes	0.484	0.710
	60 minutes	0.330	0.788

In addition to the Mean Absolute Error (MAE), the Root Mean-Square Error (RMSE) values were also calculated to evaluate the performance of the CNN-LSTM and transformers models in forecasting HVAC temperatures. The RMSE measures the square root of the average squared difference between the actual and predicted temperature values.

For the CNN-LSTM model, the RMSE values ranged from 0.06 to 0.10 across the different forecasting horizons. On the other hand, the transformer model exhibited RMSE values ranging from 0.03 to 0.07. These results indicate that the transformer model outperformed the CNN-LSTM model in terms of RMSE, showing higher accuracy in forecasting HVAC temperatures. The lower RMSE values for the transformer model suggest that its predictions were closer to the actual temperature values compared to the CNN-LSTM model. This implies that the transformer model

Footnote may be given with Font size 10 of Times New Roman.

had a smaller overall deviation from the true values, indicating a higher level of accuracy in its forecasts.

However, it is important to note that both models experienced an increase in RMSE as the forecasting horizon increased. This suggests that the accuracy of the models' predictions decreased with longer forecasting horizons. The rise in RMSE can be attributed to the increased complexity of the HVAC system's thermal behavior over extended time periods. As the forecasting horizon expands, there may be more factors and external influences that can affect temperature fluctuations, making it more challenging for the models to accurately predict the temperatures.

These findings highlight the trade-off between forecasting accuracy and the duration of the forecasting horizon. While the transformer model demonstrated superior performance in terms of lower RMSE values, it is important to consider the limitations in accuracy when forecasting over longer time intervals. It may be necessary to strike a balance between the desired forecasting horizon and the acceptable level of accuracy based on the specific requirements and constraints of the HVAC system.

The coefficient of determination (R2) scores were used to assess the predictive power and goodness of fit of the transformers and CNN-LSTM models in forecasting HVAC temperatures. The R2 score measures the proportion of the variance in the target variable that can be explained by the model's predictions. The results of the study indicate that the transformers model outperformed the CNN-LSTM model in terms of R2 scores across all forecasting horizons. For the transformers model, the R2 scores ranged from 0.936 at a horizon of 1 minute to 0.788 at a horizon of 60 minutes. On the other hand, the CNN-LSTM model had R2 scores ranging from 0.736 to 0.330 across the same horizons. The higher R2 scores obtained by the transformers model suggest that it is better able to explain the variation in the HVAC temperature data compared to the CNN-LSTM model. This indicates that the transformers model captures a larger proportion of the underlying patterns and relationships in the data, resulting in more accurate predictions.

Footnote may be given with Font size 10 of Times New Roman.

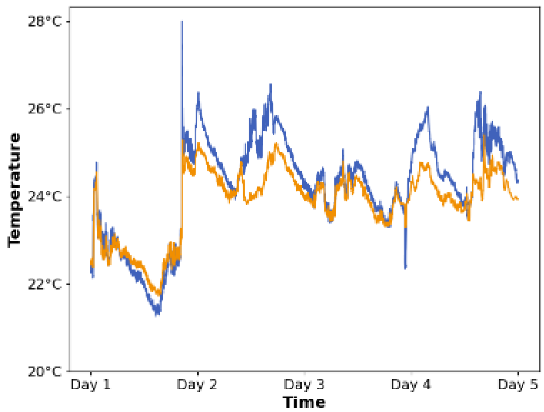

(a) CNN-LSTM at 1-minute forecasting horizon

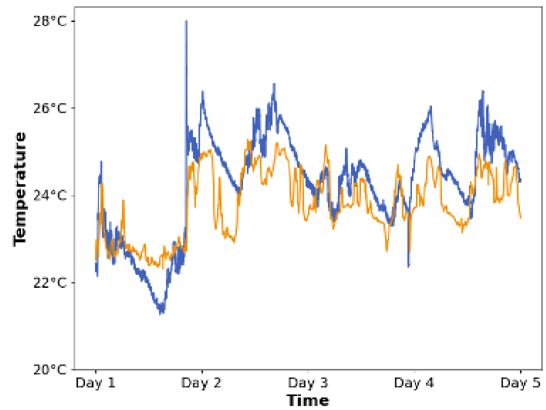

(b) CNN-LSTM at 15-minute forecasting horizon

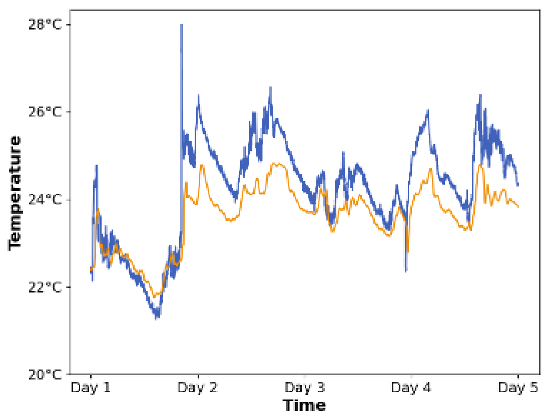

(c) CNN-LSTM at 30-minute forecasting horizon

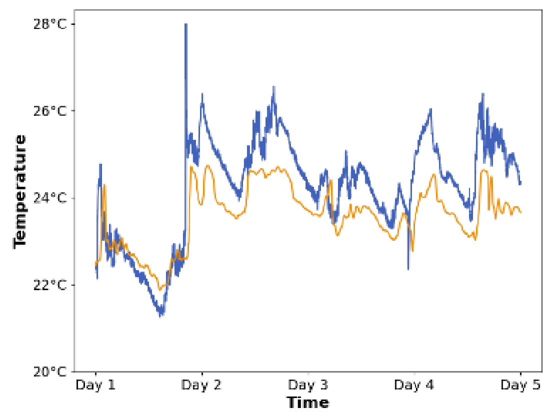

(d) CNN-LSTM at 60-minute forecasting horizon

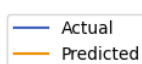

Figure 10: Actual vs Predicted temperatures using CNN-LSTM with temperature on y-axis and Total no. of predictions in terms of days are represented on x-axis.

However, it is important to note that both models experienced a decrease in R2 scores as the forecasting horizon increased. This implies that the models' ability to accurately predict HVAC temperatures diminishes over longer forecasting horizons. The decline in R2 scores can be attributed to the increased complexity and uncertainty in the HVAC system's behavior over extended time periods.

Footnote may be given with Font size 10 of Times New Roman.

These findings emphasize the trade-off between forecasting accuracy and the duration of the forecasting horizon. While the transformers model demonstrated superior performance in terms of higher R2 scores, indicating better explanatory power, it is important to consider the limitations of longer forecasting horizons. As the time interval being forecasted increases, the models may struggle to capture all the intricate dynamics and external influences affecting temperature variations, resulting in a decrease in prediction accuracy.

Figure **10**10 and Figure 11 demonstrate a notable distinction in the temperature values predicted by the CNN-LSTM and Transformer models as compared to their actual values. Figure 10 contains four subplots that illustrate the performance of the CNN-LSTM model in predicting temperature values. The blue line in each subplot represents the actual temperature values, while the yellow line represents the predicted values. Among all the tested forecasting horizons, the subplot for 1-minute forecasting horizon in the CNN-LSTM model shows the closest alignment between the predicted temperature values and the actual temperature values. Figure 11 comprises four subplots that depict the performance of Transformers in making predictions at different forecasting horizons. The subplot for the 1-minute forecasting horizon in Figure 11 provides evidence of the superior performance of Transformers in predicting one step ahead as it almost overlaps the actual values. Subplots (b), (c), and (d) of the same figure show that the performance of Transformers at forecasting horizons of 15, 30, and 60 minutes was also significantly better than that of the CNN-LSTM model. The figures suggest that the transformer model produced more accurate predictions, as the plot of the actual and predicted values almost entirely overlap at forecasting horizon of 1.

Footnote may be given with Font size 10 of Times New Roman.

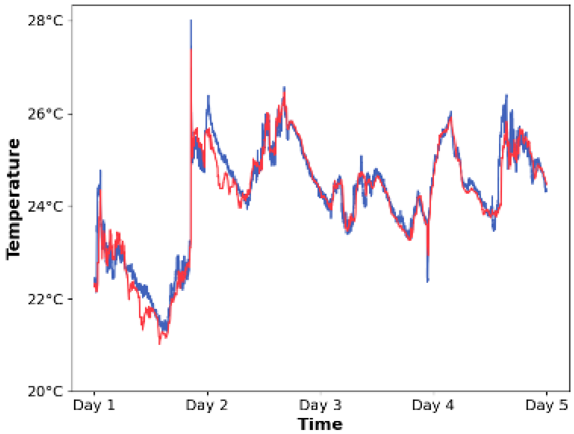

(a) Transformers at 1- minute forecasting horizon

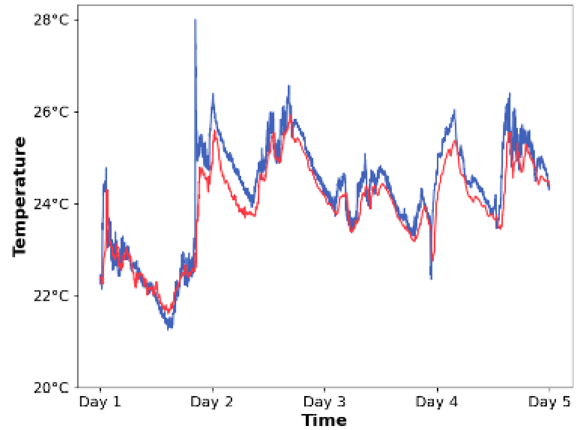

(b) Transformers at 15-minute forecasting horizon

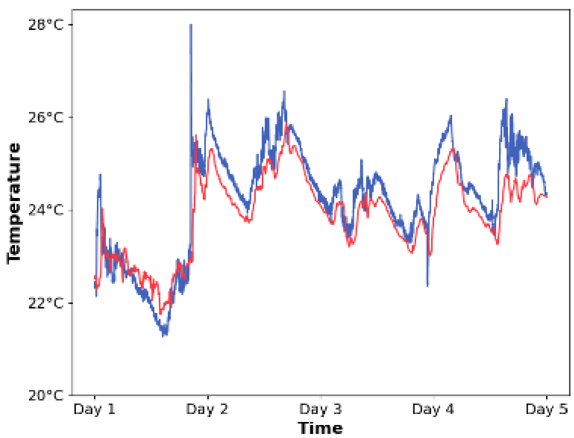

(c) Transformers at 30-minute forecasting horizon

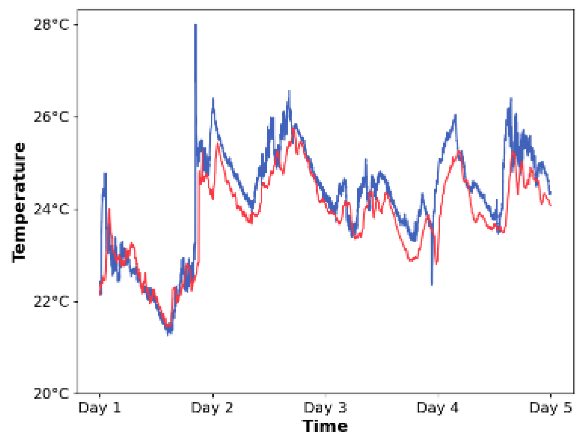

(d) Transformers at 60-minute forecasting horizon

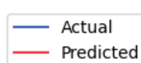

Figure 11: Actual vs Predicted Temperatures using Transformers with Temperature on y-axis and Total no. of predictions in terms of days are represented on x-axis.

5.2 Energy Optimization

The primary objective of the study was to investigate the potential of deep learning models in optimizing energy consumption in HVAC systems. The results, as illustrated in Figure 12,

Footnote may be given with Font size 10 of Times New Roman.

highlight significant differences in energy consumption patterns between different time intervals and the utilization of deep learning algorithms compared to traditional non-AI methods.

The dark magenta bars in Figure 12 represent the energy consumption under the traditional non-AI method. Comparing these bars with the yellow and green bars representing the utilization of Transformer models for 1-minute and 15-minute prediction horizons, respectively, reveals inconsistent energy consumption behaviour at different time intervals. This suggests that the deep learning models have the potential to optimize energy consumption by making more accurate predictions and adapting the HVAC system's settings accordingly.

Examining the blue line representing the Transformers' energy consumption at a 30-minute forecasting horizon, we observe periodic declines in energy consumption. This implies that the Transformers model successfully adjusts the HVAC system's settings to minimize energy consumption during specific intervals. Furthermore, the red bars representing the Transformers model at a 60-minute forecasting horizon demonstrate consistent behaviour, with energy consumption remaining stable throughout the morning and afternoon and slightly decreasing during the evening. This suggests that the Transformers model is able to optimize energy consumption over longer time intervals.

Overall, the findings indicate that deep learning models, particularly the Transformers model, can effectively optimize energy consumption in HVAC systems. By leveraging the power of deep learning algorithms, these models are capable of making accurate predictions and dynamically adjusting the HVAC system's settings to achieve energy efficiency. This has the potential to significantly reduce energy consumption and improve the overall sustainability of HVAC systems.

Figure 13 clearly illustrates the significant impact of different forecasting horizons on net power consumption in the evaluated models. The transformer model with a 1-minute horizon stands out with the lowest net power consumption of 8.19 kWh, indicated by the yellow line. In contrast, the non-AI method, represented by the dark magenta line, exhibits the highest net power consumption at 16.54 kWh.

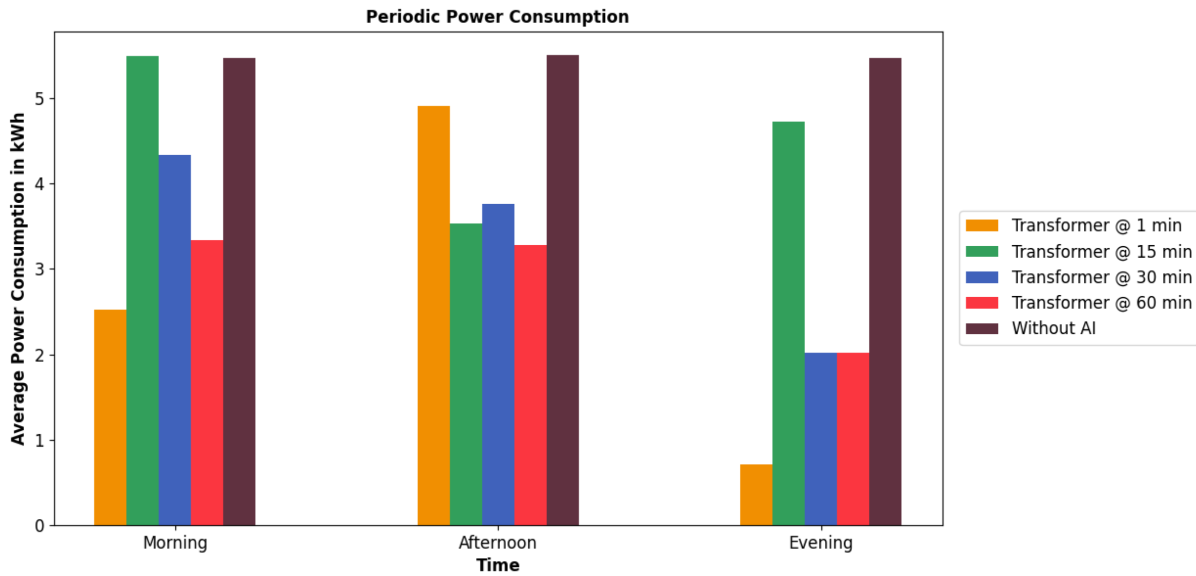

Figure 12 Bar plot representing the net power consumption over 3 time periods in a day.

These findings highlight the power of the transformer model in capturing energy usage patterns and developing optimized control strategies for HVAC systems. By analyzing historical data and making accurate predictions, the model is able to make intelligent adjustments to HVAC systems in real-time, resulting in substantial energy savings. The significantly lower net power consumption achieved by the transformer model indicates its efficacy in optimizing the operation of HVAC systems, leading to more efficient energy utilization and reduced costs.

It is worth noting that the other models tested in the evaluation also demonstrate varying net power consumption values, ranging from 8.67 to 13.83 kWh, depending on the specific time period and forecasting horizon. These variations suggest that different models and forecasting horizons may have distinct impacts on energy consumption patterns. This observation underscores the importance of carefully selecting the appropriate model and forecasting horizon to achieve optimal energy savings and system performance.

Overall, the results from Figure 13 emphasize the potential of the transformer model in capturing energy usage patterns and developing optimized control strategies for HVAC systems. The significant difference in net power consumption between the transformer model with a 1-minute horizon and the non-AI method showcases the advantages of leveraging machine learning techniques in energy optimization. These findings provide valuable insights for building managers

Footnote may be given with Font size 10 of Times New Roman.

and energy professionals seeking to implement more efficient HVAC control strategies, ultimately leading to reduced energy consumption and improved sustainability.

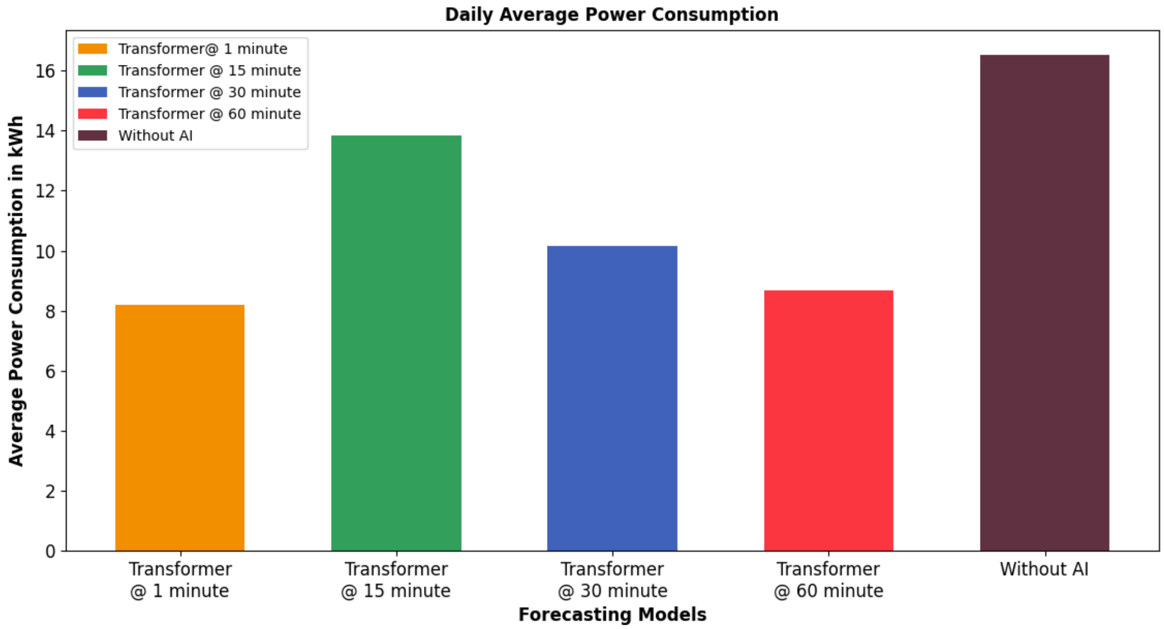

Figure 13 Bar Plot representing daily Average Power Consumption for all models.

Figure 14 provides further insights into the impact of time of day on power consumption in the evaluated models. It reveals that the evening period generally exhibits the lowest power consumption across all models, except for the Transformer model with a 15-minute forecasting horizon. This suggests that the evening period offers favorable conditions for energy optimization, potentially due to lower occupant activity or external factors such as outdoor temperature.

Interestingly, the Transformer model with a 15-minute forecasting horizon displays higher average power consumption during the morning and evening periods compared to other models. This indicates that the model may struggle to accurately predict and adjust for energy requirements during these specific time periods. Conversely, the non-AI method consistently records the highest power consumption across all periods, highlighting the limitations of traditional approaches in energy optimization compared to machine learning-based models.

Footnote may be given with Font size 10 of Times New Roman.

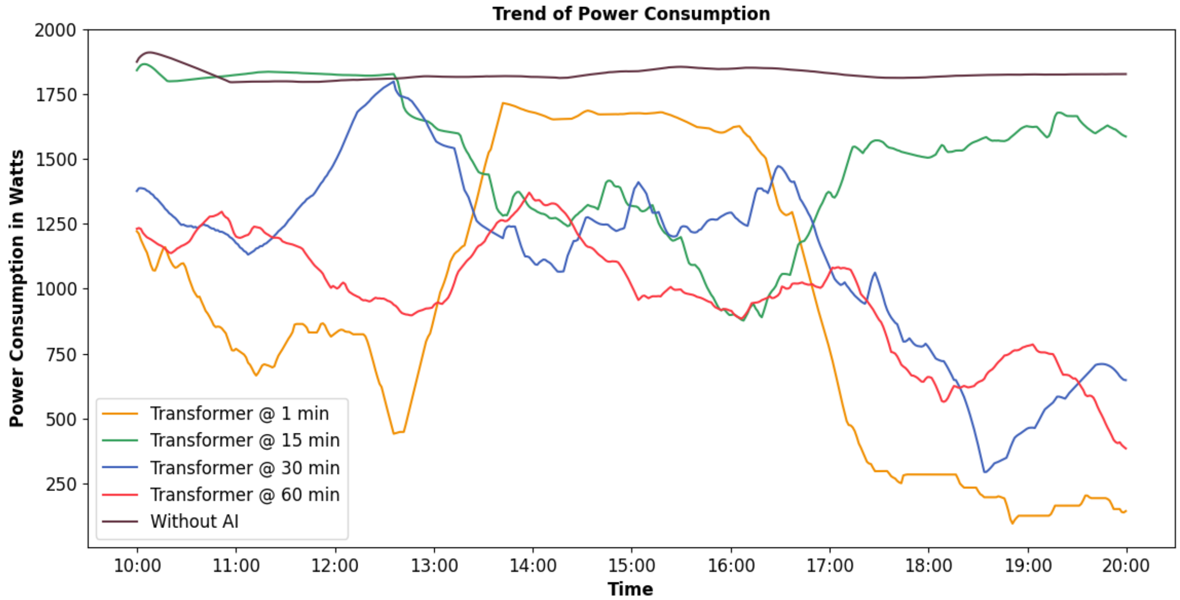

Figure 14 Trend of Power Consumption among all models

The trend of power consumption throughout the day, as depicted in Figure 14, provides valuable information about the behavior of the HVAC system. The yellow line representing the Transformer model shows greater instability at a 1-minute forecasting horizon compared to other models. This suggests that the shorter forecasting horizon may introduce more fluctuations and challenges in accurately predicting energy usage patterns, leading to less stable power consumption trends. However, as the forecasting horizons increase, the instability in the power consumption trends decreases, as shown in Figure 14. Longer forecasting horizons allow the model to capture more data and make more accurate predictions, resulting in smoother power consumption patterns. This observation underscores the importance of selecting an appropriate forecasting horizon that balances the trade-off between prediction accuracy and computational requirements.

The peaks in Figure 14 indicate instances where the HVAC system had to handle higher loads to maintain the desired conditions based on the recorded parameters used for making predictions. These peaks represent periods of increased energy demand, potentially influenced by factors such as occupancy patterns, outdoor temperature fluctuations, or specific activities within the building. On the other hand, the dips in Figure 14 suggest that the HVAC system was put into sleep mode or operated at a reduced capacity when the energy requirements fell below a certain threshold. This approach, known as demand response or load shedding, allows for energy savings

Footnote may be given with Font size 10 of Times New Roman.

during periods of lower demand. Therefore, these findings have important implications for the development of energy-efficient HVAC control strategies as well as the selection of appropriate deep learning models based on the time period and forecasting horizon.

5.3 Thermal Comfort

In addition to assessing energy consumption optimization, this study also focused on evaluating thermal comfort using three different models. The assessment of thermal comfort was conducted over a period of time, during which data on thermal comfort was collected through surveys. These surveys included evaluations of thermal comfort during three distinct time periods, accompanied by the collection of data on metabolic rate and clothing insulation.

The study involved healthy occupants, and their clothing insulation levels ranged from 0.91 to 1.25 clo, in accordance with the standards set by ASHRAE. The metabolic rate of the occupants ranged from 1 to 1.12 met units, which are indicators of their activity levels and heat production. The results presented in Figure 15 demonstrate that the model with a 60-minute forecasting horizon provided the highest level of thermal comfort to the occupants, followed by the model with a 30-minute forecasting horizon. The models with a 1-minute and 15-minute forecasting horizon yielded similar levels of thermal comfort.

Figure 15 includes donut charts that illustrate the percentage of occupants experiencing different levels of thermal comfort during each time period. It is important to note that the thermal comfort scale used in the survey includes categories such as "Neutral," "Slightly Warm," and "Slightly Cool." The outermost donut represents the morning time period, the middle donut represents the afternoon, and the innermost donut represents the evening time period. The colors used in the donut charts represent the thermal sensations reported by the occupants. Specifically, red indicates a feeling of warmth, yellow represents a slightly warm sensation, green signifies a neutral thermal sensation, and blue indicates a slightly cool sensation.

Upon comparing the indoor temperature data from Figure 6 with the thermal sensations depicted in Figure 15, several observations can be made regarding the performance of the deep learning models in relation to occupants' comfort. When the transformers model with a 1-minute

Footnote may be given with Font size 10 of Times New Roman.

forecasting horizon was employed, only one person expressed complete dissatisfaction with the model's performance in the afternoon, as shown in Figure 15(a). This dissatisfaction occurred when the average indoor temperature reached 28.671 degrees Celsius. However, it is important to note that the majority of occupants felt either neutral or slightly warm, indicating a satisfactory level of thermal comfort.

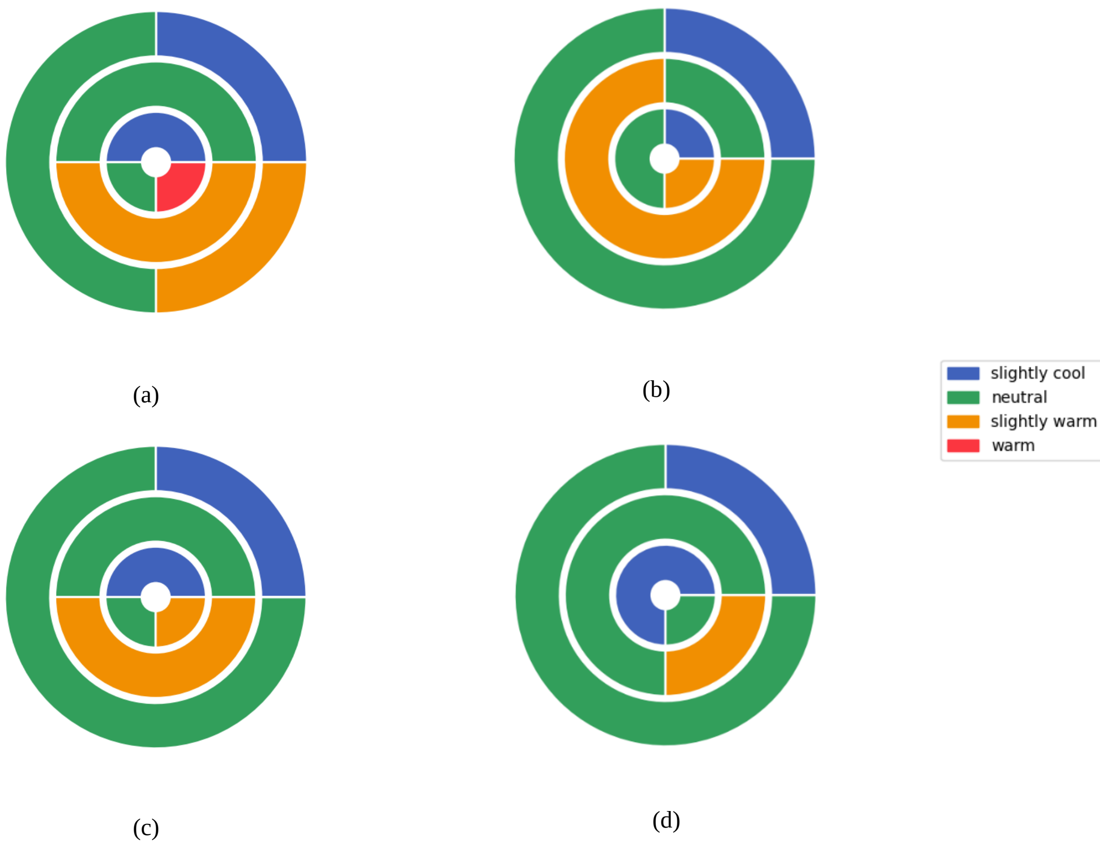

Figure 15 Nested donuts representing the thermal comfort of the occupants for models on different forecasting horizons at three multiple time periods of the day where outermost donut represents morning, middle donut represents afternoon and innermost donut repres

For the models utilizing a 15-minute and 30-minute forecasting horizon, none of the occupants reported feeling uncomfortable. Instead, most occupants reported experiencing partial comfort for 50% of the experimental duration. The average indoor temperature ranged from 26.79 to 27.74 degrees Celsius for the model with a 15-minute horizon. Similarly, for the model with a 30-minute horizon, the indoor temperature ranged from 24.161 degrees Celsius in the morning to 26.79 degrees Celsius in the afternoon and 26.1 degrees Celsius in the evening. These slight

Footnote may be given with Font size 10 of Times New Roman.

variations in average indoor temperature were found to have an impact on the occupants' thermal sensations.

Figure 15(c) demonstrates that when the average indoor temperatures changed slightly for both models, the occupants reported feeling slightly cooler compared to Figure 15(b). Throughout the entire experimental period, the indoor temperature ranged from 23.68 to 27.46 degrees Celsius when using the transformers model with a 60-minute forecasting horizon. During this period, most occupants reported feeling completely satisfied with the thermal conditions, and the sensation of partial comfort was the lowest among all the models tested, as depicted in Figure 15(d).

Footnote may be given with Font size 10 of Times New Roman.

CHAPTER 6: CONCLUSION

This study implements a predictive model based on machine learning to forecast HVAC temperature while optimizing energy consumption and thermal comfort. The goal of this research is to explore the potential of transformers and CNN-LSTM models in achieving accurate predictions for various forecasting horizons. The models were trained on a comprehensive dataset comprising air quality, power quality, and thermal comfort parameters obtained from multiple devices and surveys. The performance of the models was assessed using standard evaluation metrics, including Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and R2. The results of this study indicate that the Transformer model, consisting of 4 transformer blocks with 8 transformer heads in each multi-head attention layer, outperformed the recent CNN-LSTM model. Specifically, the Transformer model achieved an impressive R2 score of 0.936, demonstrating its ability to accurately predict HVAC temperature. The model exhibited low MAE and RMSE values of 0.02° and 0.03° respectively when forecasting 1 minute ahead. These findings highlight the effectiveness of the Transformer model in capturing the complex patterns and dynamics of HVAC temperature. The predicted temperature values generated by the Transformer model were then transmitted to the HVAC system through an infra-red controller, and the energy consumption and thermal comfort levels were monitored using surveys and power quality logs. To evaluate the efficiency of the machine learning models, the collected data was compared with consumption data obtained through a conventional, non-AI approach to HVAC control. The test results revealed that the Transformers model achieved the best balance between energy consumption and thermal comfort when projecting 60 minutes ahead. This model achieved significant energy savings of 47.5% while maintaining a high level of satisfaction among the occupants throughout the entire day. On the other hand, when using a 1-minute forecasting horizon, the energy consumption was significantly reduced to 8.19 kWh, resulting in a 50.5% energy savings. However, this came at the expense of thermal comfort, with 8.33% of occupants reporting feeling uncomfortable compared to 0% for the model with a 60-minute horizon.

Footnote may be given with Font size 10 of Times New Roman.

REFERENCES

[1] S. S. Longo, M. Cellura, M. A. Cusenza, F. Guarino, and I. Marotta, “Selecting Insulating Materials for Building Envelope: A Life Cycle Approach,” *TECNICA ITALIANA-Italian Journal of Engineering Science*, vol. 65, no. 2–4, pp. 312–316, Jul. 2021, doi: 10.18280/ti-ijes.652-426.

[2] M. Romanello *et al.*, “The 2022 report of the Lancet Countdown on health and climate change: health at the mercy of fossil fuels,” *The Lancet*, vol. 400, no. 10363, pp. 1619–1654, Nov. 2022, doi: 10.1016/S0140-6736(22)01540-9.

[3] “Indoor Air Pollution: An Introduction for Health Professionals | US EPA.” https://www.epa.gov/indoor-air-quality-iaq/indoor-air-pollution-introduction-health-professionals (accessed Dec. 24, 2022).

[4] B. Jiang, H. Gong, H. Qin, and M. Zhu, “Attention-LSTM architecture combined with Bayesian hyperparameter optimization for indoor temperature prediction,” *Build Environ*, vol. 224, p. 109536, Oct. 2022, doi: 10.1016/j.buildenv.2022.109536.

[5] B. Dong, Z. Li, S. M. M. Rahman, and R. Vega, “A hybrid model approach for forecasting future residential electricity consumption,” *Energy Build*, vol. 117, pp. 341–351, Apr. 2016, doi: 10.1016/J.ENBUILD.2015.09.033.

[6] A. Foucquier, S. Robert, F. Suard, L. Stéphan, and A. Jay, “State of the art in building modelling and energy performances prediction: A review,” *Renewable and Sustainable Energy Reviews*, vol. 23, pp. 272–288, Jul. 2013, doi: 10.1016/J.RSER.2013.03.004.

[7] D. Enescu, “A review of thermal comfort models and indicators for indoor environments,” *Renewable and Sustainable Energy Reviews*, vol. 79, pp. 1353–1379, Nov. 2017, doi: 10.1016/J.RSER.2017.05.175.

[8] F. Amara, K. Agbossou, A. Cardenas, Y. Dubé, and S. Kelouwani, “Comparison and Simulation of Building Thermal Models for Effective Energy Management,” *Smart Grid and Renewable Energy*, vol. 06, no. 04, pp. 95–112, 2015, doi: 10.4236/SGRE.2015.64009.

[9] R. Z. Homod, “Review on the HVAC System Modeling Types and the Shortcomings of Their Application,” *Journal of Energy*, vol. 2013, pp. 1–10, 2013, doi: 10.1155/2013/768632.

[10] F. Elmaz, R. Eyckerman, W. Casteels, S. Latré, and P. Hellinckx, “CNN-LSTM architecture for predictive indoor temperature modeling,” *Build Environ*, vol. 206, p. 108327, Dec. 2021, doi: 10.1016/j.buildenv.2021.108327.

[11] Y. Chen and H. Tan, “Short-term prediction of electric demand in building sector via hybrid support vector regression,” *Appl Energy*, vol. 204, pp. 1363–1374, 2017, doi: 10.1016/j.apenergy.2017.03.070.

[12] P. Vrablecová, A. Bou Ezzeddine, V. Rozinajová, S. Šárik, and A. K. Sangaiah, “Smart grid load forecasting using online support vector regression,” *Computers and Electrical Engineering*, vol. 65, pp. 102–117, Jan. 2018, doi: 10.1016/j.compeleceng.2017.07.006.

[13] J. Moon, Y. Kim, M. Son, and E. Hwang, “Hybrid short-term load forecasting scheme using random forest and multilayer perceptron,” *Energies (Basel)*, vol. 11, no. 12, Dec. 2018, doi: 10.3390/EN11123283.

Footnote may be given with Font size 10 of Times New Roman.

[14] Z. Wang, T. Hong, and M. A. Piette, “Building thermal load prediction through shallow machine learning and deep learning,” *Appl Energy*, vol. 263, Apr. 2020, doi: 10.1016/j.apenergy.2020.114683.

[15] J. Seo, S. Kim, S. Lee, H. Jeong, T. Kim, and J. Kim, “Data-driven approach to predicting the energy performance of residential buildings using minimal input data,” *Build Environ*, vol. 214, Apr. 2022, doi: 10.1016/j.buildenv.2022.108911.

[16] E. Kamel, S. Sheikh, and X. Huang, “Data-driven predictive models for residential building energy use based on the segregation of heating and cooling days,” *Energy*, vol. 206, p. 118045, Sep. 2020, doi: 10.1016/J.ENERGY.2020.118045.

[17] A. Rahman, V. Srikumar, and A. D. Smith, “Predicting electricity consumption for commercial and residential buildings using deep recurrent neural networks,” *Appl Energy*, vol. 212, pp. 372–385, Feb. 2018, doi: 10.1016/j.apenergy.2017.12.051.

[18] C. Xu, H. Chen, J. Wang, Y. Guo, and Y. Yuan, “Improving prediction performance for indoor temperature in public buildings based on a novel deep learning method,” *Build Environ*, vol. 148, pp. 128–135, Jan. 2019, doi: 10.1016/j.buildenv.2018.10.062.

[19] F. Mtibaa, K.-K. Nguyen, M. Azam, A. Papachristou, J.-S. Venne, and M. Cheriet, “LSTM-based indoor air temperature prediction framework for HVAC systems in smart buildings,” *Neural Comput Appl*, vol. 32, no. 23, pp. 17569–17585, Dec. 2020, doi: 10.1007/s00521-020-04926-3.

[20] Z. Fang, N. Crimier, L. Scanu, A. Midelet, A. Alyafi, and B. Delinchant, “Multi-zone indoor temperature prediction with LSTM-based sequence to sequence model☆,” *Energy Build*, vol. 245, Aug. 2021, doi: 10.1016/j.enbuild.2021.111053.

[21] B. Jiang, H. Gong, H. Qin, and M. Zhu, “Attention-LSTM architecture combined with Bayesian hyperparameter optimization for indoor temperature prediction,” *Build Environ*, vol. 224, p. 109536, Oct. 2022, doi: 10.1016/j.buildenv.2022.109536.

[22] M. Liu *et al.*, “SCINet: Time Series Modeling and Forecasting with Sample Convolution and Interaction”, Accessed: Jan. 06, 2023. [Online]. Available: https://github.com/cure-lab/SCINet.

[23] “ASHRAE55-version2017”.

[24] J. Kim, Y. Zhou, S. Schiavon, P. Raftery, and G. Brager, “Personal comfort models: Predicting individuals’ thermal preference using occupant heating and cooling behavior and machine learning,” *Build Environ*, vol. 129, pp. 96–106, Feb. 2018, doi: 10.1016/j.buildenv.2017.12.011.

[25] Menne, Matthew J., Imke Durre, Bryant Korzeniewski, Shelley McNeill, Kristy Thomas, Xungang Yin, Steven Anthony, Ron Ray, Russell S. Vose, Byron E.Gleason, and Tamara G. Houston (2012): Global Historical Climatology Network - Daily (GHCN-Daily), Version 3. NOAA National Climatic Data Center. doi:10.7289/V5D21VHZ Accessed March 4, 2023.

[26] cathal john murray, N. Du Bois, L. Hollywood, and D. Coyle, “State-of-The-Art Deep Learning Models are Superior for Time Series Forecasting and are Applied Optimally with Iterative Prediction Methods,” *SSRN Electronic Journal*, 2023, doi: 10.2139/SSRN.4361707.

[27] X. Shi *et al.*, “Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting”.

Footnote may be given with Font size 10 of Times New Roman.

[28] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2323, 1998, doi: 10.1109/5.726791.

[29] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Comput*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[30] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, “LSTM: A Search Space Odyssey,” *IEEE Trans Neural Netw Learn Syst*, vol. 28, no. 10, pp. 2222–2232, Mar. 2015, doi: 10.1109/TNNLS.2016.2582924.

[31] L. N. Smith, “A disciplined approach to neural network hyper-parameters: Part 1 -- learning rate, batch size, momentum, and weight decay,” Mar. 2018, doi: 10.48550/arxiv.1803.09820.

[32] S. Ahmed, I. E. Nielsen, A. Tripathi, S. Siddiqui, G. Rasool, and R. P. Ramachandran, “Transformers in Time-series Analysis: A Tutorial,” Apr. 2022, doi: 10.48550/arxiv.2205.01138.

[33] A. Vaswani *et al.*, “Attention Is All You Need,” *Adv Neural Inf Process Syst*, vol. 2017-December, pp. 5999–6009, Jun. 2017, doi: 10.48550/arxiv.1706.03762.

[34] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, “Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling,” Dec. 2014, doi: 10.48550/arxiv.1412.3555.

[35] “Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance on JSTOR.” https://www.jstor.org/stable/24869236 (accessed Mar. 06, 2023).

[36] R. J. Hyndman and A. B. Koehler, “Another look at measures of forecast accuracy,” *Int J Forecast*, vol. 22, no. 4, pp. 679–688, Oct. 2006, doi: 10.1016/J.IJFORECAST.2006.03.001.

[37] F. Emmert-Streib and M. Dehmer, “Evaluation of Regression Models: Model Assessment, Model Selection and Generalization Error,” *Machine Learning and Knowledge Extraction 2019, Vol. 1, Pages 521-551*, vol. 1, no. 1, pp. 521–551, Mar. 2019, doi: 10.3390/MAKE1010032.

Footnote may be given with Font size 10 of Times New Roman.

Instructions for Students

1. Color of the thesis binding shall be BLACK with Golden printing of text. Moreover, the title must run on the side of binding so that thesis title is visible once shelved in library.
2. All Engineering students will use referencing as given in the format; whereas Biomedical Sciences students will be use **APA** (American Psychological Association) **style** referencing.

Footnote may be given with Font size 10 of Times New Roman.