Cyprus University of Technology

杭州電子科技大学

HANGZHOU DIANZI UNIVERSITY

Master's thesis

IAQ prediction and IAQ improvement strategy based on ETSformer neural network

Zhou Zhou

Limassol, June 2025

MSc in Electronics Science and Technology

CYPRUS UNIVERSITY OF TECHNOLOGY

Faculty of Engineering and Technology

Department of Electrical Engineering, Computer Engineering, and Informatics

Master’s thesis

IAQ prediction and IAQ improvement strategy based on ETSformer neural network

Zhou Zhou

Supervisor

Michalis Michaelides

Limassol, June 2025

Approval Form

Master's thesis

IAQ prediction and IAQ improvement strategy based on ETSformer neural network

Presented by

Zhou Zhou

Supervisor: Michalis Michaelides

Member of the committee: Committee member 1

Member of the committee: Committee member 2

Cyprus University of Technology
Limassol, June 2025

Copyrights

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor, Dr. Michalis Michaelides, for his continuous guidance and support throughout the course of this research. His academic rigor, insightful feedback, and encouragement have been invaluable to the completion of this thesis.

I would also like to thank the members of the thesis committee for their time, constructive comments, and thoughtful suggestions, which significantly enhanced the quality of this work.

I am grateful to the Cyprus University of Technology and the Department of Electrical Engineering, Computer Engineering, and Informatics for providing an excellent research environment and the necessary resources for this study.

Many thanks go to my labmates and friends, whose companionship and support made this academic journey both productive and enjoyable.

A very special thank you goes to my girlfriend, Shiyu Xu. Although we have been in a long-distance relationship across countries, she has never failed to accompany and support me through countless video calls and messages. Her unwavering care, continuous encouragement, and deep understanding have been a constant source of strength during my academic journey. In times of pressure and uncertainty, her presence—even from afar—has given me the motivation and courage to keep moving forward.

Finally, I would like to express my deepest appreciation to my family for their unconditional support, patience, and love. Without them, none of this would have been possible.

ABSTRACT

This thesis proposes a novel indoor air quality (IAQ) prediction model and management strategy based on the ETSformer neural network. Unlike traditional methods that require prior data cleaning, the proposed model directly processes raw IAQ sensor data, which may contain anomalies such as noise, missing values, and outliers. The ETSformer model integrates exponential smoothing, seasonal decomposition, and damping control mechanisms within a Transformer-based architecture, enabling accurate multi-step forecasting of key IAQ parameters, including temperature, humidity, CO_2, and PM2.5 concentrations.

Based on the predictions, this thesis further designs a scenario-specific IAQ management strategy combining prediction-driven control and feedback regulation. Two distinct control objectives are defined: comfort optimization during normal periods and infection prevention during epidemic periods. Experimental results on real-world data collected from a university classroom demonstrate that the ETSformer model outperforms traditional baselines in terms of prediction accuracy and robustness. The integration of forecasting and intelligent control enables timely and efficient air quality regulation, offering practical value for health-oriented and energy-efficient indoor environment management.

Keywords: indoor air quality, time-series forecasting, ETSformer, Transformer, prediction-based control

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

1.1 Aims and Objectives

In the field of indoor air quality (IAQ) improvement, this study proposes an advanced time series prediction method based on a Transformer neural network architecture. The proposed model is specifically designed to forecast key indoor air quality indicators—such as temperature, humidity, carbon dioxide (CO_2) concentration, and PM2.5 levels—with high temporal resolution and predictive accuracy. As shown in Figure 1.1, by accurately modeling and predicting the future evolution of these environmental parameters, the system can effectively support the real-time and proactive operation of IoT-enabled air purification and ventilation devices. This predictive capability addresses the inherent delays and inefficiencies of traditional reactive systems, enabling intelligent indoor air management through timely control actions. Ultimately, the integration of Transformer-based forecasting with IoT-driven actuation lays the foundation for a more responsive, energy-efficient, and health-oriented air quality management strategy.

1.2 Research Questions

This study focuses on two core issues. First, the training of the ETSformer neural network is explored as the foundation of the proposed prediction framework. The model inherently possesses the capability to handle anomalies such as missing values, noise, and outliers through its internal smoothing and decomposition mechanisms. This allows for direct modeling of raw indoor air quality data without the need for explicit data cleaning procedures. The model is trained to capture both short-term fluctuations and long-term trends in key indicators such as temperature, humidity, CO_2 concentration, and PM2.5 levels, enabling accurate multi-step forecasting. Second, based on the predicted results, a scenario-specific indoor air quality management strategy is developed. Specifically, the strategy distinguishes between two contexts: during normal periods, the system targets comfort optimization; during epidemic periods, it prioritizes infection prevention. This dual-mode strategy allows IoT-enabled environmental control devices to respond intelligently and proactively to varying health and comfort requirements.

1.3 Significance of The Subject

In modern society, indoor environment has become the main place for people's activities. However, the impact of indoor air quality (IAQ) on health and efficiency is often overlooked. Research has shown that indoor air pollutants such as carbon dioxide (CO_2), volatile organic compounds (VOCs), particulate matter (PM2.5 and PM10), and microorganisms not only affect physical health, but also lead to fatigue, decreased attention, and weakened cognitive abilities. Long term exposure to poor air quality may cause respiratory diseases, allergies, and other health problems, while also weakening work efficiency and learning outcomes.[1]

In the work environment, when the concentration of CO_2exceeds 1000 ppm, human cognitive ability will significantly decrease; Higher concentrations can seriously affect decision-making and problem-solving

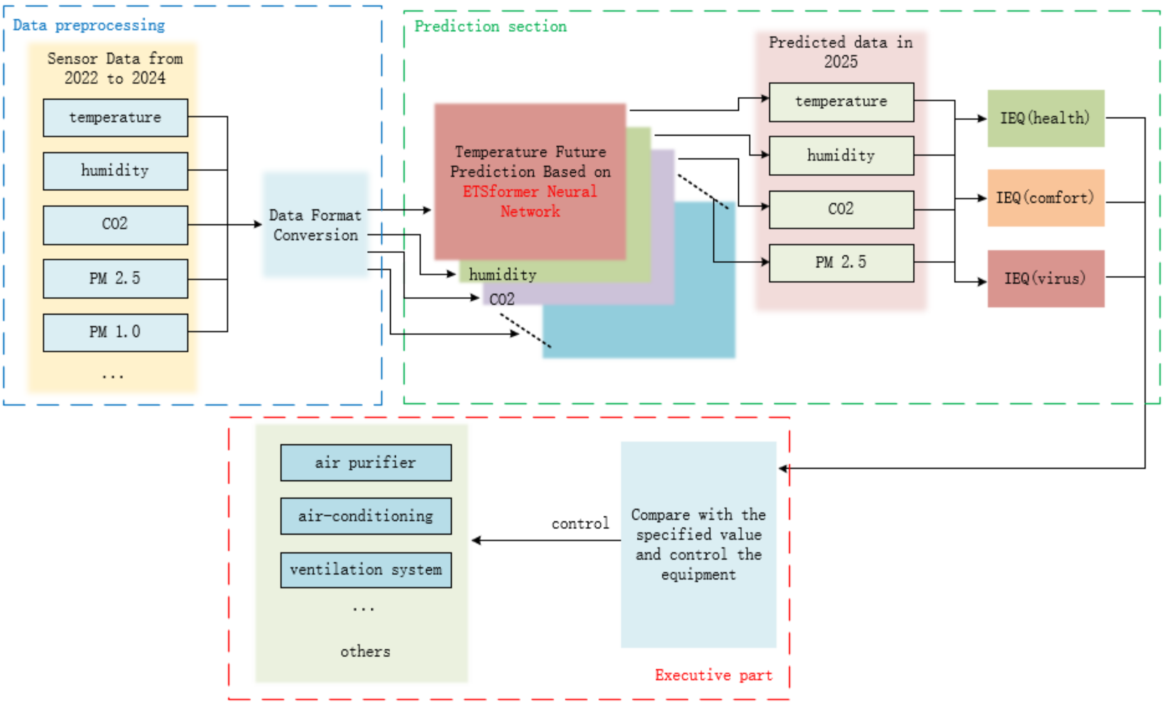

Figure 1.1: Framework of the proposed indoor air quality prediction and control strategy based on the ETSformer neural network.

abilities. Air pollution also increases the likelihood of employees taking sick leave, weakening team productivity and creativity.

In the learning environment, pollutants can reduce students' attention and memory, affecting exam scores and long-term learning outcomes. Excessive or insufficient humidity may also lead to virus transmission, exacerbating the health risks for students.[2][3]

In addition, air quality has a profound impact on psychology and emotions. Air pollution can cause fatigue and anxiety, reducing work motivation and learning focus. In contrast, fresh air can significantly improve mental health and task completion efficiency.

Indoor air quality is not only related to health, but also an important factor affecting social and economic benefits. Improving the air environment and enhancing IAQ through science and technology can significantly reduce health expenses, improve work and study efficiency, and create greater value for society.

To address the aforementioned issues, this article proposes a neural network-based indoor air quality prediction method and an IoT controlled air improvement equipment strategy.

Aiming to improve the accuracy of air quality prediction and enhance the computing power of the system to drive air quality improvement equipment and achieve air quality management. Clean past air quality data, extract effective data, and use deep learning techniques such as Transformer to predict future air quality. That is to say, it no longer relies solely on traditional sensor data, but uses machine learning to predict future air quality data, compensate for the lag in strategy of traditional air quality improvement systems, and make early predictions and improvements to indoor air. Specifically, the strategy of this paper is to transmit and process data through the Internet distributed platform. In order to further optimize the real-time performance of the air quality management system, the Transformer model is introduced. This model is superior in processing time allowed data, enabling the system to achieve more precise management of air quality under the conditions of complex indoor environmental conditions. Overall, compared to traditional system control methods, the proposed solution achieves better air quality management.

Traditional indoor air quality management mainly relies on the combination of monitoring methods and improvement strategies, aiming to create a healthier indoor environment. In terms of monitoring, traditional methods typically involve regular manual testing and laboratory analysis. These methods measure the concentrations of common indoor pollutants, including carbon dioxide (CO_2), PM2.5, PM10, and volatile organic compounds (VOCs), using chemical reagents or professional instruments. This method can provide more accurate detection results and is suitable for regular review or evaluation of indoor air quality. However, due to its intermittent nature, traditional monitoring is difficult to achieve real-time data collection and cannot dynamically reflect the rapid changes in indoor air quality. In addition, such monitoring usually relies on professional technicians and equipment, with high costs and limited coverage, making it difficult to meet the needs of modern diversification and efficient management.[4][5][6]

In terms of improving air quality, ventilation is considered one of the core means of traditional management methods. Natural ventilation allows for the exchange of indoor and outdoor air by opening windows or ventilation openings, thereby diluting pollutant concentrations and enhancing air freshness. However, in environments with high levels of urbanization and high concentrations of pollutants, the effectiveness of natural ventilation is often limited. For example, in situations where outdoor air quality is poor or building design ventilation conditions are poor, natural ventilation may not be effective. To address this issue, mechanical ventilation systems such as ventilation fans and fresh air systems are widely used. This type of system can to some extent compensate for the lack of natural ventilation by actively introducing fresh air and discharging indoor polluted air.[7][8]

In addition, air purifiers, as important auxiliary equipment, have gradually become an important tool for improving indoor air quality. These devices remove particulate matter (such as PM2.5 and PM10) from the air through efficient filtration systems, and partially adsorb or decompose harmful gases (such as formaldehyde and VOCs). The use of air purifiers is particularly important during high pollution seasons or in special environments such as children's rooms and medical facilities. It can not only improve air quality, but also effectively reduce the threat of indoor pollutants to respiratory health. But without data support, there will be no efficient improvement in air quality, and it may result in energy waste and untimely air quality regulation.[9]

Although traditional methods have alleviated indoor air pollution to some extent, their effectiveness is still limited by the intermittency of monitoring methods and the passivity of improvement strategies. The changes in indoor air quality are complex and dynamic, and traditional methods are difficult to quickly respond to sudden changes in pollutant concentrations.

In addition, modern people's demand for air quality is becoming increasingly refined, and traditional methods alone are no longer sufficient to meet efficient and comprehensive management requirements. Therefore, modern indoor air management is developing towards intelligence and automation, by introducing IoT sensors and automatic control systems to achieve real-time monitoring and dynamic regulation of indoor air. This not only improves the efficiency of air quality management, but also creates a healthier and more comfortable indoor environment for people.[10]

In summary, this article proposes an indoor air quality prediction method based on Transformer neural network and provides a prediction based air quality adjustment strategy for effectively improving current indoor air quality improvement methods.

1.4 Contribution

1. Design and train an ETSformer-based indoor air quality forecasting model that integrates trend decomposition, seasonality extraction, and damping control mechanisms. The model is capable of directly modeling raw sensor data with inherent anomalies, and achieves accurate multi-step forecasting of key environmental indicators—such as temperature, humidity, CO_2 concentration, and PM2.5 levels—across multiple temporal scales.

2. Propose a prediction-driven indoor air quality management strategy that distinguishes between normal and epidemic scenarios. The strategy targets comfort optimization under normal conditions and infection risk mitigation during epidemic periods. By integrating with IoT-enabled purification and ventilation systems, it enables proactive, automated control across diverse indoor environments.

1.5 Structure of the Thesis

In the first chapter Introduction of this article, a rough overview of the research question and contribution will be provided. In the second chapter Literature Review, the significance of The Subject will be discussed, Provide a detailed description of the current research status, Analysis of the Advantages of Transformer Networks, and analyze the current research status from the perspectives of traditional methods and machine learning. In Chapter 3 Research Methodology, specific solutions will be proposed, divided into three sections, which explain data preprocessing, model training, equipment control, and their specific formulas. In Chapters 4 and 5, we will continue to analyze the effectiveness of the method proposed in this paper, compare it with relevant algorithms in the same field, analyze the advantages and disadvantages of our algorithm, and propose directions and possibilities for future improvement.

2 Literature Review

2.1 Current research status

2.1.1 Traditional indoor air quality data processing methods

In indoor air quality (IAQ) management, data processing is an important step in monitoring, predicting, and improving air pollution. Traditional data processing methods focus on monitoring, analyzing, and predicting pollutant concentrations, relying on statistical models and classic machine learning techniques to achieve quantitative assessment and trend prediction of air pollution through systematic data cleaning, feature extraction, and modeling steps. These methods provide a scientific basis for improving indoor air quality and are widely applied in practical scenarios.Y Rybarczyk et al. reviewed the application of machine learning models in air quality modeling, emphasizing the combination of large-scale data processing and modeling methods[11]. Subsequently, Ameer et al. compared and analyzed the performance of different machine learning methods in air quality prediction, and explored their applications in data processing and pollutant concentration prediction[12]. And TV Vu et al. also successfully used machine learning methods to analyze the impact of Beijing's Clean Air Action on air quality trends, and explored the combination of statistical analysis and predictive models[13].

The first step in data processing is cleaning and preprocessing, which aims to improve data quality and eliminate noise and missing values in sensor data. Common methods include moving average filtering, which can smooth short-term data fluctuations and make the trend of pollutant concentration changes clearer. In terms of handling missing data, linear interpolation and spline interpolation are widely used to fill incomplete data and ensure data continuity and integrity. These fundamental treatments not only enhance the reliability of analysis, but also lay a solid foundation for subsequent data modeling and prediction. However, the limitation of traditional data cleaning methods is that they cannot distinguish between useful information in signals and potential noise in complex pollutant patterns. For example, a simple moving average may lead to excessive smoothing of key pollutant fluctuations and loss of important abnormal change signals[14][15].

After data preprocessing, feature extraction and dimensionality reduction techniques become key steps for efficiently processing complex indoor air quality data. Principal Component Analysis (PCA) is a common dimensionality reduction technique that reduces data dimensions, highlights key features, and effectively reduces computational complexity[16]. It is particularly suitable for high-dimensional and multivariate air quality monitoring data. In addition, time series feature extraction technology can help identify key patterns of pollutant changes, extract features such as mean, standard deviation, daily peak, etc., thereby revealing the periodic changes in pollutant concentration and providing accurate input for subsequent modeling. However, the application of PCA and other methods also has certain shortcomings, mainly manifested in their insufficient ability to model the nonlinear relationships between data features, and their inability to capture potential nonlinear interaction effects in complex dynamic environments.

Data modeling and prediction are the core components of traditional indoor air quality data processing. At this stage, statistical models and classical machine learning algorithms played an important role. The Autoregressive Integral Moving Average Model (ARIMA) is a classic time series analysis tool that can effectively capture trends and seasonal fluctuations in pollutant concentrations. This model performs particularly well for air pollutants with significant periodic changes, such as PM2.5 and VOCs. In addition, support vector machine (SVM) regression performs outstandingly in handling complex and dynamic pollutant concentration prediction tasks due to its powerful nonlinear modeling ability. However, the shortcomings of these traditional modeling methods lie in their poor adaptability to high-dimensional complex environments. And after H Yao et al. found that, ARIMA models require strict assumptions about the stationarity of data and have limited performance for long-term forecasting. Although support vector machines can capture some nonlinear relationships, they require manual definition of kernel functions, which increases the complexity of feature design[17][18][19].

In practical applications, although traditional methods perform well in specific scenarios, they face challenges in dealing with dynamic changes and processing heterogeneous data from multiple sources. For example, in formaldehyde concentration prediction, simple moving averages and ARIMA models may lack responsiveness to sudden events (such as sudden release from high pollution sources), resulting in prediction delays or significant errors. In PM2.5 concentration prediction, the combination of principal component analysis and support vector regression can improve prediction accuracy, but its ability to handle complex noise interference and potential feature interactions in sensor data is limited.

Traditional methods heavily rely on manually designed features in data cleaning, feature extraction, and modeling, which limits their ability to recognize complex pollution patterns due to their passivity. In addition, traditional methods are mostly static analysis methods that cannot dynamically respond to rapid changes in indoor environments. This limitation is increasingly prominent in the modern, multivariate, and highly dynamic indoor environment management needs[20][21].

To address these issues, the development of modern technology has provided new solutions to the shortcomings of traditional methods. For example, deep learning models can better capture the nonlinear relationships between complex data through automated feature extraction techniques, improving the robustness and accuracy of predictions. In addition, by combining real-time sensors and IoT technology, modern methods can achieve dynamic monitoring and rapid response, making indoor air quality management more efficient and accurate. Although traditional methods have laid the theoretical and practical foundation for IAQ management, the introduction of modern methods has opened up new possibilities for further optimizing air quality management. By integrating traditional and modern technologies, indoor air quality management will become more intelligent in the future, providing humans with a healthy and safe living and working environment[22][23]

2.1.2 Research Status of Deep Learning in Indoor Air Quality Prediction

In the field of indoor air quality prediction, different models have played important roles in different historical stages, from early statistical models to recent deep learning models. The continuous advancement of these technologies has promoted the precision and intelligence of air quality management. The following is an analysis of the contributions, advantages, and disadvantages of the main models in indoor air quality prediction.

In the early days, the Seasonal Autoregressive Integral Moving Average (SARIMA) model was a classic method for time series forecasting, widely used to model the periodic changes in indoor air quality. J. Dutta proposed the use of SARIMA for predicting indoor air pollutant concentrations (such as PM2.5 and

CO □) and analyzed its advantages and disadvantages in time series forecasting, indicating that SARIMA can effectively capture the seasonal characteristics of pollutant concentrations (such as PM2.5 and CO □), providing reliable results for short-term forecasting[24][25]. However, it requires high stability of the data, complex differential processing, and difficulty in modeling nonlinear relationships, which limits its performance in long-term forecasting.

With the increasing demand for dynamic system analysis, Markov models (MM) have been introduced into the field of air quality prediction. NN Zakaria uses Markov models to analyze the transition probabilities of pollution states (such as ”excellent”, ”good”, ”poor”), which is suitable for modeling the dynamic evolution process of air quality[26][27]. It has high computational efficiency and can perform long-term trend analysis, but due to the assumption that the future state depends only on the current state, this model is difficult to capture complex temporal dependencies. In addition, MM lacks the ability to model continuous values, which limits its application in concentration prediction.

Subsequently, support vector machine (SVM), as a classic machine learning method, has shown outstanding performance in air quality prediction due to its powerful nonlinear modeling ability. Zhou and Lai proposed using SVM for PM2.5 concentration prediction, emphasizing the applicability of the model in complex environments and improving its time series adaptability[28][29]. SVM can effectively capture the complex nonlinear relationship of pollutant concentration with the help of kernel functions, especially suitable for small sample data environments. However, SVM is sensitive to the selection of parameters and kernel functions, the tuning process is complex, and its performance is limited when dealing with large-scale data.

At the same time, Gaussian Process (GP), as a Bayesian non parametric method, has also been studied by Y Zhu et al. applied it to air quality prediction[30][31][32]. GP provides the ability to quantify model uncertainty and is suitable for capturing the randomness and nonlinear features in pollutant concentrations. Its advantage lies in providing a prediction interval, making it very valuable for risk assessment. However, GP has a high computational complexity, making it difficult to scale to large-scale datasets, and the model training time is relatively long.

In the process of gradually intelligentizing air quality prediction, rule-based tree models such as M5P tree model have been used by Alsultanny and Esmaeilbeiki et al. to predict the trend of air pollutant concentration due to their fast modeling ability[33][34]. M5P combines regression trees with linear regression and performs well in capturing local trends in air pollutant concentrations. Its model structure is easy to explain and suitable for real-time prediction, but its performance is slightly inadequate when dealing with nonlinear complex relationships.

The rise of artificial neural networks (ANN) and backpropagation neural networks (BPNN) has introduced more powerful nonlinear modeling capabilities for air quality prediction. ANN can simultaneously process multidimensional features and is suitable for regression tasks with complex data. However, due to the large number of model parameters, it requires a large amount of data and is prone to overfitting problems. As an implementation of ANN, BPNN improves prediction accuracy by optimizing weights through error backpropagation. Therefore, Li et al. proposed a BPNN model that combines adaptive multi-objective optimization to simultaneously predict and control indoor CO □ and PM2.5 concentrations, and analyzed the performance of BPNN in dynamic prediction[35][36]. However, BPNN has limited performance in processing time series data as it fails to effectively model temporal dependencies.

The emergence of Recurrent Neural Networks (RNNs) has made up for the above shortcomings, as they capture short-term dependencies in time series through a cyclic mechanism and perform excellently in time series modeling of indoor air quality. However, RNNs suffer from gradient vanishing problems in long-term dependency modeling, which limits their performance. This problem has been well solved in Long Short Term Memory (LSTM) networks, and Rahim studied the performance of a prediction model combining LSTM and GRU in indoor air quality dynamic monitoring, especially in capturing time series dependencies[37]. The gating mechanism of LSTM enables it to remember long-term dependencies while making flexible responses to dynamic changes in pollutant concentrations. Gated Recurrent Unit (GRU), as a simplified version of LSTM, achieves a good balance between computational efficiency and modeling capability, and is another commonly used time series prediction model. Gurumoorthy explores the combination of bidirectional GRU models and optimization algorithms for an efficient solution to air quality prediction[38].

Random Forest (RF), as an ensemble learning algorithm, improves the robustness and accuracy of predictions by combining multiple decision trees. It has strong ability to process high-dimensional features and can effectively capture nonlinear relationships. Tagliabue combines IoT networks and uses random forests to predict indoor air quality in educational facilities, demonstrating its superior performance in data-driven methods[39]. However, the interpretability of random forests is poor, and the application of the model in real-time prediction scenarios is limited to some extent.

From traditional statistical models to modern deep learning techniques, each model has its unique advantages and application scenarios in predicting indoor air quality. However, these models also have their own shortcomings, such as the limitations of SARIMA and MM in handling nonlinear data, the computational bottlenecks of SVM and GP in large-scale data, and the high demand for data volume and computing resources in deep learning models. Therefore, in future research, combining the advantages of different models to form multi model collaboration or hybrid methods may be an effective path to improve the accuracy and efficiency of air quality prediction. Through continuous technological advancements, indoor air quality management will be able to better meet the health and environmental needs of modern society[40][41].

2.2 Analysis of the Advantages of Transformer Networks

The Transformer model initially achieved significant results in the field of natural language processing, but in recent years, its application in time series prediction has gradually received attention. As one of the important scenarios for time series modeling, indoor air quality prediction has brought many advantages to this field with the introduction of Transformer. Through its unique self attention mechanism and efficient parallel computing architecture, Transformer overcomes many limitations of traditional models and provides a more accurate and flexible solution for predicting pollutant concentrations in complex indoor environments[42][43][44][45].

The core advantage of Transformer lies in its self attention mechanism. Traditional time series models, such as RNN and LSTM, gradually process time step information in a cyclic manner. Although they can capture temporal dependencies in the sequence, they are easily limited by the vanishing gradient problem when modeling long-term dependencies. The self attention mechanism of Transformer allows the model to directly calculate the correlation between any time step in the sequence, thereby efficiently capturing the long-term and short-term dependencies of pollutant concentration over time. Dong proposed a short-term air quality prediction model that combines EMD and Transformer, and combined it with bidirectional LSTM to improve modeling accuracy[46]. This characteristic is particularly important in indoor air quality prediction, as pollutant concentrations are not only influenced by short-term indoor activities such as cooking and cleaning, but also closely related to long-term trends such as seasonal changes and ventilation patterns. The self attention mechanism can dynamically adjust the level of attention to different time steps, providing a more comprehensive information basis for prediction[47].

Transformers also have significant advantages in handling multidimensional inputs and complex nonlinear relationships. Indoor air quality data typically includes multiple sensor inputs, such as CO □, PM2.5, humidity, and temperature. The relationship between these variables may exhibit highly nonlinear and time-varying characteristics. Transformer can process information from different feature dimensions in parallel through its multi head self attention mechanism, capturing complex interaction relationships between variables. Meanwhile, compared with traditional machine learning methods such as support vector machines or Gaussian processes, Transformer does not rely on manually designed features, but automatically extracts potential patterns from data through deep learning, thereby improving the accuracy of predictions[48][49].

Parallel computing is another significant advantage of Transformer over recursive models. Traditional RNN and LSTM models, due to their cyclic structure, make it difficult to fully parallelize the computation process of time series, resulting in low training efficiency. Especially when processing long time series data, the computation time will significantly increase. Transformer abandons the cyclic structure and adopts a fully parallel computing architecture, preserving the sequential information of time series through positional encoding. This design greatly improves the computational efficiency of the model, enabling it to process large-scale indoor air quality data and support real-time prediction.

In addition, the flexibility of Transformer is also reflected in its scalability and interpretability. By introducing multi-layer Transformer Encoder and Decoder structures, the model can adapt to prediction tasks of different complexities, such as short-term forecasting and long-term trend modeling. At the same time, the self attention weight matrix provides natural interpretability, which can intuitively display the degree of attention the model pays to different time steps and features, providing important references for studying the key influencing factors of indoor air pollution[50][51][52].

Overall, the Transformer model provides a new technological path for indoor air quality prediction through its self attention mechanism, efficient parallel computing, and powerful feature extraction capabilities. It can not only handle multivariate data modeling in complex environments, but also meet real-time prediction needs through efficient computing methods. In the future, with the expansion of indoor air quality data scale and the maturity of model optimization technology, the application potential of Transformer in this field will be further released, providing strong technical support for intelligent indoor environment management[53].

3 Research Methodology

In this chapter, I will provide a detailed introduction to the experimental methodology of this research. The objective of the experiment is to accomplish indoor air quality prediction based on deep learning models, with the core focus on the model training process. Specifically, the preprocessed data will be divided into training and validation sets according to a certain ratio, and then fed into the constructed Transformer-based neural network variant, ETSformer, for training. During the training process, key factors such as model parameter adjustment, optimizer configuration, and loss function selection will be carefully tuned, combined with validation set results for performance evaluation and optimization, in order to prevent overfitting and other adverse effects. Ultimately, through multiple rounds of training and validation, the goal is to obtain an indoor air quality prediction model with good generalization ability for future time windows.

3.1 Training Framework Overview

The training process of the ETSformer model proposed in this study adopts a modular design, which consists of five main stages: parameter configuration, data loading and preprocessing, model construction, training control, and model evaluation. The complete training framework is illustrated in Figure **??**, where each component collaborates to achieve effective modeling and prediction of time series data.

First, the training parameters are uniformly defined through the `argparse` interface, including the number of encoder and decoder layers, model dimensionality, number of attention heads, input and prediction sequence lengths, optimizer type, learning rate scheduling strategy, number of training epochs, and the early stopping mechanism. These parameter settings are organized in the `run.py` file, providing a unified entry point and control basis for the entire training workflow.

In the data processing stage, the framework calls the `data_provider` interface to construct a customized `Dataset_Custom` dataset based on the specified `features` mode and prediction target. This dataset class performs data loading, splits the data into training, validation, and testing sets with a ratio of 7:1:2, and applies standardization using the `StandardScaler` to the input features. Moreover, frequency-aware time feature encoding is enabled in the system, extracting periodic time information including hour, day of the week, and month.

For model construction, the `Exp_Main` class calls the `_build_model()` method to initialize the ETSformer architecture. The model consists of an input embedding layer, stacked encoders, and decoders. The encoder integrates an exponential smoothing module for growth trend modeling and applies Fourier transform for seasonal component extraction. Meanwhile, the decoder introduces a damping mechanism to achieve extrapolation for future time steps.

The training process is executed by the `exp.train()` method. In each training epoch, the model performs forward propagation and loss calculation on the training set, followed by parameter updates using a customized Adam optimizer. During training, the validation set is periodically used to monitor model performance. If the validation error shows no significant decrease over consecutive epochs, the early stopping mechanism is triggered to prevent overfitting. Both training and validation losses are fully

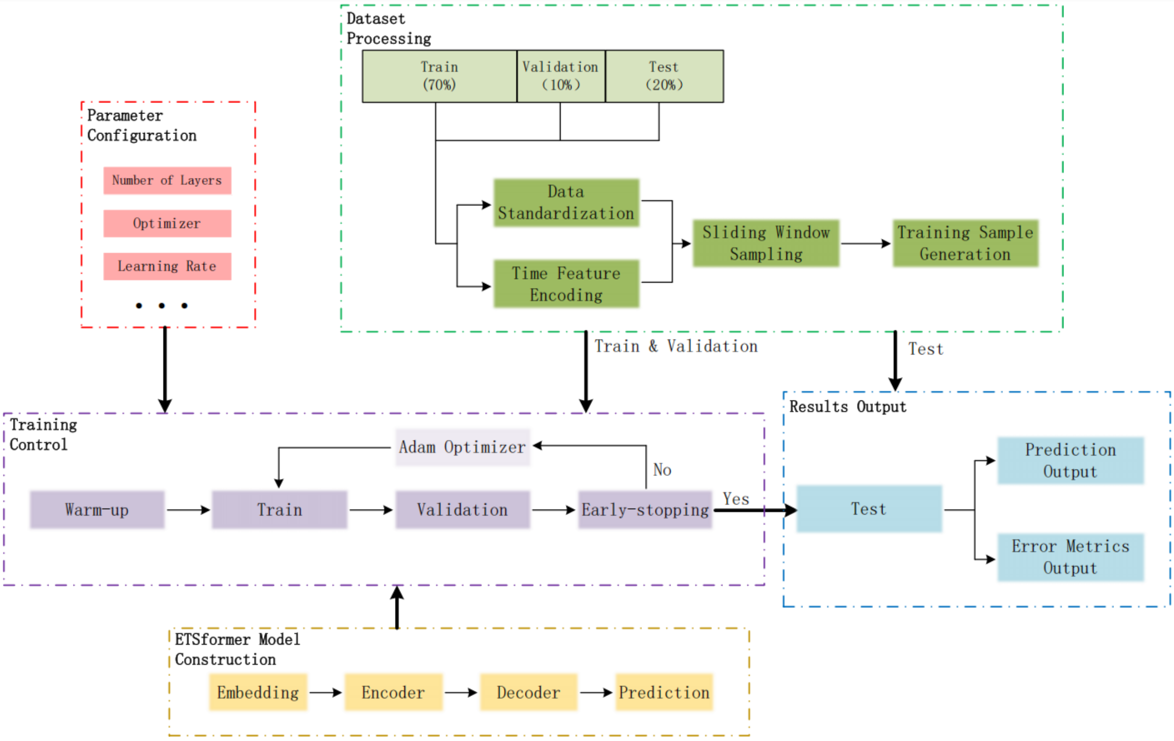

Figure 3.1: Training Process Flowchart

recorded for subsequent analysis and visualization.

After the model training is completed, evaluation is conducted on the test set. The predicted results are output, and several error metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Mean Squared Percentage Error (MSPE), are calculated to comprehensively reflect the model's prediction accuracy and generalization performance in practical tasks.

3.2 Dataset Description and Preprocessing

The air quality dataset used in this study was collected from a classroom (room001) located in the Drakos building at Cyprus University of Technology. This classroom is equipped with high-precision environmental sensors capable of real-time monitoring of several key indoor air quality parameters. The sensor system has been continuously operating 24 hours a day from September 1, 2023, to February 11, 2025, ensuring the completeness and continuity of the time series data. All collected data are recorded at fixed time intervals and stored in CSV file format, which facilitates subsequent data loading and processing.

The monitored air quality parameters include temperature, humidity, CO_2 concentration, and PM2.5 concentration (particulate matter). Each data record contains a timestamp field, which identifies the specific sampling time. These parameters cover the major physical and chemical factors affecting indoor air quality and exhibit certain periodic variation patterns, reflecting the dynamic processes influenced by human activities, ventilation conditions, and external environmental changes.

Considering the differences in physical significance and fluctuation patterns of various air quality indicators, this study constructs four independent datasets for each monitored parameter. Specifically, separate training datasets are established for temperature, humidity, CO_2 concentration, and PM2.5 concentration, with each dataset focusing on a single variable as the prediction target. This design forms univariate time-series forecasting tasks, which not only facilitate an in-depth evaluation of the ETSformer model's predictive performance across different air quality features but also avoid potential interference between variables, thereby improving the focus and generalization ability of the training process.

Through the above data collection strategy, this study obtained a high-resolution, multi-feature indoor air quality time-series dataset, providing a reliable foundation for subsequent model training and evaluation.

In the design of the prediction tasks, univariate time-series forecasting is adopted for each air quality indicator. The historical observations of a single target variable serve as the input to predict its future values over several time steps. Four independent datasets are constructed with target variables including temperature, humidity, CO_2 concentration, and PM2.5 concentration, ensuring the mutual independence of each forecasting task during data preparation and model training.

During the data loading and preprocessing stage, the forecasting task type is set as univariate prediction (`features = 'S'`), where both the input features and the prediction target belong to the same variable. Specifically, in each task, the model takes the historical sequence of the target feature within a defined time window as input and outputs the predicted values of the same feature for future time steps. The parameter `target = 'tem'` is used to specify temperature as one of the prediction targets, while humidity, CO_2 concentration, and PM2.5 concentration are designated by `target = 'hum'`, `target = 'CO2'`, and `target = 'pm2.5'`, respectively, for independent modeling.

This univariate prediction task design allows for a clear evaluation of the ETSformer model's fitting ability and forecasting performance across different air quality indicators. At the same time, it avoids feature redundancy or variable coupling issues that may arise in multivariate tasks, making the results of each task more controllable and the evaluation criteria more consistent.

In the preprocessing stage, the original time-series data are split into training, validation, and testing sets using a fixed ratio of 70%, 10%, and 20%, respectively. The splitting boundaries are determined based on chronological order to ensure that the validation and testing data strictly follow the training data, preventing future information leakage and complying with the basic assumptions of time-series forecasting tasks.

To improve the stability of model training and accelerate convergence, the input features of each dataset are standardized. The standardization parameters, including the mean and standard deviation, are calculated based on the training set using the `StandardScaler` tool. The same transformation is then applied to the validation and testing sets, ensuring the rationality of the data processing procedure and the generalization ability of the model.

Through this data splitting and standardization approach, this study ensures sufficient utilization of the training data while properly isolating the validation and testing data, providing a reliable foundation for performance evaluation and generalization assessment during the model training process.

In time-series forecasting tasks, timestamp information often contains rich periodic patterns, such as hourly, weekly, and monthly cycles, which may significantly influence the prediction target. To fully exploit the periodic features embedded in the time information, this study introduces a time feature embedding mechanism during the data preprocessing stage to enhance the model's ability to perceive temporal dynamics.

Specifically, a frequency-based time feature encoding method is adopted during data loading. This method extracts multiple time-related features from each timestamp, including hour of the day, day of the week, day of the month, and month of the year. These time features are embedded as continuous numerical inputs and combined with the original observation data to serve as the input features for the ETSformer model.

Given that the data collection scenario is situated in a university classroom, the indoor environmental parameters (such as temperature, humidity, CO_2 concentration, and PM2.5 concentration) are influenced not only by daily diurnal cycles but also by factors such as academic schedules, seasonal variations, and differences between weekdays and weekends. For example, fluctuations in occupancy levels at different times, seasonal climate changes, and usage patterns between working days and weekends may lead to complex periodic and non-stationary variations in air quality indicators. Therefore, the refined time feature encoding design enables the model to effectively capture these potential fluctuation patterns across multiple time scales, enhancing the modeling capacity and prediction accuracy for air quality dynamics.

The introduction of time feature embeddings not only strengthens the model's ability to characterize temporal dependencies but also provides an improved input representation for the ETSformer model when handling indoor air quality prediction tasks with strong periodicity and multi-scale temporal features.

3.3 Model Architecture: ETSformer

This study adopts a time-series forecasting model based on the ETSformer architecture, which enhances the traditional Transformer structure by introducing a time-series decomposition mechanism. The model explicitly decomposes the input sequence into three components: trend (growth), seasonality, and level, thereby improving its capability to model complex temporal data. The design of ETSformer is inspired by the classical Exponential Smoothing (ETS) method and integrates the powerful representation ability of deep learning, achieving a balance between forecasting stability and model interpretability.

In the field of time-series analysis, ETS decomposition models commonly use the following additive formulation to represent time-series signals:

$$y_t = l_t + s_t + g_t, \tag{3.1}$$

where y_t denotes the observed value, l_t represents the level component, s_t denotes the seasonality component, and g_t is the trend (growth) component. Based on this formulation, ETSformer implements a learnable decomposition process through deep learning and integrates it into the Transformer architecture, combining strong representation power with physical interpretability.

The overall architecture of ETSformer consists of an input embedding layer, stacked encoders, decoders, and a linear output layer. The encoder is responsible for decomposing the input historical sequence into trend and seasonal components. During the prediction phase, the decoder introduces a damping control mechanism to smooth the extrapolation of the trend component, preventing excessive amplification or oscillation of the trend during forecasting.

In particular, ETSformer integrates a Fourier transform-based seasonality modeling module and an ex-

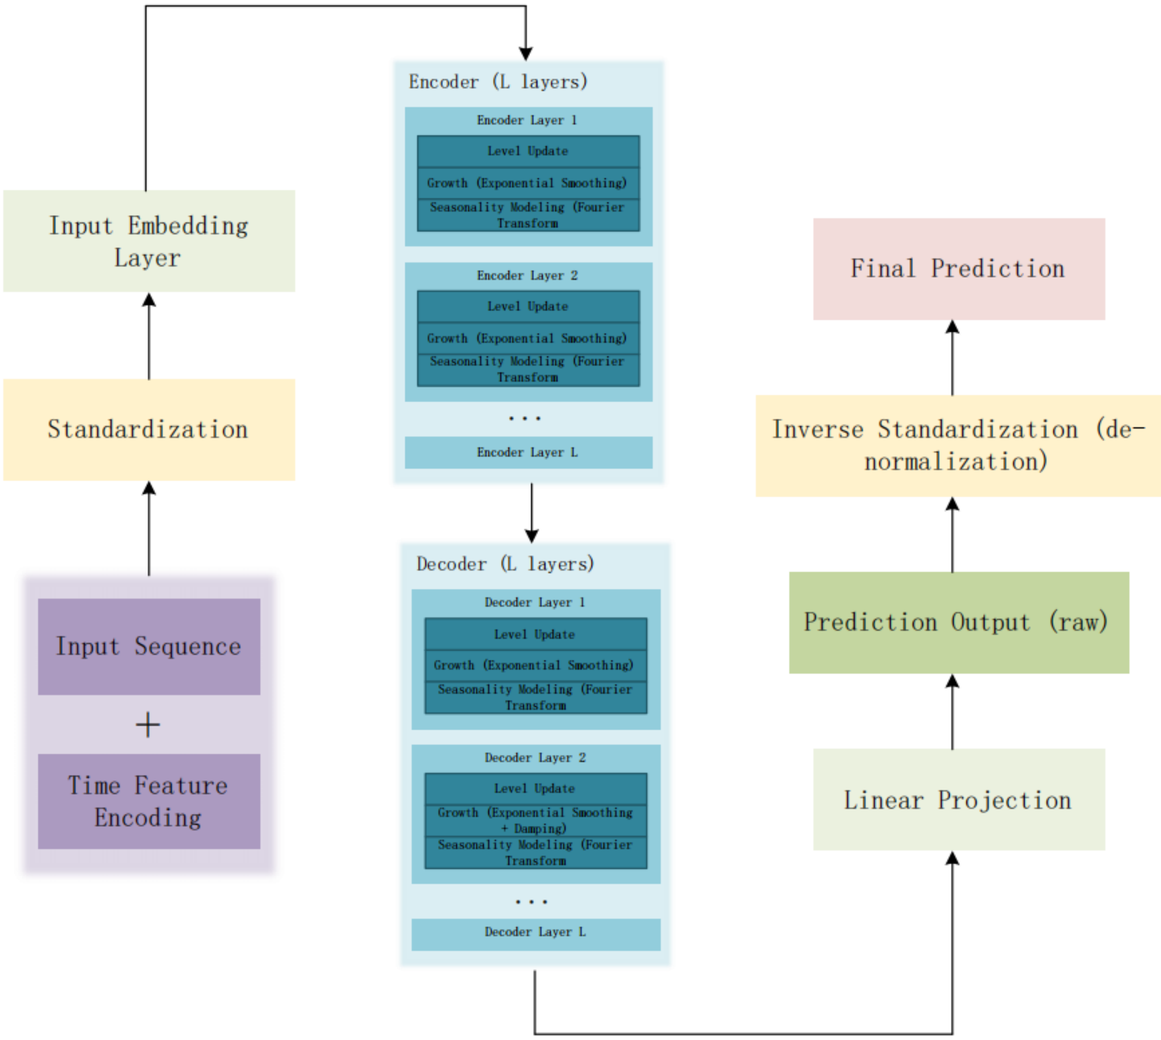

Figure 3.2: ETSformer Model Architecture Diagram

ponential smoothing-based growth trend modeling module within each encoder layer. This design effectively captures the periodic fluctuations and non-stationary changes present in time-series data.

Such an architecture not only enhances the model's capability in modeling periodic and trending signals but also improves its robustness against abnormal fluctuations and complex environmental disturbances. By maintaining the strong learning capacity of the Transformer's self-attention mechanism while incorporating principles from classical time-series analysis, ETSformer achieves a balance between model performance and interpretability.

In the ETSformer model, the input embedding layer is responsible for mapping the raw time-series data into a high-dimensional feature space, thereby enhancing the model's capability to represent the input sequence. The design of this embedding layer is based on one-dimensional convolution (Conv1D) operations, followed by a Dropout layer to prevent overfitting and improve the generalization performance of the model.

Specifically, the input embedding layer applies sliding convolutions along the temporal dimension of the sequence using one-dimensional convolution kernels. This operation transforms the single-dimensional input features (such as temperature, humidity, CO_2 concentration, or PM2.5 concentration) into high-dimensional vectors with a feature dimension of d_{model}. This process not only expands the expressive capacity of the input features but also introduces local receptive fields, enabling the model to better capture local variation patterns during the encoder stage. The convolution layer weights are initialized using the Kaiming Normal distribution to ensure stable gradient propagation at the early stage of training.

After the convolution operation, a Dropout layer is connected to randomly deactivate a portion of neuron activations, aiming to suppress overfitting. The dropout rate can be adjusted as a hyperparameter and is set to 0.3 in this study. This design further enhances the robustness of the model during the training process.

Through the embedding mechanism described above, the ETSformer model performs sufficient feature mapping and regularization before the input sequence enters the encoder module, providing a solid feature foundation for the subsequent trend decomposition and seasonality modeling.

The encoder, as a core component of the ETSformer model, is responsible for feature extraction and trend decomposition of the input sequence. This part adopts a multi-layer stacked structure, where each layer contains three functional modules: seasonality modeling, growth trend modeling with exponential smoothing, and level updating, which correspond to the FourierLayer, GrowthLayer (integrated with Exponential Smoothing), and LevelLayer, respectively. This design not only preserves the powerful learning capability of the multi-head attention mechanism in the traditional Transformer but also incorporates the idea of trend decomposition from time-series analysis, achieving an effective integration of physical modeling and deep learning.

Within the encoder of the ETSformer model, the periodic components of the input sequence are first extracted using the FourierLayer. This module is based on the Fast Fourier Transform (FFT) and selects the principal frequency components for modeling, effectively capturing the seasonal variation patterns in the time-series data. Specifically, the extraction of the seasonality component can be expressed as:

$$s_t = \sum_{k=1}^{K} A_k \cos\left(2\pi f_k t + \phi_k\right), \tag{3.2}$$

where A_k represents the amplitude, f_k denotes the frequency, ϕ_k is the phase, and K is the number of selected principal frequency components. This design allows the model to flexibly adapt to multi-scale fluctuations in air quality data, such as daily cycles and seasonal changes.

The extracted seasonality component is subsequently used to correct the residual part of the original input sequence. After this, the GrowthLayer, combined with the exponential smoothing mechanism, models the growth trend of the residual component. The exponential smoothing layer applies a recursive weighted averaging process to the historical growth changes, which captures long-term trend variations while suppressing short-term fluctuations, thus enhancing the model's adaptability to non-stationary time-series data. Additionally, this part incorporates a multi-head mechanism to improve the trend modeling capability across different subspaces. The recursive formula for the trend component is defined as:

$$\hat{g}_t = \alpha \cdot r_t + (1 - \alpha) \cdot \hat{g}_{t-1}, \tag{3.3}$$

where $r_t = y_t - s_t$ represents the deseasonalized residual, and α is the smoothing factor, which is learned by the neural network and constrained within a reasonable range using the sigmoid activation function.

Finally, the LevelLayer serves as the level updating module in each encoder layer, dynamically adjusting the level component based on the current outputs of the seasonality and trend components. This design ensures a proper balance among the seasonality, trend, and level components, preventing any single component from dominating the overall prediction results.

Through the stacking of multiple encoder layers and the intra-layer trend decomposition mechanism, ETSformer achieves joint modeling of periodicity, trend, and level components while extracting multi-scale features from time-series data. This structural design provides a solid feature foundation for the prediction extension performed by the decoder and significantly enhances the model's predictive capability and generalization performance on complex air quality data.

In the ETSformer model, the decoder is mainly responsible for extending the sequence prediction to future time steps. The design of the decoder inherits the trend decomposition approach from the encoder and introduces a damping mechanism during the extrapolation process to suppress the excessive growth of the trend component in long-term forecasting, thereby ensuring the stability and rationality of the prediction results. The damping mechanism controls the extrapolation of the trend component using the following decay formula:

$$\tilde{g}_{t+h} = \hat{g}_t \cdot \lambda^h, \tag{3.4}$$

where $\lambda \in (0, 1)$ is the damping factor, which is learned adaptively during the training process, and h denotes the prediction horizon. This design effectively prevents the uncontrolled amplification of the trend component in long-term predictions and ensures the stability of the model output.

Meanwhile, the seasonality component in the decoder directly inherits the periodic modeling results from the encoder and is truncated and extended according to the prediction horizon. The decoder processes the growth trend component and the seasonality component separately and then maps them to the final output space through linear projection. These outputs are combined with the level component generated by the encoder to form the final prediction result. By jointly considering the trend, seasonality, and level components, the final output of the ETSformer model can be expressed as:

$$\hat{y}_{t+h} = l_t + \tilde{g}_{t+h} + s_{t+h}, \tag{3.5}$$

which preserves the original ETS decomposition idea while enhancing the structural interpretability of the model.

Firstly, the encoder of ETSformer integrates the multi-head attention mechanism with the trend modeling module, enabling adaptive modeling of abnormal fluctuations. The multi-head attention mechanism captures sequence features across different subspaces, and its attention weights are calculated as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V, \tag{3.6}$$

where Q, K, and V denote the query, key, and value matrices, respectively, and d_k represents the dimensionality of the key. The multi-head mechanism applies this computation in parallel across multiple subspaces, enhancing the learning capacity across different feature dimensions. At the same time, the exponential smoothing mechanism suppresses short-term noise, improving the model's performance on non-stationary time-series data. This design prevents excessive sensitivity to outliers and enhances the robustness of the prediction results.

Secondly, the decoder introduces a damping factor in the extrapolation of the trend component to control the growth process and prevent uncontrolled amplification of the trend during long-term forecasting. The damping factor is a learnable parameter that is adaptively adjusted during training according to the data characteristics, ensuring the stability and rationality of prediction extension. Specifically, the recursive relationship for the multi-step extrapolation of the trend component is given by:

$$\tilde{g}_{t+h} = \tilde{g}_{t+(h-1)} \cdot \lambda, \tag{3.7}$$

where $\lambda \in (0, 1)$ is the damping factor that controls the growth rate during extrapolation and ensures the gradual attenuation of the trend component.

In addition, the seasonality modeling part employs the FourierLayer to extract periodic features directly in the frequency domain. Compared with window-based periodic modeling methods, this approach provides higher modeling flexibility and efficiency. This design enables the model to flexibly capture the periodic fluctuations in air quality data caused by various factors such as diurnal cycles, seasonal changes, and academic activity schedules. Specifically, the Fourier transform process is defined as:

$$X(f) = \sum_{t=0}^{T-1} x_t \, e^{-j2\pi ft/T}, \tag{3.8}$$

where $X(f)$ denotes the frequency domain representation, x_t is the time-series input, and T is the sequence length. By selecting the principal frequency components for reconstruction, the model effectively captures the seasonal characteristics of the sequence.

Finally, the overall architecture of ETSformer balances the flexibility of deep learning models with the interpretability of time-series decomposition methods. The decomposition of the trend, seasonality, and level components not only enhances the model's predictive performance but also provides explicit physical meaning for subsequent result analysis and anomaly detection. These design advantages enable ETSformer to achieve stable, accurate, and generalizable performance in air quality prediction tasks across multiple time scales and complex environments.

In summary, ETSformer enhances its forecasting capability while ensuring structural interpretability and training stability through the introduction of trend decomposition, seasonality modeling, and damping mechanisms within the encoder-decoder framework. This architecture effectively adapts to the trend variations, periodic fluctuations, and abnormal disturbances present in air quality prediction tasks, demonstrating strong generalization ability and application flexibility. Based on this design, ETSformer provides a solid model foundation for the subsequent training configuration, experimental process, and performance evaluation.

3.4 Training Configuration and Hyperparameters

In the model training process, reasonable hyperparameter configuration plays a crucial role in improving prediction performance, accelerating convergence, and preventing overfitting. In this study, a set of training hyperparameters is carefully designed and configured for the ETSformer model based on the task requirements and structural characteristics of the model. These settings cover structural parameters of the model, input and output sequence lengths, optimizer selection, learning rate scheduling strategies, training control mechanisms, and data augmentation methods. This configuration scheme not only ensures the adaptability and stability of the model across different air quality prediction tasks but also provides a consistent training environment and comparison baseline for subsequent experimental evaluation.

This section introduces the configuration details of each component mentioned above, including model structure parameters, input and output data formats, optimizer selection and learning rate adjustment strategies, training process control mechanisms, and data augmentation designs. These elements together support the effective training of the ETSformer model and the enhancement of its performance in multi-target air quality forecasting tasks.

In terms of model architecture design, ETSformer adopts a multi-layer stacked encoder-decoder structure to improve its capability for feature extraction and representation in complex time-series modeling scenarios. The core hyperparameters include the model hidden dimension (d_{model}), the number of attention heads (n_{heads}), the number of encoder layers (e_{layers}), the number of decoder layers (d_{layers}), and the number of selected frequency components in the Fourier transform module (K). In the experiments conducted in this study, these parameters are set as follows: $d_{\text{model}} = 256$, $n_{\text{heads}} = 4$, $e_{\text{layers}} = d_{\text{layers}} = 3$, and the Top-$K = 1$ principal frequency component is selected in the Fourier layer.

Among these parameters, d_{model} determines the representation capacity of each time step in the high-dimensional feature space—the higher the dimension, the richer the features the model can learn, though it also increases computational complexity. The parameter n_{heads} specifies the number of parallel attention heads in the multi-head attention mechanism, which allows the model to independently learn sequence features in different subspaces, thereby enhancing its ability to capture complex dynamic variations. Specifically, ETSformer applies the following attention computation in each attention head to perform feature weighting over the input sequence:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V, \tag{3.9}$$

where Q, K, and V represent the query, key, and value matrices, respectively, and d_k denotes the dimensionality of each attention head. The multi-head attention mechanism applies n_{heads} parallel subspaces, enabling the model to capture and fuse features across different dimensions effectively.

The parameters e_{layers} and d_{layers} correspond to the number of encoder and decoder layers, respectively. Increasing the number of layers helps the model extract multi-scale features from time-series data at deeper levels. The parameter K controls the number of principal frequencies selected in the FourierLayer for seasonal component modeling. A properly configured K value facilitates effective extraction of seasonality while avoiding interference from high-frequency noise, thus improving the model's performance.

Specifically, the FourierLayer for seasonality modeling performs frequency-domain decomposition of the time-series input using the following expression:

$$X(f) = \sum_{t=0}^{T-1} x_t \, e^{-j2\pi f t/T}, \tag{3.10}$$

where $X(f)$ denotes the frequency-domain representation, x_t is the input time-series signal, and T is the sequence length.

$$s_t = \sum_{k=1}^{K} A_k \cos\left(2\pi f_k t + \phi_k\right), \tag{3.11}$$

where A_k represents the amplitude, f_k is the frequency, and ϕ_k is the phase. This design enables the model to focus on the primary periodic components, thereby enhancing the effectiveness of seasonality modeling.

The above configuration of structural parameters comprehensively considers the model complexity, feature dimensionality of the prediction tasks, and training efficiency, providing a stable modeling foundation for ETSformer.

In time-series forecasting tasks, the configuration of the input sequence length and prediction sequence length directly affects the model's performance. If the input window is too short, it may fail to fully utilize historical information; conversely, if it is too long, it may introduce irrelevant noise and increase the learning burden of the model. Based on the time resolution and variation characteristics of the air quality monitoring data, this study reasonably configures the input and output sequence lengths to balance the utilization of historical information and the timeliness of forecasting. Specifically, the input and output window design of ETSformer satisfies the following relationship:

$$\mathbf{X} = \{x_{t-\text{seq\_len}+1}, \ldots, x_t\}, \quad \mathbf{Y} = \{y_{t+1}, \ldots, y_{t+\text{pred\_len}}\}, \tag{3.12}$$

where **X** represents the input historical window, and **Y** denotes the prediction target window. In this study, the input sequence length (`seq_len` = 60), label length (`label_len` = 0), and prediction sequence length (`pred_len` = 12) are fixed for the multi-step forecasting tasks of air quality target variables.

During the model optimization process, this study employs a custom-designed Adam optimizer to update the parameters of ETSformer. Unlike the standard single-parameter group configuration, the parameters of ETSformer are divided into three groups according to their functional differences: main network parameters (`nn`), smoothing weights (`smoothing`), and damping factors (`damping`). Separate learning rates are assigned to each parameter group. This multi-parameter-group optimization strategy enables flexible step-size adjustments according to the learning requirements of each submodule, thereby enhancing the overall training effectiveness.

Specifically, the main network parameters are updated using the base learning rate (`learning_rate`), while the smoothing weights and damping factors are assigned learning rates set to 100 times the base learning rate. This design considers the direct impact of smoothing weights and damping factors on model stability, requiring faster convergence during the early stages of training to prevent excessive fluctuations in the trend and seasonality components at initialization. The optimizer implementation is based on a custom Adam class, retaining the classical momentum mechanism and second-moment estimation of the original Adam optimizer, ensuring smooth and stable parameter updates.

The parameter update process follows the standard Adam optimization formulas:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \cdot \frac{m_t}{\sqrt{v_t} + \epsilon}, \tag{3.13}$$

where $\theta^{(t)}$ denotes the model parameters at the current iteration, η is the learning rate, m_t represents the first-order moment (the mean of gradients), v_t represents the second-order moment (the mean of squared gradients), and ϵ is a small constant added to prevent division by zero.

For learning rate scheduling, this study adopts the `exponential_with_warmup` strategy. In this approach, a certain number of warm-up epochs are set at the early stage of training, during which the learning rate increases linearly from a predefined minimum learning rate (`min_lr`) to the base learning rate. This design helps mitigate the instability caused by gradient oscillations at the beginning of training. After the warm-up phase, the learning rate decays exponentially, which enables fine-tuning of the model parameters during the later training stages and promotes stable convergence while making full use of the training data. The specific formula for learning rate adjustment is given as follows:

$$\eta_t = \begin{cases} \eta_{\text{min}} + \frac{t}{T_{\text{warmup}}} \left(\eta_0 - \eta_{\text{min}}\right), & t \le T_{\text{warmup}}, \\ \eta_0 \cdot \gamma^{(t - T_{\text{warmup}})}, & t > T_{\text{warmup}}, \end{cases} \tag{3.14}$$

where η_t denotes the learning rate at the t-th iteration, η_0 is the base learning rate, η_{min} is the minimum learning rate, γ is the exponential decay factor (typically set to 0.5), and T_{warmup} is the number of warm-up iterations.

The combination of the multi-parameter-group optimizer design with the exponential decay learning rate adjustment mechanism and warm-up scheduling ensures stable convergence and efficient learning during the training of ETSformer. This approach effectively accommodates the differentiated update speed requirements of various model components.

During the training process, to further enhance model stability and prevent overfitting, this study introduces the EarlyStopping strategy along with a data augmentation mechanism. EarlyStopping, as a commonly used training control method, monitors the changes in validation loss over consecutive epochs to determine whether the model has reached a performance saturation stage. If the validation loss does not improve within a predefined number of patience epochs, the training process is terminated early to avoid performance degradation caused by overfitting. In this study, the patience value for EarlyStopping is set to 10, meaning that if the validation loss fails to decrease for 10 consecutive training epochs, training is stopped and the best-performing model parameters are saved.

Specifically, the decision logic of EarlyStopping can be described as follows: let the validation loss at the t-th epoch be $L(t)$. Training is terminated early when the following condition is satisfied:

$$L(t) > L(t - p), \quad \text{for} \quad p = 1, 2, \ldots, \text{patience}, \tag{3.15}$$

where $L(t)$ represents the validation loss at the t-th epoch, and `patience` refers to the number of tolerated waiting epochs (set to 10 in this study). This mechanism ensures training efficiency while effectively suppressing the risk of overfitting that may occur in the later stages of training.

Regarding data augmentation, considering that real-world air quality monitoring data may contain noise disturbances and abnormal fluctuations, this study introduces standard deviation perturbation (Standard Deviation Perturbation) during the training phase. Specifically, the `Transform` operation with `sigma = 0.2` is applied for random perturbation-based augmentation. This method enhances the diversity of training samples by applying random jittering, scaling, and shifting to the input data, while maintaining the overall data distribution. This design improves the model's robustness to input noise and outliers. The perturbation process can be expressed as:

$$x'_t = (x_t \cdot (1 + \epsilon_{\text{scale}})) + \epsilon_{\text{shift}} + \epsilon_{\text{jitter}}, \tag{3.16}$$

where x_t is the original input sample, x'_t is the augmented sample, and ϵ_{scale}, ϵ_{shift}, and ϵ_{jitter} represent the scaling perturbation, shifting perturbation, and jittering perturbation, respectively. All these perturbation factors follow a normal distribution with zero mean and variance σ^2, expressed as:

$$\epsilon_{\text{scale}},\ \epsilon_{\text{shift}},\ \epsilon_{\text{jitter}} \sim \mathcal{N}(0, \sigma^2). \tag{3.17}$$

In this study, the perturbation strength parameter σ is set to 0.2. This design enhances the diversity and generalization ability of model training without disrupting the primary distribution characteristics of the data.

The combined use of EarlyStopping and data augmentation ensures training efficiency while significantly improving the generalization capability of ETSformer in air quality prediction tasks. This approach reduces the model's sensitivity to outliers and noisy data, enhancing the stability and reliability of the training process.

In summary, this study has systematically configured and designed the training process of ETSformer across several key aspects, including model structural parameters, input-output sequence lengths, optimizer selection, learning rate adjustment strategies, training process control mechanisms, and data augmentation methods. These configurations ensure that the ETSformer model achieves efficient, stable training with strong generalization performance in air quality forecasting tasks. Through the multi-parameter-group optimization strategy and dynamic learning rate scheduling, coordinated training across different modules of the model is realized. Additionally, the introduction of EarlyStopping and standard deviation perturbation-based augmentation effectively suppresses the risk of overfitting and improves the model's adaptability to complex environmental data.

In the multi-parameter-group optimization strategy, differentiated learning rates are set for the main network parameters, smoothing weights, and damping factors. This configuration allows each module to flexibly adjust its update speed according to its own learning requirements, further ensuring the stability and efficiency of the model training. The dynamic learning rate scheduling is implemented through the `exponential_with_warmup` strategy, where the learning rate increases linearly during the warm-up phase to mitigate gradient oscillations at the early stage, and then decays exponentially in the later phase to precisely control the learning rate, enhancing the training effect during convergence.

Regarding training process control, this study applies the EarlyStopping strategy to dynamically deter-

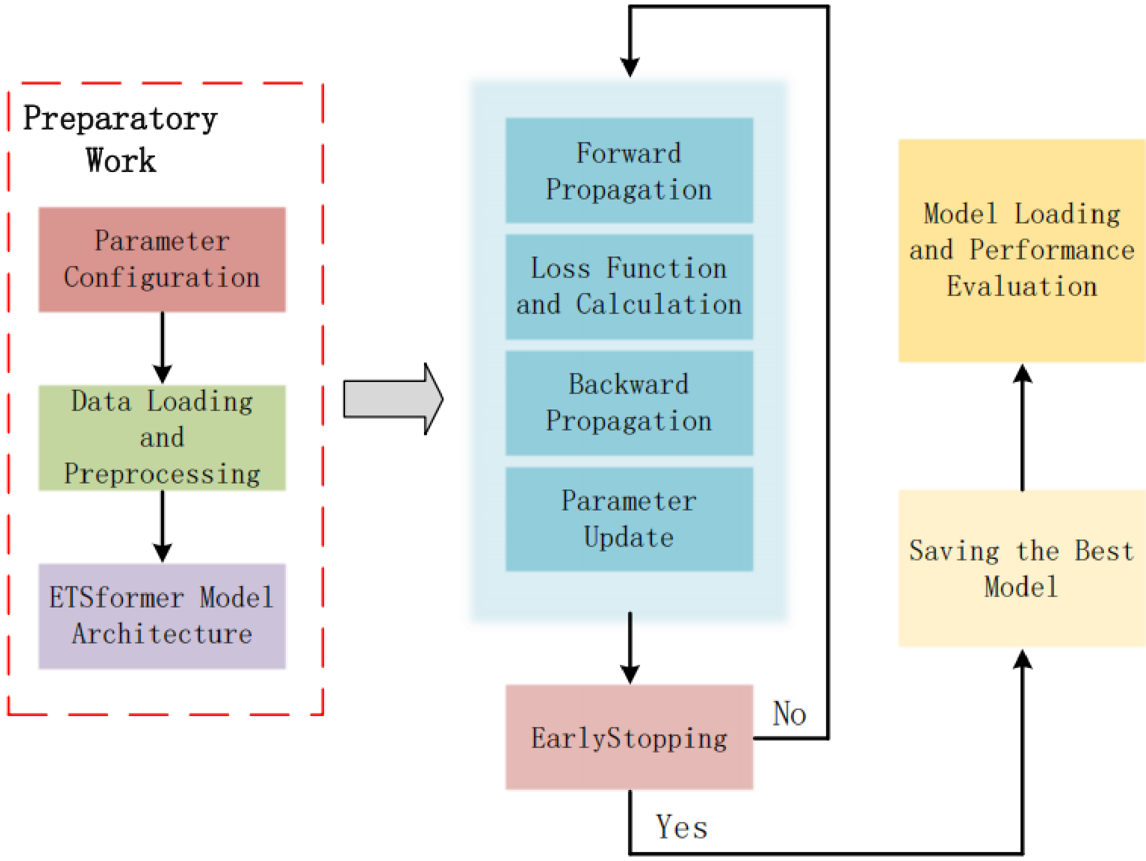

Figure 3.3: Training process flowchart of the ETSformer model

mine whether the training has reached performance saturation based on the validation loss. Training is terminated promptly when the model shows no improvement, preventing overfitting. The patience parameter is set to 10, ensuring that the model is sufficiently trained within a reasonable range while avoiding performance degradation caused by overly prolonged training.

For data augmentation, this study introduces a standard deviation perturbation-based strategy, where random perturbations are applied to the input data during training, including jittering, scaling, and shifting. The perturbation strength σ is set to 0.2. This design increases the diversity of the training samples without altering the overall data distribution, improving the model's robustness to outliers and noise, and enhancing its generalization capability under complex air quality conditions.

The above training configuration lays a solid foundation for the subsequent training process, experimental analysis, and performance evaluation, ensuring the scientific validity and reproducibility of the experimental results. Through reasonable parameter configuration and optimization design, ETSformer achieves a good balance among accuracy, stability, and training efficiency in multi-target air quality prediction tasks, providing strong support for efficient and intelligent air quality management. During the training phase, this study selects the Mean Squared Error (MSE) as the loss function to measure the deviation between the predicted values and the ground truth. The MSE loss function imposes a higher penalty on large prediction errors, making it well-suited for regression-based time-series forecasting tasks. It effectively guides the model to optimize prediction accuracy throughout the training process. The calculation formula for the MSE loss is defined as follows:

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2, \tag{3.18}$$

where N represents the total number of prediction samples, y_i denotes the true value, and $\hat{y}_i$ is the predicted value.

During the training process, both the training loss and validation loss are recorded simultaneously to monitor the convergence status of the model. By observing the loss curves of these two groups, it is possible to assess the fitting performance and identify potential overfitting issues, providing a reference for subsequent training control and evaluation.

To comprehensively evaluate the prediction performance of the ETSformer model during the testing phase, this study employs multiple error metrics for quantitative analysis of the forecasting results. Specifically, the evaluation includes the Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Mean Squared Percentage Error (MSPE). The calculation formulas for these metrics are defined as follows:

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i|, \tag{3.19}$$

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2, \tag{3.20}$$

$$\text{RMSE} = \sqrt{\text{MSE}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}, \tag{3.21}$$

$$\text{MAPE} = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \times 100\%, \tag{3.22}$$

$$\text{MSPE} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{y_i - \hat{y}_i}{y_i}\right)^2, \tag{3.23}$$

where N represents the total number of prediction samples, y_i denotes the true value, and $\hat{y}_i$ is the predicted value.

The MAE reflects the overall deviation level of the model in the prediction task, while the MSE and RMSE emphasize penalizing larger errors. RMSE has the same dimensionality as the original data, which facilitates the interpretation of the results. MAPE and MSPE measure prediction errors in relative terms, making them suitable for evaluating forecasting performance across different numerical scales. The combined use of multiple metrics provides a comprehensive evaluation of the model's predictive capability and stability from various perspectives.

To visually illustrate the convergence process and prediction performance of the model, this study adopts two result visualization approaches during the training and testing stages. First, the training and validation loss curves, as shown in Figure 3.4(a), are plotted to observe the convergence trend and evaluate the effect of early stopping. This process is automatically implemented using the `plot_loss_curve()` function in the training script, which helps determine whether the model is experiencing overfitting or underfitting.

Second, during the testing phase, the predicted values are compared with the ground truth through visual plots, as illustrated in Figure **??**(b), to directly demonstrate the fitting performance of the predictions. This visualization is achieved using the `visual()` function, which plots both the prediction curve and the actual observation curve based on the test set samples, reflecting the model's forecasting accuracy across different time steps.

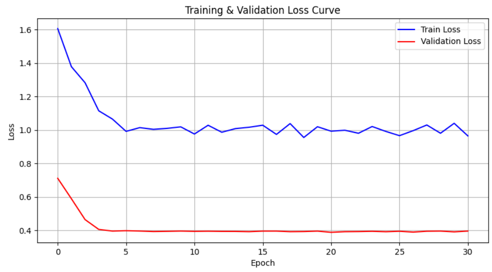

(a) Training and validation loss curves

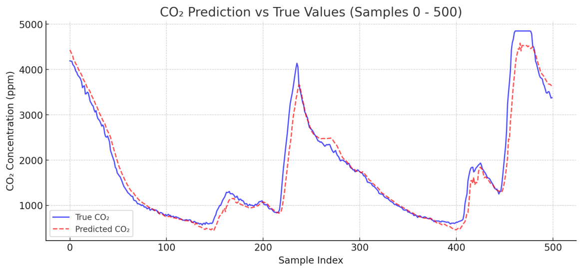

(b) Comparison between predicted results and ground truth on the test set

Figure 3.4: (a) Training and validation loss curves; (b) predicted results versus ground truth for the CO_2 concentration prediction task.

In summary, this study integrates a well-designed training process, appropriate loss function selection, comprehensive evaluation metrics, and effective result visualization methods to ensure the validity of the ETSformer model training and the verifiability of its forecasting performance. Through multi-perspective performance monitoring and result analysis, a solid foundation is established for the subsequent presentation of experimental results and comparative analysis of model performance.

3.5 Advantage Analysis of the Training Strategy

In response to the characteristics of air quality prediction tasks, including trend variations, periodic fluctuations, and abnormal disturbances, this study incorporates various optimizations and innovations in both model architecture and training strategy design. ETSformer integrates trend decomposition, seasonality modeling, and smoothing damping control, which not only enhances the model's adaptability to complex dynamic changes at the architectural level but also improves its stability and generalization performance through multi-parameter-group optimization, learning rate scheduling, data augmentation, and early stopping mechanisms in the training process.

This section analyzes and summarizes the advantages of the proposed training strategy from three perspectives: adaptive anomaly handling capability, model interpretability, and generalization ability.

In air quality monitoring data, abnormal fluctuations or noisy samples often occur due to unexpected events, environmental changes, or sensor errors. Traditional time-series forecasting methods typically rely on manual data preprocessing steps, such as anomaly removal or smoothing, to reduce the impact of such disturbances on the prediction results. However, these approaches are often based on empirical judgment and may compromise the balance between anomaly detection accuracy and data utilization. Moreover, important information might be lost during the preprocessing stage.

In contrast, ETSformer incorporates an internal anomaly adaptation mechanism through the integration of the Exponential Smoothing mechanism and the Damping Control in its structural design. The GrowthLayer within the encoder models the trend variations through a recursive smoothing process, effectively suppressing the amplification of sudden anomalies. Meanwhile, the DampingLayer in the decoder applies attenuation control on the extrapolated trend components during the prediction phase, preventing abnormal values from excessively influencing future forecasts.

This structural design enables the model to automatically regulate its response to abnormal fluctuations during training, without the need for additional anomaly filtering steps. As a result, the model achieves inherent suppression of anomalies and enhances the robustness and stability of the overall forecasting process.

Beyond achieving strong predictive performance, ETSformer also emphasizes model interpretability. The encoder and decoder utilize a multi-head attention mechanism for feature extraction, offering the potential for attention weight visualization. By analyzing the attention distribution, it is possible to identify which historical time steps the model focuses on during its prediction process, providing insight into the decision-making patterns of the model.

Compared with standard Transformer architectures that rely solely on attention distributions for interpretability, ETSformer further enhances model explainability by introducing decomposition into trend, seasonality, and level components. This design not only strengthens the model's visualization capability at the attention level but also provides physically meaningful component-based explanations through the outputs of the growth, seasonality, and level terms. This facilitates the analysis of the contribution of different variation sources to the final prediction results. This integration of attention visualization and trend decomposition within the interpretability framework enhances the usability and transparency of ETSformer in practical applications. The model is not only capable of producing accurate prediction results but also provides reasonable interpretative support for decision-making processes.

In time-series forecasting tasks, achieving high accuracy on the training set does not necessarily guarantee good performance on unseen data. To improve the generalization capability of the model, ETSformer incorporates smoothing mechanisms and damping control into its architectural design, enhancing its stability when handling various time scales and different air quality indicators.

Specifically, the GrowthLayer within the encoder utilizes exponential smoothing for recursive trend modeling, effectively suppressing high-frequency noise and short-term abnormal fluctuations. This ensures stable learning of long-term trend variations. In addition, the DampingLayer in the decoder introduces damping factors during the trend extrapolation phase to dynamically control the growth rate, preventing unreasonable amplification or oscillation of the trend components during long-term forecasting.

This collaborative design of smoothing and damping not only improves convergence efficiency during training but also demonstrates stronger generalization performance during testing. ETSformer maintains stable forecasting accuracy across different air quality indicators—including temperature, humidity, CO_2 concentration, and PM2.5 concentration—and adapts to diverse data characteristics with varying degrees of fluctuation, reducing the risk of overfitting.

In summary, through the incorporation of trend decomposition, exponential smoothing, damping control, and multi-head attention mechanisms in both model architecture and training strategy, ETSformer achieves multiple advantages, including adaptive anomaly handling, enhanced interpretability, and improved generalization capability. Unlike traditional methods that rely on manual data cleaning or static preprocessing, ETSformer provides intrinsic adaptive responses to abnormal fluctuations while preserving the physical interpretability of the prediction components, thereby enhancing its practical value in real-world air quality forecasting tasks.

These advantages enable ETSformer to balance forecasting accuracy, stability, and interpretability when dealing with complex, dynamically changing time-series data, offering reliable technical support for multi-scenario environmental monitoring and forecasting applications.

3.6 Advantage Analysis of the Training Strategy

In response to the characteristics of air quality prediction tasks, including trend variations, periodic fluctuations, and abnormal disturbances, this study incorporates multiple optimizations and innovations in both model architecture and training strategy design. ETSformer integrates trend decomposition, seasonality modeling, and smoothing damping control, which not only enhances the model's adaptability to complex dynamic changes at the architectural level but also improves its stability and generalization performance through multi-parameter-group optimization, learning rate scheduling, data augmentation, and early stopping mechanisms in the training process.

This section analyzes and summarizes the advantages of the proposed training strategy from three perspectives: adaptive anomaly handling capability, model interpretability, and generalization ability.

In air quality monitoring data, abnormal fluctuations or noisy samples often occur due to unexpected events, environmental changes, or sensor errors. Traditional time-series forecasting methods typically rely on prior data preprocessing steps, such as manually removing or smoothing anomalies, to reduce their impact on prediction results. However, such preprocessing approaches often depend on empirical judgment and may compromise the balance between anomaly detection accuracy and data utilization, with the risk of losing valuable information during the cleaning process.

In contrast, ETSformer incorporates exponential smoothing and damping control mechanisms directly into its structural design, providing inherent modeling capabilities for adaptive anomaly handling. The GrowthLayer within the encoder models trend variations through recursive smoothing, effectively suppressing the amplification of sudden anomalies. Meanwhile, the DampingLayer in the decoder applies attenuation control to the extrapolated trend components during the prediction phase, preventing abnormal values from exerting excessive influence on future forecasts.

This structural design enables the model to automatically adjust its response strength to abnormal fluctuations during the training process without relying on additional anomaly filtering steps. As a result, the model achieves intrinsic anomaly suppression, enhancing the robustness and stability of the overall forecasting process.

Beyond strong predictive performance, ETSformer also emphasizes model interpretability. Both the encoder and decoder utilize multi-head attention mechanisms for feature extraction, offering the potential for attention weight visualization. By analyzing the attention distributions, it is possible to identify which historical time steps the model focuses on during prediction, thereby providing insight into the decision-making process.

Compared with conventional Transformer architectures that rely solely on attention distributions for interpretability, ETSformer further enhances its explanatory power by introducing decomposition into trend, seasonality, and level components. This design not only strengthens the model's interpretability at the attention level but also provides physically meaningful component-based explanations through the outputs of the growth, seasonality, and level terms. Such interpretability helps analyze the contribution of different variation sources to the final prediction results.

The integration of attention visualization with trend decomposition improves the usability and transparency of ETSformer in practical applications. The model not only generates prediction results but also provides rational explanations to support decision-making processes.

In time-series forecasting tasks, high accuracy on the training set does not necessarily ensure good performance on unseen data. To enhance the generalization ability of the model, ETSformer incorporates smoothing mechanisms and damping control into its architectural design, improving stability across different time scales and various air quality indicators.

Specifically, the GrowthLayer within the encoder utilizes exponential smoothing for recursive trend modeling, effectively suppressing high-frequency noise and short-term abnormal fluctuations. This ensures stable learning of long-term trends. The DampingLayer in the decoder further introduces damping factors during the trend extrapolation stage, dynamically controlling the growth rate and preventing unreasonable amplification or oscillation of trend components in long-term forecasting.

This collaborative design of smoothing and damping not only improves convergence efficiency during training but also demonstrates stronger generalization performance during testing. ETSformer consistently maintains stable forecasting accuracy across different air quality indicators—including temperature, humidity, CO_2 concentration, and PM2.5 concentration—and adapts to diverse data characteristics with varying degrees of fluctuation, reducing the risk of overfitting.

In conclusion, ETSformer achieves multiple advantages—including adaptive anomaly handling, enhanced interpretability, and improved generalization—through the integration of trend decomposition, exponential smoothing, damping control, and multi-head attention mechanisms in both model architecture and training strategy. Unlike traditional approaches that depend on manual data cleaning or static preprocessing, ETSformer adaptively responds to abnormal fluctuations within the model itself while preserving the physical interpretability of its prediction components, thereby enhancing its practical value in real-world air quality forecasting tasks.

These advantages enable ETSformer to balance predictive accuracy, stability, and interpretability when dealing with complex and dynamically changing time-series data, providing reliable technical support for multi-scenario environmental monitoring and forecasting applications.

4 Experimental Results and Discussion

4.1 Experimental Setup

To verify the effectiveness of the ETSformer model across various time-series forecasting tasks, this study designs four groups of experiments based on four different types of time-series data, including CO_2 concentration, humidity, PM2.5 concentration, and temperature. All experiments adopt a consistent model architecture and training parameter settings to ensure the fairness of the comparative results. The specific experimental configurations are described as follows.

The four datasets used in this study are univariate time-series data with timestamps. The forecasting target is to predict the sequence values for the next 12 time steps (`pred_len` = 12). The basic information of each dataset is shown in Table **??**.

Table 4.1: Basic information of the datasets used in this study

Dataset Name	Prediction Target	Sampling Interval	Input Sequence Length (`seq_len`)	Prediction Length (`pred_len`)
CO_2	CO_2 Concentration (ppm)	5 minutes	60	12
Humidity	Humidity (%)	5 minutes	60	12
PM2.5	PM2.5 Concentration (μg/m^3)	5 minutes	60	12
Temperature	Temperature (°C)	5 minutes	60	12

The forecasting model used in this study is ETSformer (Exponential Smoothing Transformer), which integrates the concept of exponential smoothing into the Transformer architecture. Through ETS decomposition, the model explicitly handles the trend and seasonal components of time-series data, and introduces a damping mechanism to enhance its forecasting capability under sudden change scenarios.

The main parameter configurations of the ETSformer model are listed in Table **??**.

Table 4.2: Parameter configuration of the ETSformer model

Parameter Name	Value
Model dimension (`d_model`)	256
Number of attention heads (`n_heads`)	4
Number of encoder layers (`e_layers`)	3
Number of decoder layers (`d_layers`)	3
Feedforward network dimension (`d_ff`)	512
Fourier basis (`K`)	1
Activation function (`activation`)	sigmoid
Dropout rate (`dropout`)	0.3
Embedding method (`embed`)	timeF (frequency-based time embedding)

All experiments in this study were conducted under the same training configuration, using a customized Adam optimizer along with an exponential learning rate adjustment strategy. The detailed training parameter settings are summarized in Table 4.3.

Table 4.3: Training parameter configuration of ETSformer

Parameter Name	**Value**
Initial learning rate (`learning_rate`)	1×10^{-5}
Minimum learning rate (`min_lr`)	1×10^{-7}
Warmup epochs (`warmup_epochs`)	2
Batch size (`batch_size`)	256
Training epochs (`train_epochs`)	100 (with early stopping, patience set to 10)
Optimizer	Customized Adam
Loss function	Mean Squared Error (MSE)
Learning rate scheduler (`lradj`)	exponential_with_warmup

4.2 Results for Each Dataset

4.2.1 Training Loss and Validation Loss Analysis

To evaluate the training performance of the ETSformer model across different forecasting tasks, this study records the variations in training loss and validation loss on four datasets, including CO_2 concentration, humidity, PM2.5 concentration, and temperature. The corresponding results are illustrated in Figure **??**, with subfigures (a) to (d) representing the results for each forecasting task, respectively.

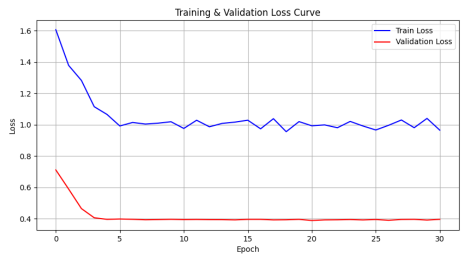

(a) CO_2 concentration forecasting task

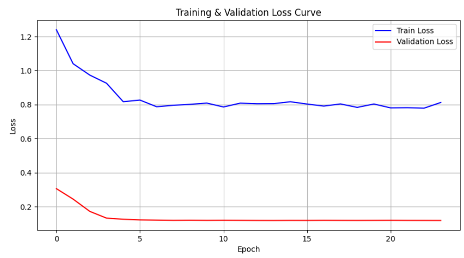

(b) Humidity forecasting task

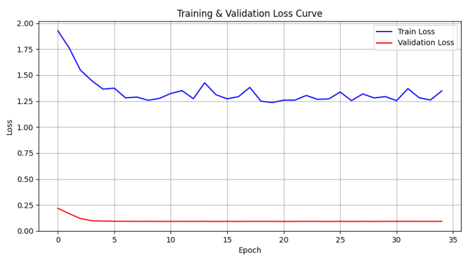

(c) PM2.5 concentration forecasting task

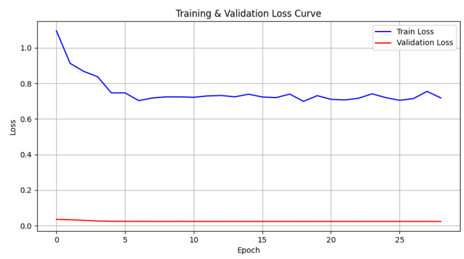

(d) Temperature forecasting task

Figure 4.1: Training and validation loss curves for the four air quality forecasting tasks using the ETSformer model

As shown in Figure **??**, the ETSformer model achieves a stable training process across all four forecasting tasks. Both the training loss and validation loss exhibit an overall decreasing trend and eventually converge.

Specifically, the loss curves on the humidity and temperature datasets show rapid decreases with relatively small fluctuations on the validation sets, indicating that the model demonstrates strong fitting capability on stable and highly periodic data.

In contrast, the CO_2 and PM2.5 datasets, which contain more abrupt changes and abnormal fluctuations, present slightly more volatile validation losses and slower convergence rates. This suggests that the model's training performance is more affected by data complexity when dealing with highly variable sequences. Nevertheless, the ETSformer model still maintains convergence without signs of overfitting.

The ETS decomposition and damping mechanisms integrated into the ETSformer architecture play a significant role during the training process, effectively enhancing the model's stability. These mechanisms are particularly beneficial in trend extraction and learning from stable segments within the time-series data.

4.2.2 Prediction Performance Evaluation

To comprehensively evaluate the forecasting performance of the ETSformer model across different tasks, this study calculates five commonly used error metrics on the four test datasets, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Mean Squared Percentage Error (MSPE). These metrics provide a comprehensive assessment of the model's prediction accuracy and stability from different perspectives, covering absolute deviation, squared deviation, and relative error measurements.

The evaluation results for each dataset are summarized in Table **??**.

Table 4.4: Performance evaluation results of the ETSformer model on different forecasting tasks

Dataset	**MAE**	**MSE**	**RMSE**	**MAPE**	**MSPE**
CO_2	55.06	20330.69	142.59	7.25%	1.79%
Humidity	0.9039	1.9196	1.3855	0.02%	0.00%
PM2.5	5.4546	95.2869	9.7615	3755.83%	16888521735.29%
Temperature	0.2423	0.1511	0.3887	0.01%	0.00%

As shown in Table **??**, the ETSformer model achieves satisfactory forecasting performance across all four prediction tasks. Among them, the humidity and temperature datasets exhibit the lowest error values, indicating that the model demonstrates strong fitting capability on sequences with stable fluctuations and clear periodic characteristics.

In contrast, the MAE and RMSE values for the CO_2 and PM2.5 datasets are relatively higher, particularly in the MAPE and MSPE metrics. This suggests that the model faces certain challenges when dealing with sequences containing abrupt changes and abnormal fluctuations.

These differences are mainly attributed to the inherent characteristics of the datasets: the humidity and temperature sequences exhibit smoother variations and stronger regularity, making it easier for the model to capture their trend and seasonal components. On the other hand, the CO_2 and PM2.5 sequences contain frequent spikes, outliers, or sudden changes, which increase the prediction difficulty and lead to higher error levels.

Overall, the ETSformer model demonstrates good stability and robustness across various prediction tasks, especially achieving high forecasting accuracy in scenarios with stable data patterns.

4.2.3 Visualization Analysis of Prediction Results

To visually demonstrate the forecasting performance of the ETSformer model across different tasks, this study selects representative samples from the test set of each dataset and plots the comparison curves between the ground truth and the predicted values. The corresponding results are presented in Figure **??**, where subfigures (a) to (d) respectively show the prediction performance on CO_2 concentration, humidity, PM2.5 concentration, and temperature forecasting tasks. From the visualization results, it can be clearly

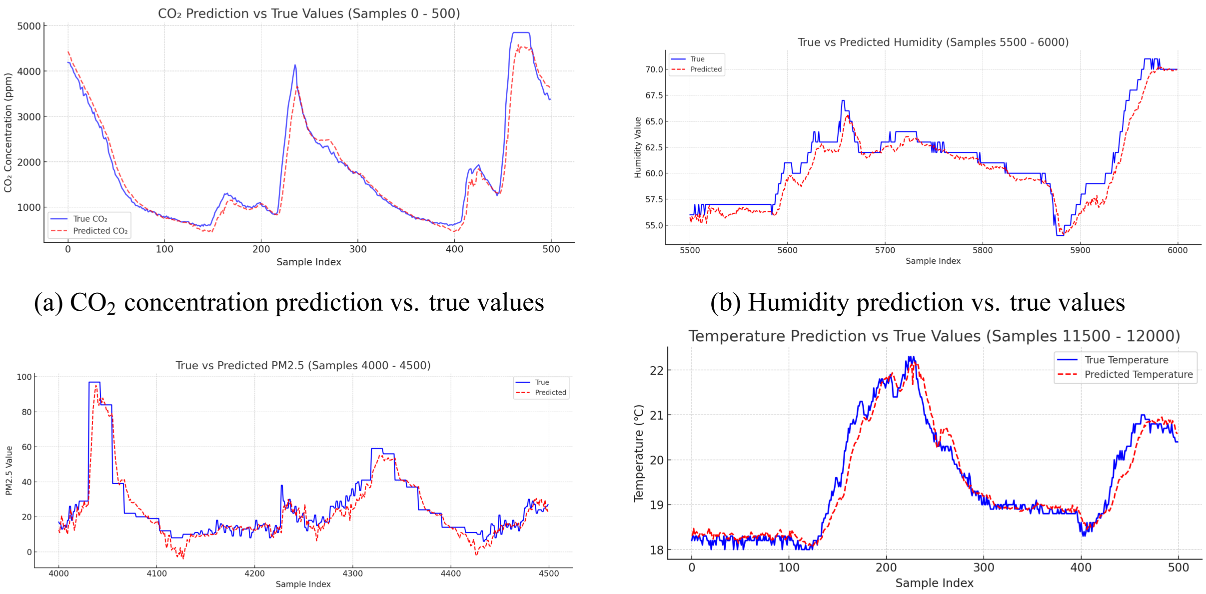

(a) CO_2 concentration prediction vs. true values

(b) Humidity prediction vs. true values

(c) PM2.5 concentration prediction vs. true values

(d) Temperature prediction vs. true values

Figure 4.2: Visualization of the predicted results and ground truth for the four forecasting tasks using the ETSformer model

observed that the ETSformer model achieves a high degree of alignment between the predicted values and the ground truth across all four forecasting tasks. The model successfully captures the underlying trends and periodic patterns of the time series. In particular, for the humidity and temperature prediction tasks, the predicted curves almost completely overlap with the actual curves, demonstrating the model's excellent fitting capability in scenarios with stable and periodic fluctuations.

In the CO_2 and PM2.5 concentration forecasting tasks, although the data exhibit certain volatility and short-term spikes, the model is still able to accurately follow the overall trend and reconstruct the long-term dynamics, avoiding significant deviations between the predicted and true trajectories. The prediction curves remain close to the actual observations throughout most of the time, which highlights the adaptability and stability of the ETSformer model across different types of time-series forecasting tasks.

Overall, the ETSformer model achieves high fitting accuracy across all the forecasting tasks in this study. It demonstrates strong capability in trend characterization and periodic signal extraction, laying a solid foundation for subsequent error analysis and performance comparison.

4.2.4 Error Distribution Analysis

To further evaluate the stability and error characteristics of the ETSformer model across different forecasting tasks, this study conducts a statistical analysis of the prediction errors (the difference between the ground truth and the predicted values) on the test sets. The corresponding error distribution results are illustrated in Figure 4.3, where subfigures (a) to (d) present the error distributions for CO_2 concentration, humidity, PM2.5 concentration, and temperature prediction tasks, respectively.

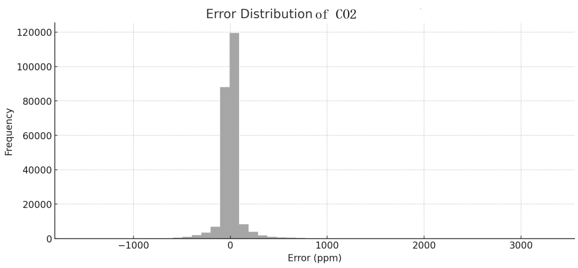

(a) Error distribution of the CO_2 concentration forecasting task

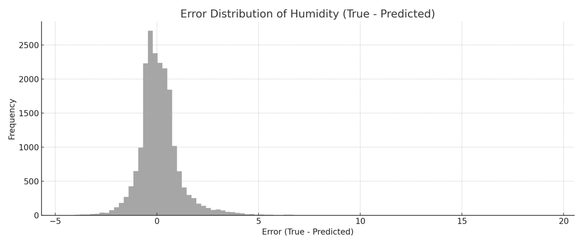

(b) Error distribution of the humidity forecasting task

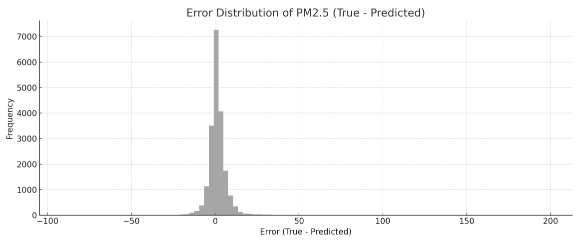

(c) Error distribution of the PM2.5 concentration forecasting task

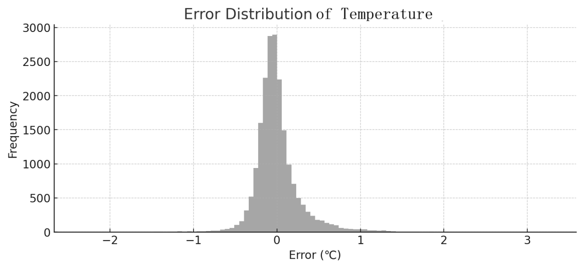

(d) Error distribution of the temperature forecasting task

Figure 4.3: Error distribution analysis of the ETSformer model across four forecasting tasks

From the visualization, it can be observed that the prediction errors of the ETSformer model exhibit a concentrated, symmetric, and approximately normal distribution across all four forecasting tasks. The majority of the errors are distributed near zero, indicating that the model maintains good robustness and consistency in various time-series prediction scenarios without showing systematic bias.

In particular, the error distributions for the humidity and temperature tasks are more compact with shorter tails, further confirming the model's excellent fitting capability on stable and periodic data. Although the CO_2 and PM2.5 concentration datasets contain higher volatility and short-term fluctuations, their error distributions also maintain good symmetry, with most errors concentrated within the range of ±500 and only a few extreme outliers.

These error distribution characteristics clearly demonstrate that the ETSformer model can achieve stable and reliable forecasting performance across different types of time-series prediction tasks. The model not only performs well in trend fitting but also shows strong balance and control over error fluctuations, maintaining consistent prediction quality in both stable and volatile intervals.

In summary, the ETSformer model exhibits stable, well-controlled, and unbiased error distribution patterns across the four forecasting tasks in this study, further validating its strong predictive performance and broad applicability.

4.3 Discussion and Future Work

Based on the experimental results of the CO_2 concentration, humidity, PM2.5 concentration, and temperature forecasting tasks conducted in this study, it can be concluded that the ETSformer model exhibits excellent performance across various prediction scenarios. Benefiting from the introduction of the ETS decomposition mechanism in the model design, ETSformer effectively separates the long-term trend, seasonality, and residual components of time-series data. This enhances the model's capability for trend modeling, resulting in forecasting outputs that closely align with the actual series in terms of overall trajectory.

Throughout the training process, the model demonstrates good convergence and stability. The training and validation loss curves are smooth, without noticeable oscillations or overfitting issues. Additionally, the error distributions on the test sets approximate a normal distribution, indicating strong generalization capability of the model.

Furthermore, ETSformer maintains stable prediction performance across different datasets, including both stable and periodic sequences (such as humidity and temperature) and highly volatile or mutation-prone sequences (such as CO_2 and PM2.5 concentrations). This result suggests that the ETSformer model possesses strong adaptability and broad application potential, capable of handling time-series forecasting tasks under complex environmental conditions. The analysis of various error metrics further validates that the model effectively controls prediction errors in most time periods, with few extreme error points, ensuring stable and reliable forecasting quality.

Despite the favorable results achieved by ETSformer across multiple forecasting tasks in this study, several issues remain to be optimized for real-world applications. Specifically, for datasets with high volatility or abrupt changes (e.g., CO_2 and PM2.5 concentrations), the model exhibits relatively limited responsiveness to sudden variations. In these cases, the forecasted values may lag behind actual changes, and deviations may occur in the peak positions and amplitudes. Moreover, the prediction accuracy slightly decreases in the presence of extreme values, where the forecasts for certain high or low outliers may deviate from the ground truth, thereby affecting the overall forecasting performance to some extent.

These issues are closely related to the core design philosophy of ETSformer, which emphasizes stable trend decomposition. While the ETS mechanism effectively extracts trend and seasonal components and prevents excessive trend extrapolation through the damping mechanism, this suppressive effect may also limit the model's dynamic adjustment capability when facing mutation signals, resulting in insufficient adaptability in sudden-change scenarios.

To enhance the forecasting capability of ETSformer for highly volatile sequences and mutation scenarios, several optimization directions can be considered in future research. First, introducing multi-scale feature modeling mechanisms, such as multi-scale convolution or pyramid structures, could strengthen the model's perception of features at different time scales and improve its ability to capture complex signals with coexisting trends and mutations. Second, integrating anomaly detection techniques to identify and handle mutation points or abnormal segments during the forecasting process may improve the model's responsiveness and accuracy in these situations.

Additionally, combining ETSformer with other models known for mutation handling capabilities, such as LSTM, GRU, or CNN architectures, may allow for complementary strengths between trend extraction and high-frequency variation modeling. Furthermore, designing mutation-aware attention mechanisms that dynamically adjust the model's focus when mutation signals appear could further enhance the model's adaptability in highly variable segments.

In summary, while ETSformer has demonstrated excellent stability and prediction performance in this study, the exploration of the above directions may further improve its applicability and effectiveness, especially for complex time-series forecasting tasks involving abrupt changes and high volatility.

5 Indoor Air Quality Management Strategy Based on Forecasting

In the field of indoor air quality (IAQ) management, the establishment of reasonable control targets serves as the fundamental prerequisite for achieving effective regulation. Fluctuations in air quality parameters not only impact human health and comfort but may also significantly influence virus transmission pathways, especially during special periods such as respiratory infectious disease outbreaks. Therefore, the design of air quality management strategies should fully consider the differences in application scenarios and formulate targeted control objectives to balance both human comfort and epidemic prevention requirements.

Based on this consideration, this study proposes the concept of **"dual-mode management"** in the design of air quality control strategies. Specifically, during normal periods, the control objectives prioritize human comfort, while in epidemic periods, the focus shifts toward suppressing aerosol transmission and reducing the risk of viral infection. This dual-mode design enables dynamic adjustment of air quality parameter settings according to environmental demands, achieving more flexible and scientific indoor environment management.

5.1 Scenario-Based Air Quality Control Objectives

5.1.1 Comfort-Oriented Air Quality Control Targets During Normal Periods

In normal periods, the primary goal of indoor air quality (IAQ) management is to ensure human comfort and maintain a healthy indoor environment that supports cognitive performance, work efficiency, and learning productivity. Numerous studies have demonstrated that a well-maintained indoor air environment not only reduces health risks associated with air pollution but also positively affects cognitive function, sleep quality, and emotional well-being. Therefore, the reasonable definition of control ranges for air quality parameters forms the basis for the development of indoor air management strategies.

From the perspective of comfort, key air quality parameters typically include temperature, relative humidity, carbon dioxide (CO_2) concentration, and fine particulate matter (PM2.5). Among these, temperature and humidity directly influence thermal comfort perception, while CO_2 concentration serves as an indicator of indoor ventilation efficiency, indirectly reflecting occupant density and fresh air supply levels. PM2.5 concentration represents one of the critical pollutants affecting respiratory health.

According to authoritative guidelines such as the World Health Organization (WHO) *Indoor Air Quality Guidelines*, the ASHRAE Standard 55 from the American Society of Heating, Refrigerating and Air-Conditioning Engineers, and the Chinese national standard *GB/T 18883-2022 Indoor Air Quality Standard*, the comfort-oriented target ranges for major air quality parameters during normal periods can be summarized as follows: Among these parameters, the recommended CO_2 concentration range is based on the review by Mendell et al. (2024), which summarizes 43 global indoor air quality guidelines. In most residential, office, and educational environments, 1000 ppm is widely adopted as the reference upper limit for adequate ventilation. However, to further enhance indoor air sensory satisfaction and reduce

Table 5.1: Recommended control ranges for indoor air quality parameters during normal periods

Air Parameter	**Recommended Control Range**
Temperature	22–26 °C
Relative Humidity	40%–60%
CO_2 Concentration	$\leq$ 1000 ppm (preferably $<$ 800 ppm)
PM2.5 Concentration	$\leq$ 15 g/m^3 (24-hour average)

the risk of pollutant accumulation, some studies recommend setting the target concentration below 800 ppm, especially for high-occupancy or long-duration spaces [54, 55].

Thermal comfort refers to the integrated human perception of ambient temperature and humidity conditions. The conventional thermal comfort standard, such as the PMV-PPD model, is derived from the principle of heat balance and predicts human responses to environmental conditions. However, Brager and de Dear (2001), through large-scale field studies across diverse climate zones, proposed the ”adaptive comfort model,” which emphasizes that thermal comfort preferences vary with regional climate, cultural background, and individual expectations. This model suggests that in naturally ventilated spaces, the acceptable temperature range is generally broader compared to mechanically ventilated environments, and moderate fluctuations in indoor conditions can improve user comfort and satisfaction [56].

The adaptive comfort model also highlights the role of behavioral adjustments and psychological adaptation, such as opening windows or changing clothing, in expanding the acceptable temperature and humidity range. Therefore, the air quality management strategy during normal periods should fully consider space types and user habits, supporting a flexible adjustment window for temperature and humidity to balance energy efficiency and occupant comfort.

Unlike pollutants that directly pose health risks (e.g., PM2.5 or CO), CO_2 is typically not regarded as harmful at normal concentrations. Instead, it serves as an important proxy for occupant density and fresh air supply. Studies have shown that maintaining CO_2 concentrations within 600–800 ppm generally meets the requirements for pollutant dilution and comfort. When concentrations approach or exceed 1000 ppm, there may be a risk of minor discomfort, reduced attention, and impaired cognitive performance.

Based on these findings, this study sets the target CO_2 concentration below 800 ppm as a key constraint for the ”comfort-priority” mode under normal conditions.

As for PM2.5, a fine particulate pollutant that can penetrate deep into the respiratory tract, it poses long-term risks to the respiratory and cardiovascular systems. Although short-term low levels of PM2.5 may have limited impact on comfort perception, continuous accumulation in enclosed or semi-enclosed environments can still cause discomfort or health concerns. According to the WHO 2021 air quality guidelines, the recommended annual average concentration for PM2.5 should not exceed 5 g/m^3, and the 24-hour average should not exceed 15 g/m^3 [57]. Considering indoor-outdoor particulate exchange and secondary indoor pollution sources, this study adopts 15 g/m^3 as the PM2.5 control target during the comfort-priority period.

In summary, the air quality management strategy for normal periods focuses on ensuring occupant comfort by integrating considerations of thermal conditions, ventilation, and pollutant levels. The recommended control ranges are established based on international guidelines, national standards, and recent research findings, providing a scientifically grounded and adaptable reference for subsequent management strategies.

5.1.2 Air Quality Control Objectives with Priority on Infection Prevention during Epidemic Periods

During respiratory infectious disease outbreaks, indoor air quality management not only plays a role in maintaining comfort but also becomes a critical factor affecting the transmission of pathogens through aerosols. Since the outbreak of COVID-19, numerous studies have confirmed that the SARS-CoV-2 virus is primarily transmitted via airborne aerosols in indoor environments, especially in poorly ventilated and densely occupied spaces, where the risk of airborne transmission is significantly increased. Therefore, during epidemic periods, the air quality management objectives should shift from comfort-oriented control to infection prevention priority, adjusting air parameters to minimize the risk of virus transmission.

According to the ventilation guidelines issued by the ASHRAE Epidemic Task Force, WHO recommendations on COVID-19 airborne transmission control, and the Chinese "Design Code for Infectious Disease Hospital Buildings," air quality control during epidemic periods should focus on the following key parameters:

Table 5.2: Recommended air quality control ranges during epidemic periods

Air Quality Parameter	Recommended Control Range (Epidemic Priority)
CO_2 concentration	≤ 600 ppm (Enhanced ventilation)
Relative humidity	50%–60% (Properly increased humidity)
Temperature	20–24°C (Adjusted to suppress virus stability)
PM2.5	$\leq 10\ \mu g/m^3$ (Reduce airborne virus carriers)

Among the key air quality parameters, CO_2 concentration not only serves as an indirect indicator of occupancy density and ventilation effectiveness but also directly correlates with the potential aerosol load and infection risk in the air. Bazant et al. (2021) proposed using CO_2 concentration monitoring as a practical alternative to complex aerosol sampling methods, establishing an exposure time limit model based on CO_2 exceedance. This provides an operational risk assessment tool for air quality management during epidemic control periods [58].

In terms of humidity control, Marr et al. (2019) reviewed the impact of relative humidity on the survival and transmission of airborne viruses. They concluded that maintaining a relative humidity range of 50%–60% can effectively reduce virus viability while influencing aerosol particle size, thereby disrupting their aerodynamic behavior and deposition characteristics [59]. This recommendation has been supported by multiple animal studies and epidemiological investigations.

Furthermore, Riddell et al. (2020) systematically demonstrated that the surface stability of SARS-CoV-2 varies significantly under different temperature conditions. Their findings showed that higher temperatures can greatly accelerate viral inactivation, with virus survival on most surfaces being limited to less than 24 hours at 40 °C [60]. Therefore, moderately lowering indoor temperatures can help suppress the environmental stability of viruses in both aerosol and surface states, enhancing infection control effectiveness.

In addition, Nor et al. (2021) reported from monitoring hospital wards during epidemic outbreaks that PM2.5 particles can act as important carriers of viral aerosols. Viral RNA was detected in PM2.5 collection samples, indicating that controlling fine particulate matter loading is also a critical component of air quality management during epidemics [61].

In summary, air quality management during epidemic periods should prioritize minimizing aerosol transmission risks. The control targets should follow the principle of ”ventilation priority, appropriate humidity, temperature suppression, and particulate reduction.” Rational parameter settings rely not only on environmental regulation needs but also on a deep understanding of pathogen survival mechanisms and transmission dynamics. These epidemic control ranges provide a scientific basis for the design of multi-mode management strategies and support the dynamic response and automated adjustment capabilities of air quality management systems during infectious disease control.

5.2 Air Quality Management Strategy Design: Integration of Prediction-Driven and Feedback Regulation

5.2.1 Overall Strategy Design Concept

In the process of indoor air quality management, traditional passive response control strategies rely on real-time sensor detection to trigger regulation actions only after pollutant concentrations exceed preset thresholds. However, this approach often suffers from delayed responses and untimely adjustments, making it difficult to effectively handle rapid changes in pollutant levels, especially in scenarios with high occupancy or unstable ventilation conditions.

To overcome these limitations, this study proposes introducing a feedforward control mechanism based on time-series forecasting into the indoor air quality management system, in combination with traditional feedback control methods, thereby constructing a dual-mechanism interactive air quality regulation strategy.

Specifically, the proposed approach utilizes the trained ETSformer prediction model to perform multi-step time-series forecasting of key air quality parameters (such as CO_2 concentration, PM2.5 concentration, temperature, and humidity), providing the expected trends of these parameters over the next period. The prediction results serve as the feedforward signals, offering proactive decision support for the regulation system. This allows the management system to adjust device operation states in advance, before pollutant concentrations exceed the thresholds, enabling active regulation of air quality.

Meanwhile, the system simultaneously collects real-time monitoring data from sensors, forming the feedback signals. When discrepancies between the actual measurements and the forecasted results occur—such as due to external disturbances (e.g., sudden window opening or occupancy spikes)—the feedback mechanism promptly corrects the control actions to ensure the accuracy and robustness of the regulation strategy. This design combines the foresight of prediction-driven control with the real-time responsiveness of feedback control, forming a dynamically adaptive air quality regulation system.

To achieve effective integration of prediction and feedback, the proposed management strategy adopts a prediction-feedback closed-loop control architecture, with the core workflow consisting of the following steps:

1. **Multi-step prediction based on ETSformer:** Periodically obtain the forecasted trends of air qual-

ity parameters within the prediction interval and identify potential exceedance risks.

2. **Feedforward decision-making:** When the prediction indicates an upcoming exceedance of thresholds, proactively adjust the operation states of the ventilation system, air purification devices, and HVAC system to implement preventive regulation.
3. **Real-time feedback monitoring:** Continuously collect air quality sensor data to assess whether the actual operating state meets the target control ranges.
4. **Deviation compensation and regulation correction:** If prediction errors or external disturbances cause air quality deviations from the target ranges, the feedback mechanism promptly intervenes to dynamically adjust device parameters and compensate for prediction uncertainty.
5. **Mode adaptation and strategy updating:** Depending on the current management mode (comfort-priority or epidemic-control-priority), apply the corresponding target parameter settings and control logic to ensure the achievement of management objectives.

Through the dual-mechanism design that integrates prediction-driven feedforward control with real-time feedback regulation, the air quality management system achieves more flexible and efficient control responses. This design not only enhances the stability of air quality parameters but also supports energy optimization, device health maintenance, and adaptability across various application scenarios.

5.2.2 Control Measures and Multi-Device Coordination Design

In a multi-parameter air quality management system, the rational configuration of control measures and the coordinated operation of multiple devices are the fundamental basis for ensuring the effective implementation of regulation strategies. Based on the air quality management requirements, this study selects natural ventilation, mechanical ventilation, fresh air systems, air purifiers, and air conditioning (including dehumidification and humidification functions) as the main control components. A multi-device coordination mechanism is constructed with a focus on energy efficiency and safety assurance.

Multi-dimensional regulation of air quality requires the flexible scheduling of various devices according to different air parameters and pollution sources, allowing complementary actions to be achieved. The functional mechanisms of each device and their roles in the management strategy are summarized in Table 5.3.

Table 5.3: Control measures and applicable scenarios in air quality management

Control Measure	Main Function	Applicable Scenario
Natural Ventilation	Introduce outdoor fresh air, dilute indoor pollutants	Suitable weather conditions, good outdoor air quality
Mechanical Ventilation	Enhance airflow organization, supplement insufficient natural ventilation	When natural ventilation is not feasible or ventilation demand increases
Fresh Air System	Precisely control fresh air volume, filter outdoor particulate matter	When particulate pollution is severe and large air exchange is required
Air Purifier	Filter PM2.5, remove microbial aerosols	Local supplementary purification, prioritized in densely populated areas
Air Conditioning (Cooling/Heating, Humidification/Dehumidification)	Maintain thermal comfort, regulate humidity levels	Comfort preservation, humidity control for virus transmission suppression

Among these control measures, natural ventilation is the preferred option in the air quality management strategy due to its advantage of requiring no additional energy consumption. When natural ventilation is insufficient or ineffective, mechanical ventilation and fresh air systems serve as supplementary methods to ensure the basic ventilation requirements. Air purifiers act as rapid response tools in localized areas where particulate matter exceeds the standard, while the air conditioning system is responsible for maintaining the stability of the thermal and humidity environment within the target range.

Based on the dual-mechanism framework of prediction-driven and feedback control proposed in this study, the actions of each control device follow the following priority principles:

1. **Energy-saving priority**: Under the premise of meeting air quality targets, natural ventilation is prioritized. The unnecessary operation of high-energy-consuming devices (such as fresh air units and air conditioning systems) is avoided as much as possible.
2. **Safety supplementation**: When natural ventilation cannot meet the air quality goals, mechanical ventilation and fresh air systems are activated in a timely manner to ensure sufficient pollutant dilution and air exchange rates.
3. **Compliance assurance**: If the pollutant prediction trends or real-time monitoring results indicate a risk of exceeding the standard, multiple devices are coordinated to enhance purification capacity (e.g., activating air purifiers, adjusting fresh air volume, modifying humidification or dehumidification strategies of the air conditioning system) to maintain each parameter within the preset control range.

This control logic takes into account the energy efficiency characteristics and functional mechanisms of different devices, ensuring both energy-saving operation and effective regulation during strategy implementation.

Although the control logic remains consistent under different management modes (comfort-priority / epidemic-prevention-priority), the target ranges for air quality parameters differ between these modes. This results in corresponding adjustments to the action thresholds and device scheduling sequences (see Table **??**). For example, under the epidemic-prevention-priority mode:

1. The CO_2 control threshold is stricter (≤ 600 ppm), requiring an early increase in fresh air volume or the intensity of mechanical ventilation.
2. The humidity control target is shifted upward to 50%–60% to suppress the transmission of viral aerosols.
3. Air purifiers are given higher priority, especially in areas with limited ventilation or high occupant density.

By flexibly adjusting the parameter ranges, this design enables the same set of control logic to be efficiently adapted to different application scenarios. This approach enhances the generalizability and emergency response capability of the system, allowing air quality management to dynamically switch between normal operation and epidemic prevention requirements.

5.3 Strategy Implementation Process

In order to achieve efficient control of indoor air quality, this study designs a complete closed-loop control process based on the dual-mechanism concept of prediction-driven and feedback regulation. This process integrates the feedforward control capability provided by air quality forecasting with the feedback correction function from real-time monitoring data, forming a dynamically adaptive air management decision-making system.

The core idea of this closed-loop process includes the following key steps:

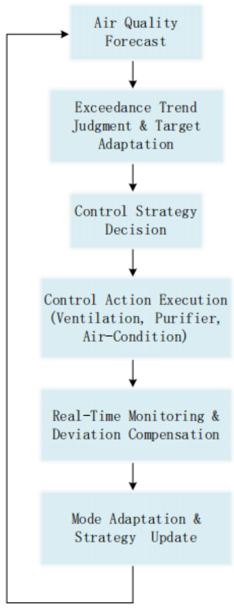

Figure 5.1: Decision Flowchart of Air Quality Control Strategies

1. **Air Quality Parameter Prediction (Predict):** The ETSformer forecasting model is used to perform multi-step prediction on key air quality indicators (CO_2, PM2.5, temperature, humidity) to obtain the trend of air quality changes over the next period. The prediction cycle and steps can be flexibly set according to specific application needs.
2. **Exceedance Trend Determination and Management Target Adaptation (Judge):** The prediction results are compared with the air quality target ranges set under the current management mode (comfort-priority / epidemic-prevention-priority). This step determines whether there is a future risk of exceeding the standard and identifies the regulation targets accordingly.
3. **Control Strategy Decision-Making (Decide):** Based on the judgment results, considering the current operation status of each device, priority logic, and energy consumption balance, a specific control action plan is formulated. This includes adjustments of ventilation volume in the fresh air system, power settings of air purifiers, and humidity parameter modifications of the air conditioning system.
4. **Real-time Feedback Monitoring and Deviation Compensation (Feedback & Correct):** Continuous real-time air quality data collection from indoor sensors is carried out to assess the deviation between predictions and actual conditions. When there are prediction errors or external disturbances that cause pollutant levels to deviate from the target range, the feedback mechanism dynamically corrects control actions to ensure that air quality remains stable within the target range.
5. **Mode Adaptation and Strategy Update (Adapt):** The system supports dynamic switching of management modes based on external epidemic signals or manual intervention, allowing automatic updates of control parameters and adjustments of device strategies. This ensures good adaptability and robustness of the system across different application scenarios.

The overall decision-making process is illustrated in Figure **??**.

5.4 Theoretical Advantages: Exploration of Prediction-Based Air Quality Management Design

In the field of indoor air quality management, traditional control strategies are mostly based on passive responses. These approaches rely on real-time monitoring data and activate control devices only after pollutant concentrations reach or exceed preset thresholds to restore air quality to the target range. However, due to sensor response delays, system inertia, and the rapid accumulation of pollutants, traditional passive control methods often suffer from delayed reactions, large fluctuations, and insufficient energy efficiency, making it difficult to meet air quality assurance needs in high-density occupancy spaces or complex application scenarios.

To address these issues, this study proposes a prediction-based air quality management strategy. It incorporates the future air quality trends predicted by the ETSformer model as feedforward signals into the control system's decision-making process, achieving a fusion of proactive intervention and dynamic regulation. From a theoretical perspective, this strategy offers potential advantages over traditional passive control in terms of response speed, control stability, and energy optimization.

The prediction-based air management strategy can anticipate air quality trends before pollutant concentrations exceed the threshold. It initiates control measures proactively when a future exceedance risk is predicted. Compared with traditional strategies that reactively adjust after the pollutant concentration surpasses the limit, the prediction-driven design has a clear feedforward control advantage. In theory, it can effectively reduce the probability of exceedance events, shorten the air quality recovery time, and improve overall environmental stability.

Traditional passive response control, due to the lack of predictive capability, often leads to sharp fluctuations in control actions and frequent on-off cycling of equipment, increasing both energy consumption and the risk of equipment wear. The prediction-driven strategy designed in this study takes into account the forecasted trends and regulation buffer zones when making decisions on device operation, supporting gradual adjustments of control actions. This design facilitates smooth transitions in the control process while achieving air quality targets, reducing mechanical stress and maintenance costs caused by frequent equipment cycling, and optimizing energy utilization efficiency.

Indoor air quality is influenced by multiple dynamic factors, including seasonal changes, fluctuations in occupancy density, differences in activity types (e.g., meetings, classes, rest periods), weather conditions, and variations in outdoor pollutant levels. The prediction-based management strategy can flexibly adjust the forecasting interval length and control response intensity, demonstrating adaptability to various complex application scenarios. Especially in environments with high and variable foot traffic, this strategy is expected to achieve more precise pollutant control and more reasonable energy scheduling.

It is worth noting that this study primarily focuses on the design and methodology of air quality management strategies. Direct experimental comparisons between the prediction-driven strategy and traditional passive response strategies have not yet been conducted. However, drawing on existing literature regarding the application of predictive control in building energy optimization, industrial process control, and air pollution management, it is generally recognized that prediction-driven approaches provide advantages in response speed, control effectiveness, and energy savings.

For example, in the optimization control of heating, ventilation, and air conditioning (HVAC) systems, predictive models are widely used for load forecasting and feedforward scheduling, and have been proven to enhance energy efficiency and reduce operational fluctuations [62]. In the field of air pollution control, research on model predictive control (MPC) also indicates that using feedforward prediction can significantly reduce the time that pollutant concentrations exceed the standards and improve pollutant reduction efficiency [63]. Therefore, although this study has not yet conducted field validation, the proposed prediction-driven management strategy is theoretically reasonable and holds great application potential.

BIBLIOGRAPHY

[1] D. Chen, G. Huebner, E. Bagkeris, M. Ucci, and D. Mumovic, “Effects of short-term exposure to moderate pure carbon dioxide levels on cognitive performance, health symptoms and perceived indoor environment quality,” *Building and Environment*, vol. 245, p. 110967, 2023.

[2] J. Vidal and A. Gulati, “Impact of indoor air quality on cognitive performance? an evaluation of co2 levels in south florida classrooms,” *An Evaluation of CO2 Levels in South Florida Classrooms (July 20, 2024)*, 2024.

[3] C. Wang, F. Zhang, J. Wang, J. K. Doyle, P. A. Hancock, C. M. Mak, and S. Liu, “How indoor environmental quality affects occupants’ cognitive functions: A systematic review,” *Building and environment*, vol. 193, p. 107647, 2021.

[4] L. Zhao, Y. Yang, and Z. Wu, “Review of communication technology in indoor air quality monitoring system and challenges,” *Electronics*, vol. 11, no. 18, p. 2926, 2022.

[5] S. Yun, “Assessment of monitoring strategies for inhalation exposure and occupancy in office environments,” EPFL, Tech. Rep., 2023.

[6] S. Sun, X. Zheng, J. Villalba-Díez, and J. Ordieres-Meré, “Indoor air-quality data-monitoring system: Long-term monitoring benefits,” *Sensors*, vol. 19, no. 19, p. 4157, 2019.

[7] J. Chen, G. S. Brager, G. Augenbroe, and X. Song, “Impact of outdoor air quality on the natural ventilation usage of commercial buildings in the us,” *Applied energy*, vol. 235, pp. 673–684, 2019.

[8] C. C. Vassella, J. Koch, A. Henzi, A. Jordan, R. Waeber, R. Iannaccone, and R. Charrière, “From spontaneous to strategic natural window ventilation: Improving indoor air quality in swiss schools,” *International Journal of Hygiene and Environmental Health*, vol. 234, p. 113746, 2021.

[9] P. Fermo, B. Artíñano, G. De Gennaro, A. M. Pantaleo, A. Parente, F. Battaglia, E. Colicino, G. Di Tanna, A. G. da Silva Junior, I. G. Pereira *et al.*, “Improving indoor air quality through an air purifier able to reduce aerosol particulate matter (pm) and volatile organic compounds (vocs): Experimental results,” *Environmental research*, vol. 197, p. 111131, 2021.

[10] R. Mumtaz, S. M. H. Zaidi, M. Z. Shakir, U. Shafi, M. M. Malik, A. Haque, S. Mumtaz, and S. A. R. Zaidi, “Internet of things (iot) based indoor air quality sensing and predictive analytic—a covid-19 perspective,” *Electronics*, vol. 10, no. 2, p. 184, 2021.

[11] Y. Rybarczyk and R. Zalakeviciute, “Machine learning approaches for outdoor air quality modelling: A systematic review,” *Applied Sciences*, vol. 8, no. 12, p. 2570, 2018.

[12] S. Ameer, M. A. Shah, A. Khan, H. Song, C. Maple, S. U. Islam, and M. N. Asghar, “Comparative analysis of machine learning techniques for predicting air quality in smart cities,” *IEEE access*, vol. 7, pp. 128 325–128 338, 2019.

[13] T. V. Vu, Z. Shi, J. Cheng, Q. Zhang, K. He, S. Wang, and R. M. Harrison, “Assessing the impact of clean air action on air quality trends in beijing using a machine learning technique,” *Atmospheric Chemistry and Physics*, vol. 19, no. 17, pp. 11 303–11 314, 2019.

[14] C. Fan, M. Chen, X. Wang, J. Wang, and B. Huang, "A review on data preprocessing techniques toward efficient and reliable knowledge discovery from building operational data," *Frontiers in energy research*, vol. 9, p. 652801, 2021.

[15] Z. Xiao, W. Gang, J. Yuan, Z. Chen, J. Li, X. Wang, and X. Feng, "Impacts of data preprocessing and selection on energy consumption prediction model of hvac systems based on deep learning," *Energy and Buildings*, vol. 258, p. 111832, 2022.

[16] Y. Zhang, Y. Wang, M. Gao, Q. Ma, J. Zhao, R. Zhang, Q. Wang, and L. Huang, "A predictive data feature exploration-based air quality prediction approach," *IEEE Access*, vol. 7, pp. 30 732–30 743, 2019.

[17] H. Yao, X. Shen, W. Wu, Y. Lv, V. Vishnupriya, H. Zhang, and Z. Long, "Assessing and predicting indoor environmental quality in 13 naturally ventilated urban residential dwellings," *Building and Environment*, vol. 253, p. 111347, 2024.

[18] J. Saini, M. Dutta, and G. Marques, "Indoor air quality prediction systems for smart environments: A systematic review," *Journal of Ambient Intelligence and Smart Environments*, vol. 12, no. 5, pp. 433–453, 2020.

[19] N. R. Kapoor, A. Kumar, A. Kumar, A. Kumar, and H. C. Arora, "Prediction of indoor air quality using artificial intelligence," *Machine Intelligence, Big Data Analytics, and IoT in Image Processing: Practical Applications*, pp. 447–469, 2023.

[20] H. Liu, G. Yan, Z. Duan, and C. Chen, "Intelligent modeling strategies for forecasting air quality time series: A review," *Applied Soft Computing*, vol. 102, p. 106957, 2021.

[21] K. Mehmood, Y. Bao, W. Cheng, M. A. Khan, N. Siddique, M. M. Abrar, A. Soban, S. Fahad, R. Naidu *et al.*, "Predicting the quality of air with machine learning approaches: Current research priorities and future perspectives," *Journal of Cleaner Production*, vol. 379, p. 134656, 2022.

[22] B. Jo and R. M. A. Khan, "An internet of things system for underground mine air quality pollutant prediction based on azure machine learning," *Sensors*, vol. 18, no. 4, p. 930, 2018.

[23] H. Omidvarborna, P. Kumar, J. Hayward, M. Gupta, and E. G. S. Nascimento, "Low-cost air quality sensing towards smart homes," *Atmosphere*, vol. 12, no. 4, p. 453, 2021.

[24] J. Dutta and S. Roy, "Indoor air pollutant prediction using time series forecasting models," in *Emerging Technologies in Data Mining and Information Security: Proceedings of IEMIS 2020, Volume 1*. Springer, 2021, pp. 499–510.

[25] ——, "Indoorsense: context based indoor pollutant prediction using sarimax model," *Multimedia Tools and Applications*, vol. 80, pp. 19 989–20 018, 2021.

[26] N. N. Zakaria, M. Othman, R. Sokkalingam, H. Daud, L. Abdullah, and E. Abdul Kadir, "Markov chain model development for forecasting air pollution index of miri, sarawak," *Sustainability*, vol. 11, no. 19, p. 5190, 2019.

[27] S. Taheri and A. Razban, "Learning-based co2 concentration prediction: Application to indoor air quality control using demand-controlled ventilation," *Building and Environment*, vol. 205, p. 108164, 2021.

[28] Y. Zhou, F.-J. Chang, L.-C. Chang, I.-F. Kao, Y.-S. Wang, and C.-C. Kang, “Multi-output support vector machine for regional multi-step-ahead pm2. 5 forecasting,” *Science of the Total Environment*, vol. 651, pp. 230–240, 2019.

[29] X. Lai, H. Li, and Y. Pan, “A combined model based on feature selection and support vector machine for pm2. 5 prediction,” *Journal of Intelligent & Fuzzy Systems*, vol. 40, no. 5, pp. 10 099–10 113, 2021.

[30] Y. Zhu, X. Yan, C. Chen, Y. Li, G. Huang, and Y. Li, “Analysis of industry-air quality control in ecologically fragile coal-dependent cities by an uncertain gaussian diffusion-hurwicz criterion model,” *Energy Policy*, vol. 132, pp. 1191–1205, 2019.

[31] H. Rong, A. Teixeira, and C. G. Soares, “Ship trajectory uncertainty prediction based on a gaussian process model,” *Ocean Engineering*, vol. 182, pp. 499–511, 2019.

[32] R. M. A. Velásquez and J. V. M. Lara, “Gaussian approach for probability and correlation between the number of covid-19 cases and the air pollution in lima,” *Urban Climate*, vol. 33, p. 100664, 2020.

[33] Y. Alsultanny, “Machine learning by data mining reptree and m5p for predicating novel information for pm 10,” *Cloud Computing and Data Science*, pp. 40–48, 2020.

[34] S. Devi, F. Esmaeilbeiki, S. M. Karimi, and P. Sihag, “Prediction of sulphur dioxide (so2) in air by using bagging, ann and m5p: a case study, gaya and tirupati, india,” *Arabian Journal of Geosciences*, vol. 15, no. 7, p. 613, 2022.

[35] L. Li, Y. Zhang, J. C. Fung, H. Qu, and A. K. Lau, “A coupled computational fluid dynamics and back-propagation neural network-based particle swarm optimizer algorithm for predicting and optimizing indoor air quality,” *Building and Environment*, vol. 207, p. 108533, 2022.

[36] L. Li, Y. Fu, J. C. Fung, K. T. Tse, and A. K. Lau, “Development of a back-propagation neural network combined with an adaptive multi-objective particle swarm optimizer algorithm for predicting and optimizing indoor co2 and pm2. 5 concentrations,” *Journal of Building Engineering*, vol. 54, p. 104600, 2022.

[37] M. S. Abd Rahim, F. Yakub, M. Omar, R. Abd Ghani, S. A. Z. S. Salim, S. Masuda, and I. Dhamanti, “Prediction of indoor air quality using long short-term memory with adaptive gated recurrent unit,” in *E3S Web of Conferences*, vol. 396. EDP Sciences, 2023, p. 01095.

[38] S. Gurumoorthy, A. K. Kokku, P. Falkowski-Gilski, and P. B. Divakarachari, “Effective air quality prediction using reinforced swarm optimization and bi-directional gated recurrent unit,” *Sustainability*, vol. 15, no. 14, p. 11454, 2023.

[39] L. C. Tagliabue, F. R. Cecconi, S. Rinaldi, and A. L. C. Ciribini, “Data driven indoor air quality prediction in educational facilities based on iot network,” *Energy and Buildings*, vol. 236, p. 110782, 2021.

[40] S. A. Shahriar, I. Kayes, K. Hasan, M. A. Salam, and S. Chowdhury, “Applicability of machine learning in modeling of atmospheric particle pollution in bangladesh,” *Air Quality, Atmosphere & Health*, vol. 13, no. 10, pp. 1247–1256, 2020.

[41] L. Mampitiya, N. Rathnayake, Y. Hoshino, and U. Rathnayake, “Forecasting pm10 levels in sri lanka: A comparative analysis of machine learning models pm10,” *Journal of Hazardous Materials Advances*, vol. 13, p. 100395, 2024.

[42] M. Yu, A. Masrur, and C. Blaszczak-Boxe, “Predicting hourly pm2. 5 concentrations in wildfire-prone areas using a spatiotemporal transformer model,” *Science of The Total Environment*, vol. 860, p. 160446, 2023.

[43] N. Wu, B. Green, X. Ben, and S. O’Banion, “Deep transformer models for time series forecasting: The influenza prevalence case,” *arXiv preprint arXiv:2001.08317*, 2020.

[44] A. Y. Yıldız, E. Koç, and A. Koç, “Multivariate time series imputation with transformers,” *IEEE Signal Processing Letters*, vol. 29, pp. 2517–2521, 2022.

[45] B. Wu, W. Cai, F. Cheng, and H. Chen, “Simultaneous-fault diagnosis considering time series with a deep learning transformer architecture for air handling units,” *Energy and Buildings*, vol. 257, p. 111608, 2022.

[46] J. Dong, Y. Zhang, and J. Hu, “Short-term air quality prediction based on emd-transformer-bilstm,” *Scientific Reports*, vol. 14, no. 1, p. 20513, 2024.

[47] Y. Lu, J. Wang, D. Wang, C. Yoo, and H. Liu, “Incorporating temporal multi-head self-attention convolutional networks and lightgbm for indoor air quality prediction,” *Applied Soft Computing*, vol. 157, p. 111569, 2024.

[48] H. Long, J. Luo, Y. Zhang, S. Li, S. Xie, H. Ma, and H. Zhang, “Revealing long-term indoor air quality prediction: An intelligent informer-based approach,” *Sensors*, vol. 23, no. 18, p. 8003, 2023.

[49] Y. Feng, J.-S. Kim, J.-W. Yu, K.-C. Ri, S.-J. Yun, I.-N. Han, Z. Qi, and X. Wang, “Spatiotemporal informer: A new approach based on spatiotemporal embedding and attention for air quality forecasting,” *Environmental Pollution*, vol. 336, p. 122402, 2023.

[50] Y. Gao and Y. Ruan, “Interpretable deep learning model for building energy consumption prediction based on attention mechanism,” *Energy and Buildings*, vol. 252, p. 111379, 2021.

[51] Z. Zhang and S. Zhang, “Modeling air quality pm2. 5 forecasting using deep sparse attention-based transformer networks,” *International Journal of Environmental Science and Technology*, vol. 20, no. 12, pp. 13 535–13 550, 2023.

[52] W. Cao, W. Qi, and P. Lu, “Air quality prediction based on time series decomposition and convolutional sparse self-attention mechanism transformer model,” *IEEE Access*, 2024.

[53] A. Hameed, J. Violos, A. Leivadeas, N. Santi, R. Grünblatt, and N. Mitton, “Toward qos prediction based on temporal transformers for iot applications,” *IEEE Transactions on Network and Service Management*, vol. 19, no. 4, pp. 4010–4027, 2022.

[54] O. Ramalho, G. Wyart, C. Mandin, P. Blondeau, P.-A. Cabanes, N. Leclerc, J.-U. Mullot, G. Boulanger, and M. Redaelli, “Association of carbon dioxide with indoor air pollutants and exceedance of health guideline values,” *Building and Environment*, vol. 93, pp. 115–124, 2015.

[55] M. J. Mendell, W. Chen, D. R. Ranasinghe, R. Castorina, and K. Kumagai, "Carbon dioxide guidelines for indoor air quality: a review," *Journal of Exposure Science & Environmental Epidemiology*, vol. 34, no. 4, pp. 555–569, 2024.

[56] G. S. Brager and R. De Dear, "Climate, comfort, & natural ventilation: a new adaptive comfort standard for ashrae standard 55," 2001.

[57] L. Morawska and W. Huang, "Who health guidelines for indoor air quality and national recommendations/standards," *Handbook of indoor air quality*, pp. 1–20, 2022.

[58] M. Z. Bazant, O. Kodio, A. E. Cohen, K. Khan, Z. Gu, and J. W. Bush, "Monitoring carbon dioxide to quantify the risk of indoor airborne transmission of covid-19," *Flow*, vol. 1, p. E10, 2021.

[59] L. C. Marr, J. W. Tang, J. Van Mullekom, and S. S. Lakdawala, "Mechanistic insights into the effect of humidity on airborne influenza virus survival, transmission and incidence," *Journal of the Royal Society Interface*, vol. 16, no. 150, p. 20180298, 2019.

[60] S. Riddell, S. Goldie, A. Hill, D. Eagles, and T. W. Drew, "The effect of temperature on persistence of sars-cov-2 on common surfaces," *Virology journal*, vol. 17, pp. 1–7, 2020.

[61] N. S. M. Nor, C. W. Yip, N. Ibrahim, M. H. Jaafar, Z. Z. Rashid, N. Mustafa, H. H. A. Hamid, K. Chandru, M. T. Latif, P. E. Saw *et al.*, "Particulate matter (pm2. 5) as a potential sars-cov-2 carrier," *Scientific Reports*, vol. 11, no. 1, p. 2508, 2021.

[62] A. J. Khabbazi, E. N. Pergantis, L. D. R. Premer, P. Papageorgiou, A. H. Lee, J. E. Braun, G. P. Henze, and K. J. Kircher, "Lessons learned from field demonstrations of model predictive control and reinforcement learning for residential and commercial hvac: A review," *arXiv preprint arXiv:2503.05022*, 2025.

[63] L. Sangiorgi and C. Carnevale, "A model predictive control methodology to integrate short and long term air quality objectives," *IEEE Access*, vol. 12, pp. 10 760–10 768, 2024.