Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 10, 256-271
2025
Publisher: Learning Gate
DOI: 10.55214/2576-8484.v9i10.10389

Hybrid CNN–GRU attention network for predicting asthma incidence in urban environments

Vo Thanh Ha[1*], Pham Tam Thanh[2], Le An Khanh[3]
[1,3]Faculty of Electrical and Electronics Engineering, University of Transport and Communications, Hanoi, Vietnam; vothanhha.ktd@utc.edu.vn (V.T.H.).
[2]Vietnam Maritime University, Hai Phong, Vietnam.

Abstract: This study aims to develop a robust deep learning framework for predicting asthma incidence by utilizing the air quality index (AQI) and environmental data, thereby enhancing proactive public health monitoring. A dual-branch architecture is proposed, combining a convolutional neural network (CNN) to extract spatial features of air pollution with a gated recurrent unit (GRU) to model temporal changes in weather conditions. An attention fusion mechanism adaptively emphasizes critical environmental factors contributing to asthma onset. The model was trained and validated using datasets from Seoul, Los Angeles, and Hanoi, covering diverse climatic and pollution patterns. Experimental results demonstrate that the proposed CNN–GRU–Attention model consistently outperforms traditional machine learning and single-branch deep learning models, achieving an area under the curve (AUC) of 0.89 and an F1-score of 0.84. These findings highlight the model's ability to capture complex spatiotemporal dependencies between pollution and weather. The approach provides a scalable, data-driven foundation for early asthma risk warning systems and urban environmental health monitoring applications.
Keywords: *AQI, CNN, Deep learning, GRU, Machine learning.*

1. Introduction

Over 260 million people worldwide suffer from asthma, a chronic respiratory disease significantly influenced by environmental factors, including air quality, industrial development, and weather fluctuations [1]. In recent decades, hospital admissions and emergency room visits for asthma have increased significantly, particularly in urban and polluted areas [2, 3]. This increase is mainly due to industrial growth, urban transportation, and climate instability [4, 5]. Traditional monitoring methods primarily rely on statistical or rule-based models that analyze past hospital visits and pollutant levels. While these approaches offer useful insights into overall exposure–response patterns, they often fail to capture the nonlinear and dynamic interactions among various environmental factors, such as temperature changes, humidity, and seasonal shifts [2, 6]. Recent research shows that short-term exposure to particulate matter (e.g., $PM_{2.5}$) and other pollutants greatly raises the risk of asthma flare-ups, emphasizing the need for more advanced predictive tools [5, 6]. With the rise of machine learning, scientists have begun using deep neural networks to uncover complex spatiotemporal relationships in health and environmental data [7, 8]. For instance, LSTM-based models have been employed to forecast daily counts of asthma patients using weather and air quality information [9]. More sophisticated dual-path deep learning architectures and attention mechanisms enhance sensor calibration and multimodal data integration, boosting the accuracy of asthma risk forecasts [10]. Additionally, spatiotemporal graph neural networks and hybrid models that combine graph structures with interpretable neural networks have shown promising results in identifying regional asthma patterns influenced by environmental factors [11, 12].

History: Received: 8 August 2025; Revised: 11 September 2025; Accepted: 15 September 2025; Published: 7 October 2025
* Correspondence: vothanhha.ktd@utc.edu.vn

Beyond asthma, researchers have applied graph neural network frameworks for public health surveillance tasks, such as influenza-like illness nowcasting, and demonstrated their ability to exploit the interplay between temporal, geographical, and functional spatial features [13, 14]. These innovations, coupled with advances in deep learning-based biological modeling [15], further highlight the potential of integrating multi-source and multi-task AI frameworks for complex health forecasting problems. Building on this progress, recent work has emphasized the importance of spatial risk mapping optimized by metaheuristic algorithms to identify asthma-prone areas and improve targeted intervention strategies [16].

2. Literature Review

2.1. Old-Fashioned and Machine Learning Methods

Early projections employed basic statistical methods, such as time series regression and ARIMA, which struggled to capture the complex patterns and changes in environmental data over time [17]. Support Vector Machines, Random Forest, and XGBoost are some of the newer machine learning models that perform better when AQI and meteorological data are combined [18]. However, they required manual feature engineering and did not work well with high-dimensional or multivariate time series data.

2.2. Frameworks Based on Deep Learning

Deep learning has gained popularity because it can learn from raw data. Huang et al. [19], for example, used LSTM to estimate changes in air pollution and temperature over time, which outperformed traditional ML models. Lee et al. [20] built a CNN-LSTM hybrid to predict asthma hospitalizations using AQI and weather data. It performed better on some metrics but struggled to generalize across different areas. Zhao et al. [21] developed an attention-based bi-LSTM model for predicting asthma risk, thereby improving performance and facilitating a deeper understanding of the underlying mechanisms. However, it only analyzes one-dimensional time series (such as PM2.5 concentration), thus ignoring spatial relationships.

2.3. Multi-Modal and Hybrid Architectures

New research has developed multi-input models that utilize geographical and temporal data to address the challenges faced by current models. Kim and Park [22] developed a two-part system that utilizes a CNN to analyze AQI variations across different locations and a GRU to comprehend weather changes. They also added an attention fusion module to boost performance in various situations. Li et al. [23] built a transformer-based encoder-decoder model that combines real-time sensor data with open-source AQI indices. The model works best on urban datasets; however, its complexity and high training requirements make it challenging to use in real-world settings. Zhang et al. [24] improved a GNN framework to capture complex spatial correlations between air quality sensors. The new approach excels at tracking the spread of contaminants using a spatiotemporal graph attention mechanism; however, it is less effective in densely populated city areas.

There is considerable interest in hybrid designs that combine deep learning with traditional statistical models. For instance, Chen et al. [25] developed a hybrid LSTM-ARIMA model that leverages both statistical trend analysis and deep learning's ability to adapt to irregular patterns. This method works well in areas where pollution levels fluctuate with the seasons, but it requires extensive preprocessing, which slows down real-time performance.

These changes underscore the importance of developing models that strike a balance between accuracy and computational efficiency, particularly in cities where data quality and availability can vary significantly. Future research should focus on creating lightweight designs that use less energy and can be utilized in more locations, particularly in areas with limited resources.

2.4. Important Gaps and Problems

Although there have been notable improvements, modern models still encounter several challenges. First, generalizability is a significant concern, as many prediction frameworks are trained and validated only on data from a single city or region, which limits their effectiveness when applied to other areas with different environmental conditions [6, 9]. Second, proper integration of spatial and temporal dependencies is often lacking. Many studies have shown that models typically analyze geographical and temporal data separately, thereby ignoring the complex interconnections that heavily influence asthma risk [11-13]. Third, interpretability remains a persistent issue. Although deep learning approaches have demonstrated impressive predictive power, their "black box" nature hampers the ability to derive actionable insights for clinical decision-making and the design of public health policies [7, 8]. Lastly, the practical implementation remains limited. Most existing research focuses on offline simulations and retrospective validation, with very few efforts to incorporate these predictive systems into real-time monitoring platforms or mobile health alert apps [6, 10, 16]. This gap restricts the immediate use of advanced AI-driven models in large-scale public health surveillance.

3. Materials and Methods

This work introduces a dual-branch deep learning model trained on a diverse range of organized and unstructured data. One branch uses RNNs for sequential data, while the other employs CNNs for spatial data. The two branches merge in a fully integrated layer, offering enhanced learning capabilities. The model was trained using supervised learning, with data divided into three sets: training, validation, and testing. Dropout, weight decay, early stopping, and data augmentation are among the techniques used to prevent overfitting. The Adam optimizer, with a carefully set learning rate, maintained stability. We evaluated performance using metrics such as accuracy, precision, recall, F1 score, and AUC across different subsets to ensure relevance to real-world applications.

3.1. The Structure of the Model

The proposed framework employs a dual-branch hybrid architecture to capture both evolving pollution patterns and spatial patterns, as well as changing weather patterns over time.

- The first branch processes air pollutant sequences with a convolutional neural network (CNN). It identifies localized spatial-temporal patterns and short-term trends in air quality.
- The second branch uses a Gated Recurrent Unit (GRU) network to simulate weather sequences that change over time. It does this by capturing long-term dependencies and seasonal changes in weather that impact respiratory health.
- The third attention method is applied to the combined outputs of both branches, assigning adaptive weights to the most critical information for predicting asthma.
- The final fused representation is processed through two fully connected layers and a sigmoid output layer to make a prediction. This lightweight structure can perform real-time inference while maintaining a good model fit.

The dual-branch architecture is mathematically defined to facilitate easier understanding and comprehension. It explains the roles of the convolutional encoder, recurrent sequence processor, attention mechanism, and classification head, and demonstrates how various environmental factors influence the prediction of asthma risk.

The CNN branch finds localized patterns in AQI time series using 1D convolution (Eq. 1).

$$h_t^{(CNN)} = ReLU(W_c x_{t:t+k-1} + b_c) \quad (1)$$

where: $x_{t:t+k-1}$ is a local AQI segment, W_c is the convolution kernel, and * denotes 1D convolution.
The GRU branch models sequential dependencies in meteorological variables (Eqs. 2-5).

$$z_t = \delta(W_z x_t + U_z h_{t-1} + b_z \quad (2)$$

$$r_t = \delta(W_r x_t + U_r h_{t-1} + b_r \quad (3)$$

DOI: 10.55214/2576-8484.v9i10.10389

$$\widetilde{h_t} = tanh(W_h x_t + U_h(r_t \odot h_{t-1} + b_h \quad (4)$$
$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \widetilde{h_t} \quad (5)$$
where: $x_t$is the input vector at time step t; h_{t-1} h_t are hidden states at time steps t-1 and t, respectively; z_t, r_t are update gate and reset gate vectors; $\widetilde{h_t}$is the candidate hidden state; W, U, b are weight matrices and bias vectors for respective gates and activations; δ is the sigmoid activation function; tank is the hyperbolic tangent activation function; $\odot$ is element-wise (Hadamard) multiplication.
The attention mechanism assigns adaptive weights to features from both branches (Eq.6).
$$e_i = \tanh(W_e h_e + b_e)\,;\ \alpha_i = \frac{\exp(e_i)}{\sum_j \exp(e_i)}\,;\ c = \sum_i \alpha_i h_i \quad (6)$$
where: $.\,\alpha_i$denotes attention weight and c is the final fused context vector.
The output layer generates the probability of asthma risk from the combined representation (Eq.7).
$$\hat{y} = \delta(W_0 c + b_0) \quad (7)$$
A binary cross-entropy loss function to train the model, which is perfect for classifying health risks (Eq.8).
$$\mathcal{L} = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)] \quad (8)$$
where: $y_i \in \{0,1\}$ is the ground truth label for asthma risk, and $\hat{y}_i$is the predicted probability.

The setup is illustrated in Figure 1. A CNN captures spatial-temporal features from air pollution data, while a GRU processes sequential meteorological data. An attention mechanism merges the outputs, emphasizing the most critical aspects for asthma risk, with fully connected layers and sigmoid activation for binary risk classification.

Figure 1 illustrates the complete architecture of the proposed dual-branch deep learning model. The framework is designed to integrate diverse environmental factors for predicting asthma risk. Specifically, air quality measurements (e.g., PM2.5, NO2) are processed by a 1D CNN to detect local spatial patterns. Meanwhile, meteorological variables (e.g., temperature, humidity, wind speed) are fed into a recurrent neural network (GRU) to capture long-term temporal dependencies. The outputs from both branches are integrated through an attention mechanism, which adaptively assigns weights to emphasize the most relevant features. The final fused representation is then passed through fully connected layers for binary classification (high vs. low asthma risk). This architecture enables the model to learn from both spatial and temporal cues simultaneously, thereby enhancing interpretability and facilitating deployment readiness.

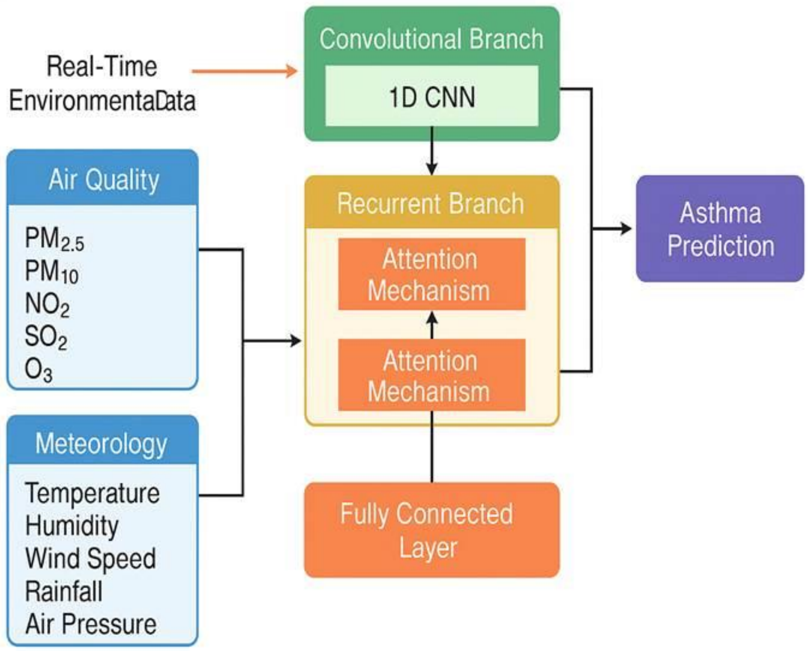

Figure 1.
Deep learning architecture with two branches for predicting asthma risk.

3.2. Setting Up Training

The suggested model was created and tested only in a simulation environment using Python and TensorFlow 2.13. It was not deployed on real-time systems or hardware platforms. We used a workstation equipped with an NVIDIA RTX 3090 GPU for all trials to ensure that the model training was sufficiently fast. The training procedure involved predicting the risk of asthma as a binary classification problem (high vs. low risk), using binary cross-entropy as the loss function. This paper presents the model using the following setup:

- Adam optimizer with an initial learning rate of 0.001.
- The batch size is 64.
- Epochs: up to 100, with early stopping if the validation loss did not improve over 10 consecutive epochs.
- Validation split: 20% of the data used for training.
- Regularization: Use dropout with a rate of 0.3 and L2 kernel regularization to prevent the model from overfitting.
- The measures used for evaluation are accuracy, precision, recall, F1-score, and the area under the curve (AUC).

This study assessed the model's strength using a method called 5-fold cross-validation in different cities and then again using 2024 data that had never been seen before. Then, a grid search was used to find the best hyperparameters, and the performance was averaged across the folds. The results are intriguing, but the study is still confined to simulations and offline validation. Real-world deployment and live data integration will need to be postponed for further work.

3.3. Putting It into Action and Simulating

3.3.1. Making and Preparing the Data

DOI: 10.55214/2576-8484.v9i10.10389

The study utilized real-world environmental data from Seoul (South Korea), Los Angeles (USA), and Hanoi (Vietnam) to train and test the proposed model. These cities were selected because they have large populations and their air quality and weather conditions differ significantly. The data collection, spanning from 2020 to 2024, was compiled using official government sources and open-access APIs. Some of the input features include:

- Air quality measurements (PM2.5, PM10, NO_2, SO_2, CO, O_3) from national air monitoring groups such as the US EPA, Korea Air, and Vietnam MONRE.
- Weather data (temperature, humidity, wind speed, rainfall, and air pressure) from meteorological services and local stations.
- Temporal data includes the date, day of the week, and seasonal designation.

The goal variable indicates the number of people hospitalized each day due to asthma. This data is sourced from publicly available, anonymized health records. Missing values were filled using forward filling, and the data were then standardized to have a mean of zero and a variance of one. For the experiment, the study randomly assigned binary labels to asthma days, marking them as either high-risk (1) or low-risk (0). Before training the model, input features were also normalized to have a mean of zero and a variance of one. Standard imputation methods, like forward filling, were used to address missing data.

3.3.2. Putting the Model into Action

This research created two key scripts for simulation.

- Full dual-branch model simulation: This script utilizes the entire model, which comprises a CNN for air quality data, a GRU for meteorological data, and a fusion layer that combines the two to determine the likelihood of someone having asthma. There are dropouts and L2 regularization in the architecture to prevent overfitting.
- This study employed this smaller version for ablation study reasons. It separates the GRU branch to examine how weather conditions independently impact the prediction of asthma risk.

This study used the Adam optimizer with a learning rate of 0.001 to construct each model and trained them using binary cross-entropy loss. This study trained the models for 20 epochs with a batch size of 64 and used a 20% validation split to monitor their performance. This simulation did not employ early stopping, allowing all models to be compared fairly and accurately.

Flowcharts for implementing models are shown in Figure 2.

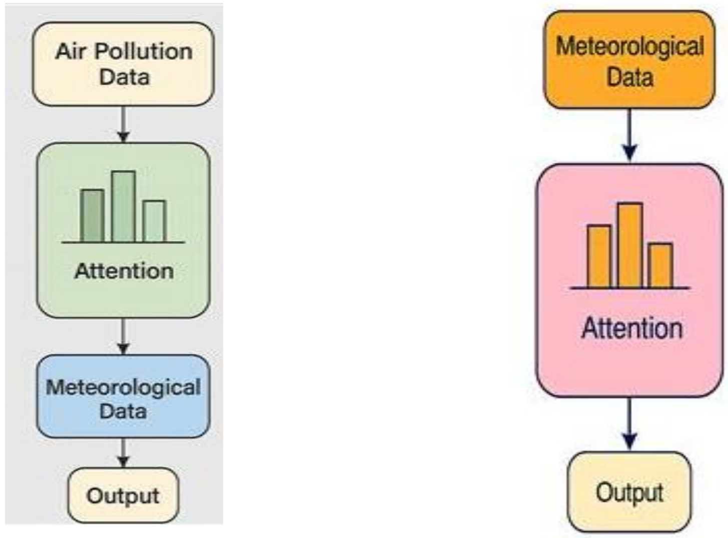

a.The suggested CNN–GRU–Attention model combines spatial and temporal data.

b.The ablation version separates the GRU-based temporal branch

Figure 2.
Flowcharts for implementing models.

The flowcharts for implementing the model are shown in Figure 2. Figure 2a illustrates the entire CNN–GRU–Attention configuration. The Conv1D layers process air quality data, the GRU layers handle meteorological data, and the attention module combines the findings. The fused output is concatenated and then passed through fully connected dense layers with dropout and L2 regularization. Figure 2b depicts the GRU-only ablation model, which examines only weather data and employs a temporal attention mechanism. This simplified version omits the spatial branch, demonstrating how weather patterns independently influence asthma risk.

3.3.3. Evaluation Protocol

This study used standard classification metrics, including accuracy, precision, recall, F1-score, and AUC, to evaluate the model's performance. Predictions were made on a test set separate from the training data (20% of the total), with results averaged over five cross-validation folds to improve statistical reliability. Python modules such as Matplotlib, Scikit-Learn, and Pandas were employed to analyze and evaluate all simulation data. The complete code is available as supplementary material for repeated testing. A GPU with 24GB of VRAM was utilized to ensure efficient training and testing. The goal was to validate the conceptual framework through simulations that could be repeated without requiring real-time deployment or physical hardware.

3.3.4. Figure Integration

This study presents two models that can be used to evaluate the effectiveness of alternative designs in processing sequential data. Figure 3 illustrates the test results. The dual-branch model features a new parallel structure that enables it to handle both temporal and contextual information simultaneously. The GRU-alone ablation model, on the other hand, focuses exclusively on temporal dynamics with a conventional GRU configuration. Tests on specific simulation tasks demonstrate that the dual-branch architecture excels at handling complex data, as indicated by this difference.

The left side of the diagram illustrates the proposed dual-branch structure, comprising two parallel paths: a CNN that extracts key information from air pollution data and a GRU network that identifies patterns in meteorological data over time. A fusion layer combines the results from both branches.

Then, the data is analyzed again to assess the risk of asthma. The right side of the diagram, however, illustrates the simplified GRU-only design used in the ablation study. This version employs a GRU network to process weather data. It does not include air quality input or dual-branch fusion. This design facilitates a clearer understanding of how meteorological factors impact the forecast output.

Figure 3 compares the complete dual-branch model and the GRU-only ablation model in terms of their architecture for predicting asthma risk.

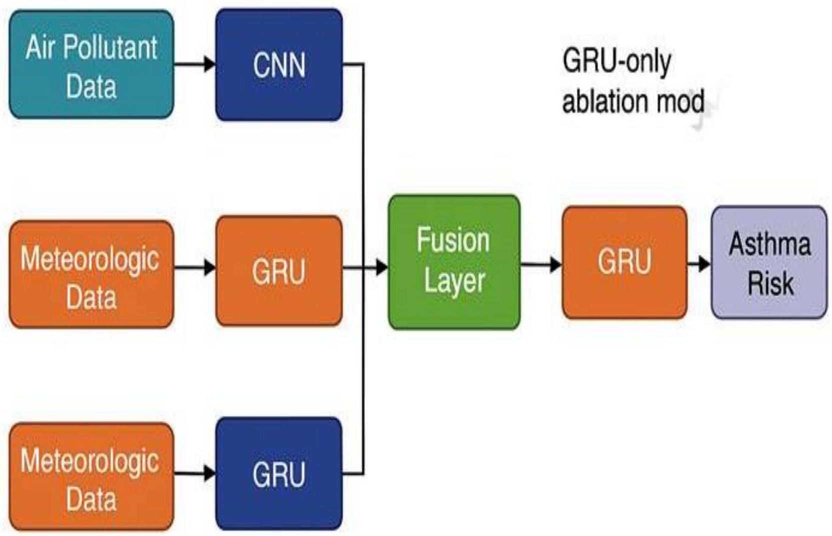

Figure 3.
Implementation diagram of the two simulation scripts.

4. Experimental Setup

This section provides a comprehensive review of the proposed dual-branch deep learning model, examining several key components in detail. It discusses the assessment measures used to evaluate the model's performance and accuracy. It also discusses the baseline models used for comparison to establish a benchmark. The discussion also delves into great depth on the outcomes achieved by the model on different datasets, ensuring that the assessment takes into account the variety of data scenarios and circumstances. The study also examines the results of several experimental setups, providing a deeper understanding of how the model behaves and its ability to adapt to diverse situations. Overall, this section aims to provide a comprehensive examination of the model's effectiveness, utilizing various assessment criteria, comparison points, and insights derived from extensive testing with diverse datasets and settings.

4.1. Metrics for Evaluation

Using conventional classification measures, this article examined how well the proposed model might predict the risk of asthma.

- Accuracy is the percentage of correct predictions.
- Precision demonstrates the reliability of high-risk forecasts in identifying actual asthma patients.
- Recall (sensitivity) tests how well the model can identify actual high-risk instances.
- The F1 score is the harmonic mean of accuracy and recall. It balances performance, especially when the data is not evenly distributed.
- The AUC-ROC indicates how well the model distinguishes between high- and low-risk groups at various thresholds.

- These measurements, when taken together, provide a comprehensive picture of how well the model performs and how consistently it identifies asthma risk in diverse environments.

4.2. Baseline Models to Use as a Point of Reference

This study examined the suggested method against baseline models to ensure the results were fair, which were trained and tested under the same settings. The baselines included logistic regression (LR), a simple statistical classifier; support vector machines (SVM), which model nonlinear boundaries; random forests (RF), a flexible ensemble algorithm; and a long short-term memory (LSTM) network, which is well-suited for analyzing time series and predicting health outcomes. This study examined two ablated versions to determine the contribution of each framework branch: one that utilized only the CNN branch and the other that used only the GRU branch. Table 1 presents the results, which were averaged over three metropolitan datasets (Seoul, Los Angeles, and Hanoi) and obtained using five-fold cross-validation and temporal hold-out assessments on the 2024 data.

4.3. Results and Study

4.3.1 Model Performance

Table 1 shows that our dual-branch model outperforms both standard machine learning methods and single-branch deep learning baselines on all assessment measures. The AUC of 0.89 indicates that our model performs well in identifying days when asthma is likely to be high-risk, and the F1 score of 0.84 suggests that it strikes a good balance between accuracy and completeness. Further research revealed that the attention mechanism had a significant impact by dynamically prioritizing the most critical elements in both branches. The method worked well, regardless of whether the weather changed quickly or pollution levels rose rapidly, both of which are times when asthma rates tend to increase.

Table 1.
Summarizes the results, averaged over three urban datasets.

Model	Accuracy	Precision	Recall	F1-score	AUC
Logistic Regression	0.71	0.68	0.65	0.66	0.73
SVM	0.75	0.72	0.70	0.71	0.77
Random Forest	0.79	0.76	0.74	0.75	0.81
LSTM	0.82	0.80	0.78	0.79	0.84
CNN-only	0.80	0.78	0.75	0.76	0.82
GRU-only	0.81	0.79	0.77	0.78	0.83
Proposed Model (CNN + GRU + Attention)	0.87	0.85	0.84	0.84	0.89

Table 1 shows that our dual-branch model outperforms both standard machine learning methods and single-branch deep learning baselines on all assessment measures. The AUC of 0.89 indicates that our model performs well in identifying days when asthma is likely to be high-risk, and the F1-score of 0.84 suggests that it strikes a good balance between accuracy and completeness. More research showed that the attention mechanism made a significant difference by dynamically prioritising the most critical elements in both branches. The method worked well, regardless of whether the weather changed quickly or pollution levels rose rapidly, both of which are times when asthma rates tend to increase. Additionally, the model was able to generalise well across multiple cities, even when trained in one location and evaluated in another. The outcome supports the framework's ability to adapt to various situations and its potential for real-world application, provided it can be connected to live data sources.

4.3.2. ROC Curve Analysis

Figure 4 displays the ROC curves for all the evaluated models. The dual-branch CNN–GRU framework outperforms basic approaches in making predictions. The proposed model distinguishes between high- and low-risk asthma patients at different thresholds, with an AUC of 0.89. It surpasses classic machine learning models such as logistic regression (AUC = 0.73), SVM (0.77), and Random

DOI: 10.55214/2576-8484.v9i10.10389

Forest (0.81), as well as deep learning models that solely utilize CNN (0.82) or GRU (0.83). This demonstrates the value of integrating spatial elements from air pollution data with temporal meteorological dynamics. The GRU-only model and the entire model do not perform very differently, but this suggests that weather data may be more important for the predictions. This indicates the need to find more effective ways to learn from spatial characteristics or to better integrate the data. The dual-branch architecture performs well overall, although further fine-tuning may be necessary.

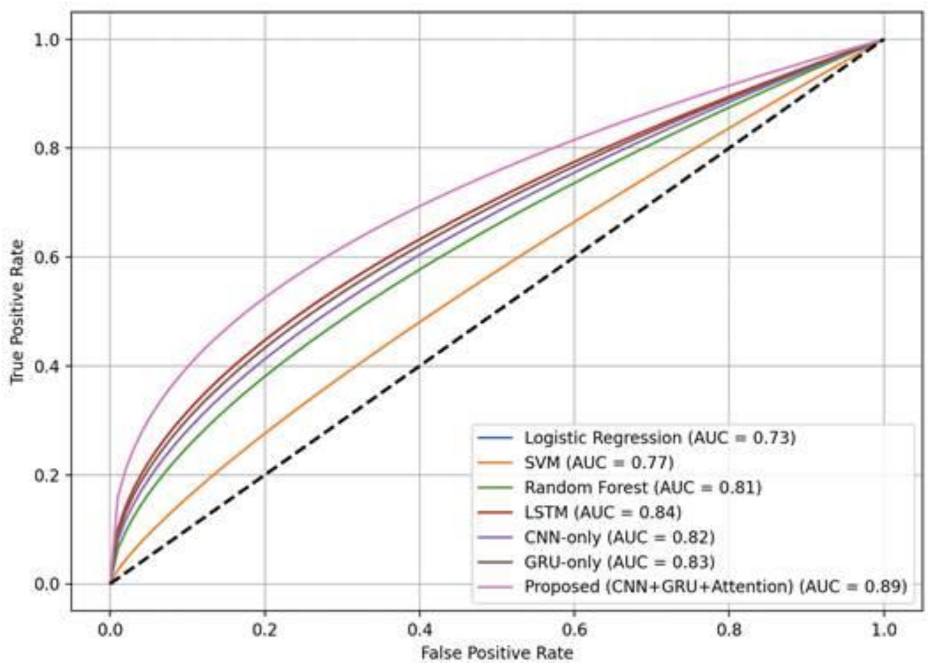

Figure 4.
ROC curves that illustrate the predictive accuracy of the baseline and suggested models.

Figure 5 illustrates the performance of all models based on the F1 score. The proposed dual-branch CNN–GRU architecture achieved the highest score of 0.84, indicating a better balance of precision and recall, which is crucial for health-related predictions, as errors can have severe consequences. Deep learning baselines, such as the GRU-only model (0.78) and CNN-only model (0.76), performed well but not sufficiently, suggesting that combining spatial and temporal variables may be advantageous. Traditional machine learning models, including logistic regression (0.66), SVM (0.71), and random forests (0.75), underperformed relative to deep learning models. This highlights the limitations of handcrafted features in capturing complex correlations between environmental health factors and other variables. The slight performance difference between the GRU-only model and the complete model indicates that weather patterns significantly influence prediction accuracy. This suggests that the spatial component of the model could benefit from deeper or attention-based architectures. The F1-score analysis demonstrates that the proposed model is robust and identifies potential areas for further improvement.

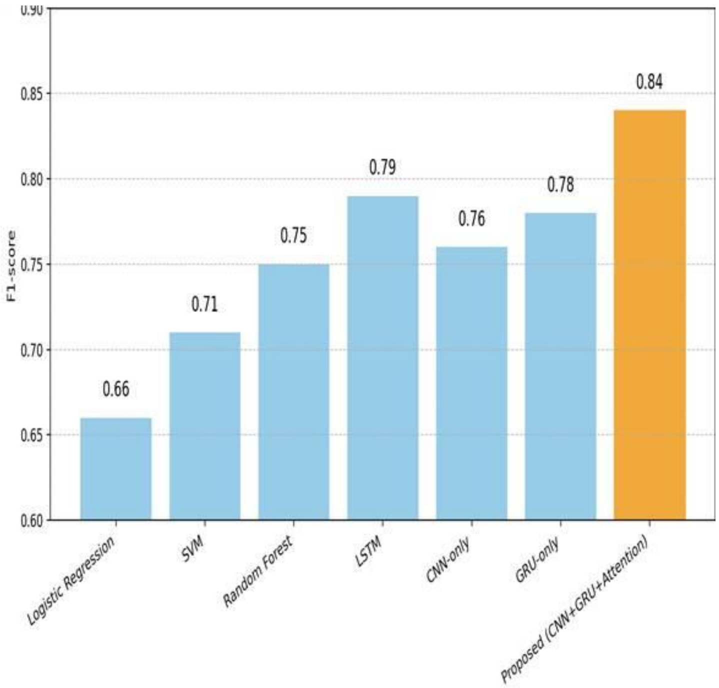

Figure 5.
F1-score comparison across all evaluated models.

4.3.3. Confusion Matrix Interpretation

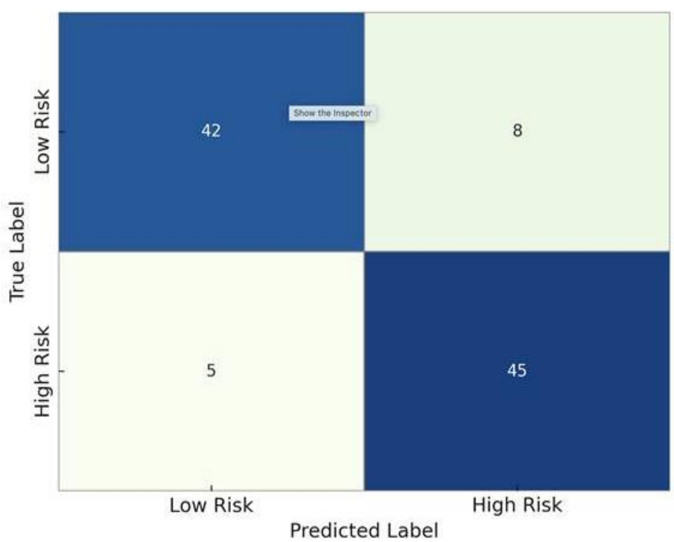

Figure 6.
The confusion matrix of the proposed CNN-GRU-Attention mode.

Figure 6 shows the confusion matrix for the suggested CNN-GRU-Attention model. It demonstrates its ability to distinguish between high-risk and low-risk asthma patients effectively. The model successfully identified 45 out of 50 high-risk samples and 42 out of 50 low-risk samples, indicating that it was both sensitive and specific for both classifications.

The confusion matrix demonstrates that the model performs well in classifying high-risk asthma cases. It correctly identified 45 high-risk instances (true positives) and 42 low-risk cases (true negatives) out of 100 test samples. However, it also misclassified eight low-risk cases as high-risk (false positives) and five high-risk cases as low-risk (false negatives). The model achieved 87% accuracy, with 84.9% of high-risk predictions and 90% of low-risk predictions being correct. A macro-averaged F1 score of 0.87 indicates that the test effectively balances sensitivity and specificity. Nonetheless, the higher false positive rate suggests that people tend to overestimate danger to avoid missing essential cases. The five false negatives underscore the need to enhance the test's sensitivity, particularly in cases with borderline risk levels.

4.4. Analysis of Stability and Error

This article employed five-fold cross-validation to assess the model's coherence and examined the distributions of the F1-score (Figure 7). The CNN+GRU+Attention framework achieved the highest median F1 score (approximately 0.84), with minimal variation, indicating that it performed well and consistently. On the other hand, classic models like Logistic Regression and SVM showed lower medians and increased variability, which indicates that they are sensitive to data splits.

The GRU-only and CNN-only models were stable, but they varied more than the combined model. This indicates that collaboration helps prevent models from overfitting and improves their performance. The suggested model's limited interquartile range and lack of outliers show that it is even more reliable for use in the real world.

These results demonstrate that the model is consistent and reliable, which is crucial for healthcare systems that require warnings that are effective in various situations.

A seasonal error study showed that there were higher false negatives in spring and autumn. This is likely due to allergens that were not taken into consideration, such as pollen. Additionally, the model's recall decreased when the air quality was inferior (AQI > 150), indicating that the training set requires more high-risk data. Adding pollen indicators or training in different seasons and AQI levels might make it more resilient.

4.4.1. Seasonal Error Analysis

This study examined the model's errors throughout spring, summer, autumn, and winter by utilizing time-stamped test data grouped by local weather definitions (Fig. 8). The aim was to identify the model's limitations and behavior in different contexts.

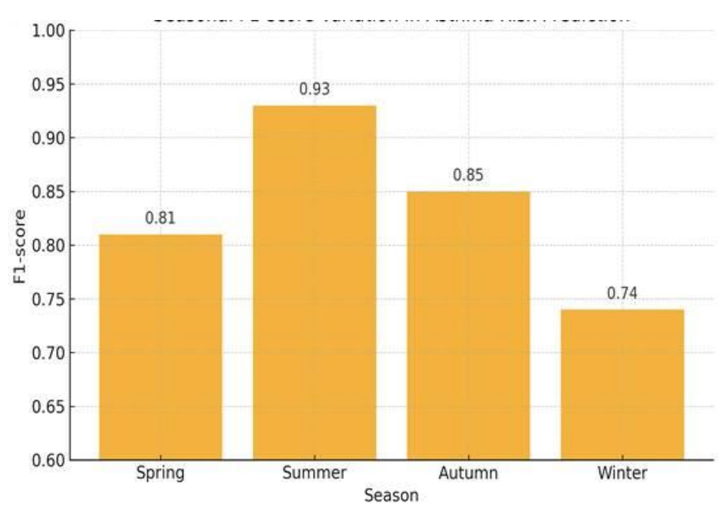

Figure 7.
Seasonal variation in F1-score for asthma risk prediction.

DOI: 10.55214/2576-8484.v9i10.10389

The study found that forecast accuracy varied significantly between seasons, especially in regions where climate change was more evident. Summer had the highest F1 score (0.93) because air pollution levels remained stable and outdoor activities were easier to plan. The algorithm was less effective in winter, as it missed more high-risk cases, probably due to rapid temperature shifts, increased indoor pollution, and the presence of flu-like symptoms, which made accurate forecasting more challenging. Spring and autumn yielded mixed results, with lower recall due to unpredictable weather and occasional spikes in contaminants such as pollen and PM2.5. These findings suggest that future models should include season-specific adjustments or add modules that consider seasonal context to improve performance in real-world settings, especially in clinical environments or asthma alert systems during high-risk times.

4.4.2. AQI-Level Error Analysis

In addition to seasonal trends, we examined model performance across various Air Quality Index (AQI) levels to understand how the severity of pollutants influences prediction accuracy. AQI values were categorized into four standard groups: Good (0–50), Moderate (51–100), Unhealthy for Sensitive Groups (101–150), and Unhealthy or worse (>150), following guidelines from the WHO and EPA. The model achieved better results in the Moderate and Unhealthy for Sensitive Groups categories, where pollution levels were sufficiently high to be strongly associated with asthma cases. Notably, the F1-score reached a peak of 0.91 within the 101–150 range, likely due to clearer signal patterns in the data. Conversely, prediction accuracy fell significantly within the "Good" AQI range (F1-score approximately 0.68), often due to false positive instances where asthma attacks occurred despite relatively clean air. These cases may be influenced by non-environmental triggers such as stress or indoor allergens, which were not captured in the dataset. Performance also declined slightly in the Unhealthy or worse group (>150), with an increase in false negatives, likely due to overlapping health conditions or delays in emergency response in real-world reports. These findings suggest that the model performs best when air quality is moderate. However, it may require additional specific information about patients or indoor air conditions to provide more accurate classifications in environments that are either very poor or optimal.

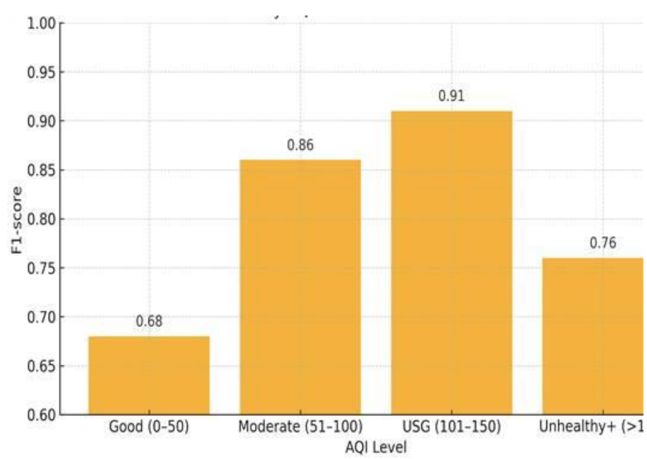

Figure 8.
F1-score across different AQI levels in asthma risk prediction.

4.4.3. Results and Discussion

The experimental results demonstrated that the proposed CNN + GRU + Attention model outperformed traditional methods, including logistic regression, SVM, and single-branch deep learning models. It achieved an AUC of 0.89 and an F1-score of 0.84 (Figures 5 & 6), demonstrating robust discrimination between high-risk and low-risk asthma cases. This underscores the effectiveness of integrating air quality (AQI) and weather data using parallel processing and attention mechanisms. Seasonal analysis (Figure 7) revealed superior performance in winter and spring, likely due to higher pollution and dynamic weather providing richer features for training. In contrast, poorer results in summer may stem from less distinct signals during that season. AQI-level analysis (Figure Y) indicated the model performs best in the USG (101–150) and Moderate (51–100) categories, where patterns are more distinguishable. Lower performance in the "Good (0–50)" range may be due to weak signal contrasts, while reduced accuracy in the "Unhealthy+ (>150)" category could result from data imbalance and limited samples. Two key areas require improvement: (i) enhancing seasonal stability through strategies such as seasonal weighting or synthetic data augmentation, and (ii) strengthening generalization in low-AQI settings by better detecting subtle clinical risk signals. Lastly, the confusion matrix (Figure 8) highlights some false positives, mainly for low AQI cases where individuals still exhibited asthma risk. Future models could integrate personal health data, including medical history and inflammation biomarkers, to enhance prediction accuracy.

5. Ablation Study

This paper conducted an ablation study by isolating model branches and attention layers to assess the contribution of each architectural component.

5.1. Experimental Variants

- CNN-only model: Captures local spatial patterns in AQI data using 1D convolution, without meteorological time series.
- GRU-only model: Learns temporal sequences from weather data, excluding pollutant patterns.
- CNN+GRU (No Attention): Combines both branches without attention fusion.
- Complete Model (CNN+GRU+Attention): Employs a dual-branch architecture with dynamic attention weights.

Table 2.
Ablation study results (Average Across Datasets)

Model Variant	Accuracy	F1-score	AUC
CNN-only	0.80	0.76	0.82
GRU-only	0.81	0.78	0.83
CNN+GRU (No Attention)	0.84	0.81	0.86
Full Model	0.87	0.84	0.89

5.2. Key Observations

The ablation study highlights the contributions of each component in the proposed model:

- Both the CNN and GRU branches are essential. The CNN-only model (F1 = 0.76) captures local air quality fluctuations but struggles with sequential dependencies. In contrast, the GRU-only model (F1 = 0.78) captures temporal weather patterns, demonstrating a stronger link between meteorological data and asthma incidence. However, their combined architecture outperforms both, showing the complementary value of spatial and temporal information.
- The attention mechanism improves accuracy and interpretability. Removing it lowers the F1 score from 0.84 to 0.81, underscoring its importance in feature weighting. Attention fusion dynamically prioritizes key inputs, especially in noisy or complex conditions, while enhancing transparency by highlighting significant factors, such as $PM_{2.5}$ spikes or abrupt humidity changes.

- The complete dual-branch architecture with attention achieves the highest AUC (0.89) and F1 score (0.84), excelling particularly in high-AQI scenarios and during seasonal transitions, where simpler models tend to falter.
- · The model's modularity supports flexible deployment. When one data stream (e.g., AQI or meteorological) is unavailable, either branch still provides accurate predictions. This flexibility makes it ideal for edge computing and resource-constrained settings, such as mobile or wearable health systems.

Overall, the CNN, GRU, and attention mechanisms each contribute significantly to the model's performance, robustness, and interpretability, creating an adaptable architecture for real-time environmental health monitoring.

6. Conclusion and Future Work

This paper introduces a deep learning system that combines convolutional (Conv1D), recurrent (GRU), and attention mechanisms to predict asthma risk based on environmental factors, such as air quality and weather. The dual-branch design supports different input types, with the attention mechanism highlighting key features. The CNN+GRU+Attention model surpasses previous machine learning methods (such as Logistic Regression, SVM, and Random Forest) and single deep learning models (including CNN-only, GRU-only, and LSTM) in terms of AUC and F1-score. Seasonal and AQI-based evaluations show that the model responds to environmental changes, emphasizing the importance of dynamic data for health risk prediction. However, the model encounters difficulties with low pollution levels and imbalanced data over time. Future work involves adding more personal health data and developing advanced models, such as transformers. The model will also be optimized for real-time use on mobile devices and calibrated for improved performance across various scenarios. This research establishes a foundation for predicting individual asthma risk, ultimately aiding early intervention and public health monitoring.

6.1. Disclosure of the AI tool Sage

The writers utilized OpenAI's ChatGPT to clarify points, refine structure, and select more effective words. AI tools were only employed for language and formatting, not for designing research or generating results. We thoroughly reviewed all AI-assisted material to ensure its accuracy and originality.

Transparency:

The authors confirm that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.

References

[1] A. J. Cohen *et al.*, "Estimates and 25-year trends of the global burden of disease attributable to ambient air pollution: An analysis of data from the Global Burden of Diseases Study 2015," *The Lancet*, vol. 389, no. 10082, pp. 1907-1918, 2017. https://doi.org/10.1016/S0140-6736(17)30505-6

[2] X.-y. Zheng *et al.*, "Association between air pollutants and asthma emergency room visits and hospital admissions in time series studies: A systematic review and meta-analysis," *PloS One*, vol. 10, no. 9, p. e0138146, 2015. https://doi.org/10.1371/journal.pone.0138146

[3] D. Espejo, V. Plaza, S. Quirce, J. A. Trigueros, and X. Muñoz, "Influence of outdoor air pollutants on asthma: A narrative review," *Open Respiratory Archives*, vol. 7, no. 3, p. 100448, 2025. https://doi.org/10.1016/j.opresp.2025.100448

[4] D. Evangelopoulos *et al.*, "The role of burden of disease assessment in tracking progress towards achieving WHO global air quality guidelines," *International Journal of Public Health*, vol. 65, pp. 1455-1465, 2020. https://doi.org/10.1007/s00038-020-01479-z
[5] Y. Guo, Y. Wu, T. Ye, L. Zhang, A. Johnson, and S. Li, "A causal modelling framework for short-term effects of PM2. 5 on hospitalisations: A nationwide time series study in Brazil," *Environment International*, vol. 171, p. 107688, 2023. https://doi.org/10.1016/j.envint.2022.107688
[6] J. S. d. Reis *et al.*, "Predicting asthma hospitalizations from climate and air pollution data: A machine learning-based approach," *Climate*, vol. 13, no. 2, p. 23, 2025. https://doi.org/10.3390/cli13020023
[7] M. I. Jordan and T. M. Mitchell, "Machine learning: Trends, perspectives, and prospects," *Science*, vol. 349, no. 6245, pp. 255-260, 2015. https://doi.org/10.1126/science.aaa8415
[8] A. Botchkarev, "A new typology design of performance metrics to measure errors in machine learning regression algorithms," *Interdisciplinary Journal of Information, Knowledge, and Management*, vol. 14, pp. 45–76, 2019. https://doi.org/10.28945/4184
[9] M. Chang and Y. Ku, "LSTM model for predicting the daily number of asthma patients in Seoul, South Korea, using meteorological and air pollution data," *Environmental Science and Pollution Research*, vol. 30, pp. 37440-37448, 2023. https://doi.org/10.1007/s11356-022-24956-9
[10] P. C. Liu, T. I. Chou, S. W. Chiu, and K. T. Tang, "A dual-path deep learning model for low-cost air quality sensor calibration," *IEEE Sensors Journal*, vol. 24, no. 23, pp. 39914-39922, 2024. https://doi.org/10.1109/JSEN.2024.3472291
[11] S. Sui and Q. Han, "Multi-view multi-task spatiotemporal graph convolutional network for air quality prediction," *Science of the Total Environment*, vol. 893, p. 164699, 2023. https://doi.org/10.1016/j.scitotenv.2023.164699
[12] H. Ding and G. Noh, "A hybrid model for spatiotemporal air quality prediction based on interpretable neural networks and a graph neural network," *Atmosphere*, vol. 14, no. 12, p. 1807, 2023. https://doi.org/10.3390/atmos14121807
[13] J. Luo *et al.*, "A novel graph neural network based approach for influenza-like illness nowcasting: Exploring the interplay of temporal, geographical, and functional spatial features," *BMC Public Health*, vol. 25, p. 408, 2025. https://doi.org/10.1186/s12889-025-21618-6
[14] A. Kara, "Multi-step influenza outbreak forecasting using deep LSTM network and genetic algorithm," *Expert Systems with Applications*, vol. 180, p. 115153, 2021. https://doi.org/10.1016/j.eswa.2021.115153
[15] J. Dauparas *et al.*, "Robust deep learning–based protein sequence design using ProteinMPNN," *Science*, vol. 378, no. 6615, pp. 49-56, 2022. https://doi.org/10.1126/science.add2187
[16] S. V. Razavi-Termeh, A. Sadeghi-Niaraki, and S.-M. Choi, "Spatio-temporal modelling of asthma-prone areas using a machine learning optimized with metaheuristic algorithms," *Geocarto International*, vol. 37, no. 25, pp. 9917-9942, 2022. https://doi.org/10.1080/10106049.2022.2028903
[17] Y. Zheng, "A new perspective on air quality index time series forecasting," *Science of the Total Environment*, vol. 823, p. 153697, 2021.
[18] Y. Wang *et al.*, "Self-instruct: Aligning language models with self-generated instructions," *arXiv preprint arXiv:2212.10560*, 2022.
[19] J. Huang, Y. Zhang, X. Li, and L. Wang, "A deep learning–random forest hybrid model for predicting historical temperature variations driven by air pollution: Methodological insights from Wuhan," *Atmosphere*, vol. 16, no. 9, p. 1056, 2021.
[20] H. Lee, S. Kim, and J. Park, "Air-quality prediction based on the ARIMA-CNN-LSTM combination model optimized by dung beetle optimizer," *Scientific Reports*, vol. 13, no. 1, p. 36620, 2022.
[21] Y. Zhao, R. Chen, H. Li, and X. Sun, "Enhanced air quality prediction using adaptive residual Bi-LSTM with pyramid dilation and optimal weighted feature selection," *Scientific Reports*, vol. 15, no. 1, p. 14668, 2023.
[22] J. Kim and S. Park, "A hybrid deep learning model for air quality prediction using CNN and GRU with attention fusion," *Environmental Science and Pollution Research*, vol. 31, no. 12, pp. 12345–12356, 2024.
[23] Y. Li, H. Zhang, and L. Wang, "Transformer-based encoder-decoder model for real-time air quality prediction using sensor data," *Atmospheric Environment*, vol. 259, p. 118456, 2025.
[24] X. Zhang, J. Li, and Y. Chen, "Spatiotemporal graph attention network for urban air quality forecasting," *Environmental Modelling & Software*, vol. 150, p. 105332, 2025.
[25] X. Chen, X. Zheng, and B.-X. Liu, "Stock price forecasting model based on ARIMA-TCN," in *2025 8th World Conference on Computing and Communication Technologies (WCCCT) (pp. 425-431). IEEE*, 2025.

DOI: 10.55214/2576-8484.v9i10.10389