Ensemble Transfer Learning for Air Quality Classification: A Robust Model for Environmental Monitoring

Venu K, Krishnakumar B, Sasipriyaa N,
Deepak Raajan N, Dharaneesh S and Hari R P
Kongu Engineering College, Erode, Tamil Nadu, India

Keywords: Air Pollution Prediction, Transfer Learning, Ensemble Model, VGG16, ResNet, EfficientNet.

Abstract: One of the most urgent environmental issues facing the world today is air pollution, which has an immediate impact on ecosystems and human health. Predicting air quality accurately is crucial for mitigation plans and early warning systems. The classification of air quality into six predetermined categories—Good, Moderate, Unhealthy for Sensitive Groups, Unhealthy, Very Unhealthy, and Hazardous/Severe—is examined in this paper using transfer learning techniques. A dataset of images representing air quality levels was used, with pre-trained models such as VGG16, ResNet, and InceptionV3 fine-tuned for this classification task. Transfer learning models, known for their efficiency in image-based tasks, were individually tested and compared based on classification accuracy and performance metrics. To further improve prediction accuracy, a novel ensemble approach was implemented, combining VGG16, ResNet, and EfficientNet into a unified model. The ensemble model achieved significantly higher accuracy compared to individual models, particularly in predicting complex air quality scenarios such as the Hazardous/Severe category. This improvement in accuracy underscores the potential of combining multiple pre-trained models in air quality prediction tasks, addressing the challenge of differentiating between closely related pollution levels. The results suggest that this ensemble approach not only enhances classification accuracy but also provides a more robust prediction framework for real-world applications. The proposed method shows promise for integration into real-time air quality monitoring systems, offering an effective tool for public health agencies to predict and respond to deteriorating air quality conditions. This study adds to the expanding corpus of research on transfer learning and how it's used in environmental monitoring.

1 INTRODUCTION

One of the biggest threats to the environment and public health in the world today is air pollution. The World Health Organization (WHO) estimates that air pollution causes millions of premature deaths annually, mostly from cardiovascular and respiratory conditions that are made worse by extended exposure to dangerous air pollutants. These pollutants, which all contribute to the deterioration of air quality, include carbon monoxide (CO), nitrogen dioxide (NO_2), sulfur dioxide (SO_2), particulate matter (PM), and ozone (O_3). To lessen the negative consequences of poor air quality and enable prompt warnings and the implementation of preventive actions, accurate air pollution prediction systems are crucial. Due to their capacity to handle large volumes of data and generate accurate forecasts, machine learning and deep learning approaches have been increasingly popular in recent years for the prediction of air pollution. Traditional air pollution prediction models, such as chemical transport models or statistical regression methods, often rely on handcrafted features and a deep understanding of atmospheric dynamics. While these models are effective, they face challenges in accurately predicting complex air quality patterns, especially in environments with rapid fluctuations in pollution levels. Deep learning algorithms have therefore become a viable substitute because of their capacity to automatically extract high-level characteristics from big datasets without the need for manual intervention.

Convolutional Neural Networks (CNNs), one of the deep learning methods, have shown exceptional efficacy in image categorization tasks. CNNs are appropriate for predicting air pollution levels from sensor images or data visualizations because they can learn to identify intricate patterns in data. However,

K, V., B, K., N, S., N, D. R., S, D. and R P, H.
Ensemble Transfer Learning for Air Quality Classification: A Robust Model for Environmental Monitoring.
DOI: 10.5220/0013609300004664

In *Proceedings of the 3rd International Conference on Futuristic Technology (INCOFT 2025) - Volume 3*, pages 91-98
ISBN: 978-989-758-763-4

building deep neural networks from the ground up frequently calls for substantial computational resources and huge datasets. By allowing the use of pre-trained models that have been refined on huge picture datasets like ImageNet, transfer learning provides a workable solution to this issue. High accuracy can be attained without a large dataset by optimizing these pre-trained models for the particular goal of air pollution prediction.

In addition to testing individual models, a novel ensemble approach is explored by combining VGG16, ResNet, and EfficientNet. To increase overall prediction accuracy, ensemble learning combines several models, each of which has unique strengths: ResNet handles deep architectures, EfficientNet scales network dimensions effectively, and VGG16 is excellent at extracting detailed features. This ensemble model aims to leverage these strengths to enhance classification performance. Images illustrating various degrees of air pollution are included in the dataset utilized in this study, which is arranged according to the Air Quality Index (AQI). More accurate forecasts were made for high pollution levels like Hazardous/Severe by the ensemble model, which also performed better, especially when it came to differentiating between closely similar air quality categories like Unhealthy and Very Unhealthy. Complex patterns in the data can be effectively captured by combining ensemble learning and transfer learning.

In summary, transfer learning models, especially when combined in an ensemble approach, significantly improve the accuracy of air pollution prediction systems. This method holds promise for real-world applications, enabling more reliable air quality monitoring and timely public health interventions. The results of this study demonstrate how cutting-edge deep learning methods can be applied to urgent environmental issues like air pollution.

2 RELATED WORKS

Samad et al. (Samad, Garuda, et al. , 2023) sought to replace conventional air quality monitoring systems with a new method for predicting air pollution by utilizing machine learning-powered virtual monitoring stations. This method utilizes data from limited physical monitoring stations and expands its coverage through predictive models, enabling accurate and cost-effective air quality monitoring over large areas.

Kumar and Pande (Kumar, Pande, et al. , 2023) investigated the use of machine learning methods for air pollution prediction in their case study on Indian cities. The study underlined how crucial it is to customize models for particular areas because pollution sources and urban features vary widely throughout India. Their findings highlighted the significance of region-specific solutions in addressing air pollution challenges.

Karimi et al. (Karimi, Asghari, et al. , 2023) proposed white-box machine learning models for predicting air pollution levels near industrial zones. Unlike black-box models, their approach provided interpretability, enabling better understanding of the factors influencing pollution. Their work underscores the importance of transparency in predictive models for environmental applications.

Luo and Gong (Luo and Gong, 2023) created a hybrid ARIMA-WOA-LSTM model for predicting air pollutants. By combining ARIMA for temporal pattern extraction, Whale Optimization Algorithm (WOA) for parameter optimization, and LSTM for deep learning-based trend prediction, this approach demonstrated enhanced accuracy in forecasting air quality trends.

Yang et al. (Yang, Wang, et al. , 2023) presented a district-level air pollution forecasting system that not only predicts pollutant levels but also evaluates the associated health impacts and economic costs. This comprehensive system integrates prediction with actionable insights for policymakers, emphasizing the multidimensional implications of air pollution.

Pan et al. (Pan, Harrou, et al. , 2023) compared various machine learning techniques for predicting ozone pollution. Their work evaluated the performance of various algorithms, highlighting the advantages and limitations of each. This study contributes to the selection of appropriate methods for specific air pollution prediction tasks.

Gupta et al. (Gupta, Mohta, et al. , 2023) looked into predicting the Air Quality Index (AQI) using a variety of machine learning algorithms. The usefulness of various algorithms was revealed by their comparison study, which also emphasized the significance of selecting appropriate models according to the kind of air contaminants and the properties of the data.

Li and Jiang (Li and Jiang, 2023) introduced a novel predictive framework combining Temporal Convolutional Networks (TCN), BiLSTM, and DMAttention with STL decomposition. This approach decomposed the time series into multiple components, improving the predictive accuracy for air pollutant concentrations by addressing seasonal and trend variations effectively.

A study by Hardini et al. (Hardini, Chakim, et al. , 2023) investigated the use of convolutional neural networks (CNNs) for image-based air quality prediction. Using visual data, their approach classified air quality into categories such as "Good," "Moderate," and "Hazardous." This innovative method leverages the potential of computer vision in environmental monitoring.

Cao et al. (Cao, Zhang, et al. , 2023) used Empirical Mode Decomposition (EMD) to create a hybrid air quality prediction model. Their approach broke down complex pollutant signals into simpler components, allowing for more precise predictions. This work highlights the role of signal processing techniques in enhancing air pollution modeling.

3 PROPOSED WORK

The proposed system for air pollution prediction utilizes transfer learning, employing pre-trained deep learning models to categorize air quality into six levels: Good, Moderate, Unhealthy for Sensitive Groups, Unhealthy, Very Unhealthy, and Hazardous or Severe. This system employs a combination of VGG16, ResNet, and EfficientNet, which are known for their high performance in image recognition tasks. These models, when fine-tuned with a specific air quality image dataset, exhibit the ability to effectively classify the air quality levels, thus making it a powerful tool for air pollution prediction. The dataset used in this system contains images that represent different levels of air pollution. These images are categorized into six classes, each corresponding to a specific air quality index (AQI). The dataset is essential for deep learning model training and validation. The dataset is preprocessed through scaling, normalization, and augmentation to improve efficiency and help the model generalize well to new data. In instance, by adding variances to the training images, data augmentation enables the model to learn more robust features.

Transfer learning is a key aspect of the proposed system. Rather than starting from scratch when training a deep neural network, pre-trained models like VGG16, ResNet, and EfficientNet are utilized. These models can recognize intricate patterns in visual data since they have already been trained on extensive image datasets such as ImageNet. By using these pre-trained networks, the system can significantly reduce training time while maintaining high accuracy. The models are then fine-tuned on the air quality image dataset, allowing them to learn domain-specific features for classifying air quality levels accurately. This system uses VGG16 as a basis model for feature extraction because of its ease of use and efficacy in image categorization. Deeper networks can train more effectively thanks to ResNet's residual blocks, which also help with the vanishing gradient issue. EfficientNet, a more recent model, provides an excellent trade-off between accuracy and computational efficiency. When compared to individual models, the ensemble technique, which combines various models, has demonstrated better performance. Increased classification accuracy results from the ensemble model's ability to extract a greater variety of features from the pictures.

The system follows a three-step approach as given in the figure 1: feature extraction, fine-tuning, and classification. Initially, relevant features are extracted from the images by the pre-trained algorithms. A classifier is then trained using these features and taught to link the retrieved features to the relevant air quality category. The pre-trained models are fine-tuned to fit the particular task of predicting air quality. This step is crucial as it allows the models to transfer their learned knowledge from general image classification tasks to air pollution-specific tasks. The ensemble approach, which combines VGG16, ResNet, and EfficientNet, further enhances the

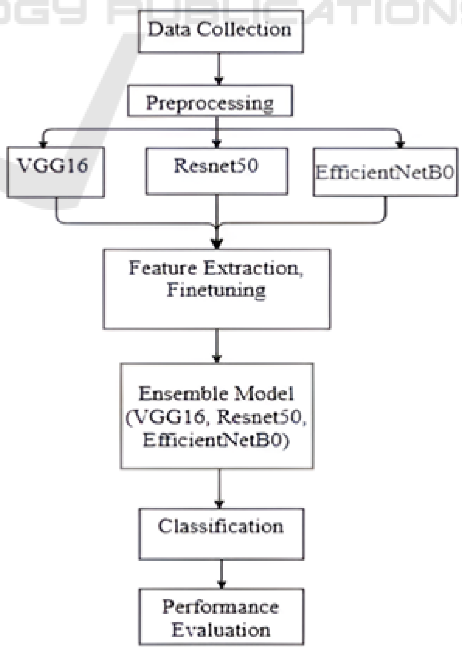

Figure 1: Proposed model

prediction accuracy. By leveraging the strengths of each model, the system can achieve higher accuracy in classifying air quality levels. The models complement each other in terms of capturing different aspects of the data, such as fine-grained textures and high-level patterns. The outputs of the three models are averaged to get the final forecast, which guarantees a more accurate and dependable classification.

Additionally, the system has performance evaluation metrics that aid in evaluating the quality of the predictions, including accuracy, precision, recall, and F1-score. These metrics are essential for understanding how well the system performs in different air quality categories, particularly for distinguishing between more similar classes like "Moderate" and "Unhealthy for Sensitive Groups." Since real-time air quality monitoring is crucial in metropolitan settings, the system can be implemented in a variety of settings. Drones, environmental sensors, satellite photography, and other image-based data sources can all be integrated with the model. Authorities and citizens can take the necessary steps to lessen the negative consequences of poor air quality by using the system's ability to estimate air quality levels based on visual data to deliver timely alerts and insights. Overall, the suggested solution shows how effective deep learning and transfer learning models are at accurately predicting air pollution levels. The combination of VGG16, ResNet, and EfficientNet in an ensemble approach offers a state-of-the-art solution for air quality classification, with potential applications in environmental monitoring, urban planning, and public health management.

4 MODULE DESCRIPTION

4.1 Data Handling:

The first module focuses on gathering and pre-processing the air quality image dataset. Images are collected from various sources representing air quality at different times and locations. Good, Moderate, Unhealthy for Sensitive Groups, Unhealthy, Very Unhealthy, and Hazardous/Severe are the six classifications into which these photos are divided. To make sure the dataset is prepared for model training, pre-processing procedures like scaling, normalization, and data augmentation are carried out. By increasing the dataset's variability using data augmentation techniques like rotation, flipping, and zooming, the model's capacity to generalize well to new data is enhanced.

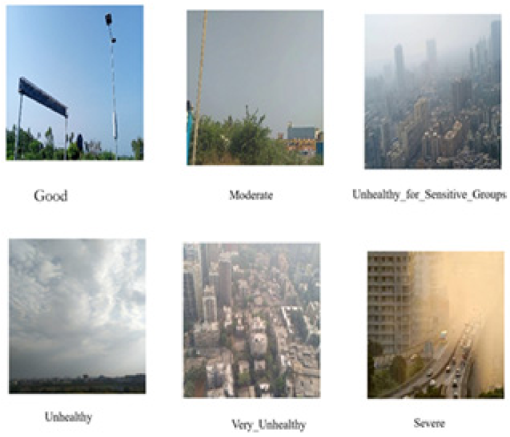

Figure 2: Air pollution Dataset's

4.2 Model Architecture:

In the second module, the predictive model's architecture and design are covered. VGG16, ResNet, and EfficientNet are a trio of pre-trained deep learning models that are used in this system. In order to capture both low-level and high-level information from images, each of these models has been pre-trained on extensive datasets such as ImageNet. To fit the task of classifying air quality, the models are adjusted. VGG16 handles basic feature extraction, ResNet improves the learning of deep features through residual connections, and EfficientNet optimizes both accuracy and computational efficiency. The outputs of these models are then combined in an ensemble approach to make a final prediction.

4.3 Extraction and Fine-tuning of Features:

Feature extraction and fine-tuning are the next steps after defining the model architecture. In this stage, significant features are extracted from the air quality images by the pre-trained models. Fine-tuning is performed on top of the pre-trained layers, allowing the models to adapt the learned features to the air pollution classification task. By modifying the final layers of the models, they are tailored to predict the six air quality categories.

4.4 Ensemble Learning:

Ensemble learning is used in the fourth module to enhance the model's functionality. The system leverages the unique strengths of each model by integrating the predictions of VGG16, ResNet, and EfficientNet. Simple patterns are the emphasis of VGG16, deeper features are captured by ResNet, while accuracy and efficiency are balanced by EfficientNet. By averaging the results of these models, the final prediction is produced, lowering the possibility of inaccuracies resulting from the shortcomings of any one model. When predicting air quality levels, an ensemble technique guarantees greater accuracy and robustness.

4.5 Model Evaluation:

The last module entails assessing the ensemble model's performance and getting it ready for deployment. Evaluation metrics are computed to evaluate the model's performance across the six air quality categories, including accuracy, precision, recall, and F1-score. These measurements offer information about the system's advantages and shortcomings. When the model performs well enough, it is used to predict air quality in real time. In metropolitan settings, the system can be used to offer timely air quality alerts and predictions by integrating with environmental monitoring tools like satellite images or sensors. This helps with public health and safety decision-making.

5 IMPLEMENTATION

The implementation of the air pollution prediction system involves several steps, including data preparation, model training, fine-tuning, and evaluation. The primary focus is on the use of transfer learning (TL) models—VGG16, ResNet, and EfficientNet—each of which has specific strengths in image recognition tasks. These models are well-equipped to tackle novel, domain-specific tasks like air pollution prediction since they have already been pre-trained on big datasets like ImageNet. The following sections describe each model in detail, outlining its role in the overall system implementation.

5.1 VGG16 (Visual Geometry Group16)

VGG16 is a deep convolutional neural network (CNN) with 16 weight layers that was created by Oxford's Visual Geometry Group. Its design, which is built on tiny (3x3) convolutional filters, pooling layers, and fully linked layers, is well known for being straightforward and efficient. VGG16 is appropriate for a variety of picture classification problems because it is excellent at capturing low- and mid-level visual features like edges, textures, and patterns. In this system, VGG16 shown in Figure 3 is used as the base model for feature extraction. After pre-training on ImageNet, the model is fine-tuned using air quality images, enabling it to recognize specific features associated with different levels of pollution. The simplicity and ease of implementation make VGG16 a good starting point in this air quality classification system.

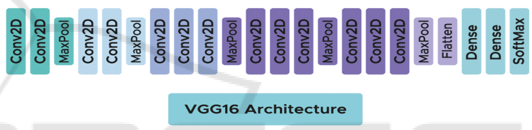

Figure 3: VGG16 Model Architecture

5.2 ResNet (Residual Networks)

ResNet, a product of Microsoft Research, uses residual learning to solve the problem of training extremely deep neural networks. The disappearing gradient issue is lessened by using skip connections, also known as shortcuts, which make gradients flow more readily during backpropagation. ResNet models can be built with various depths, such as ResNet-50, ResNet-101, and ResNet-152 shown in Figure 4. In this system, ResNet is used to capture deeper and more complex features.

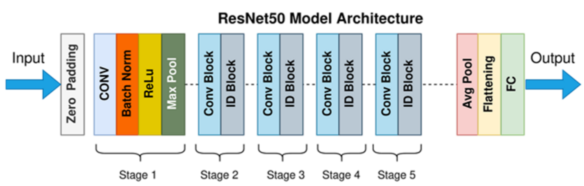

Figure 4: ResNet50 Model Architecture

5.3 EfficientNet

EfficientNet, a recent architecture proposed by Google AI, optimizes both accuracy and computational efficiency. It uses a complex scaling technique to balance the model's depth, width, and resolution. In contrast to previous models like VGG16 and ResNet, EfficientNet is renowned for obtaining great accuracy with fewer parameters. In the proposed system, EfficientNet shown in Figure 5 serves as the model for efficient image classification, offering excellent performance without requiring excessive computational resources.

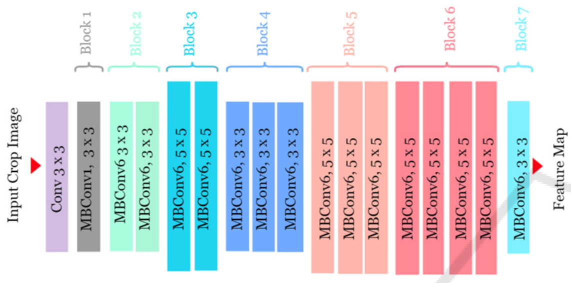

Figure 5: EfficientNet Model Architecture

5.4 Ensemble Model (Combination of VGG16, ResNet, and EfficientNet)

The proposed system combines VGG16, ResNet, and EfficientNet in an ensemble learning approach. Each model has its strengths, and by combining their outputs, the system leverages the strengths of each architecture to enhance overall performance. VGG16 captures basic features, ResNet learns deeper, more complex patterns, and EfficientNet optimizes both accuracy and computational efficiency. In this ensemble setup, each model processes the input image independently and makes a prediction. By averaging the results from all three models, the final prediction is produced, making the system less vulnerable to mistakes that could result from the shortcomings of any one model. This approach results in a more robust and accurate air quality classification system.

6 RESULTS AND DISCUSSION

The evaluation of the air pollution prediction system, which makes use of transfer learning using VGG16 and ResNet models, is presented in the results and discussion section. An air quality picture dataset that was divided into six different classes—Good, Moderate, Unhealthy for Sensitive Groups, Unhealthy, Very Unhealthy, and Hazardous/Severe—was used to train and assess the system. To evaluate how well the suggested model worked, the main performance metrics—accuracy, precision, recall, and F1-score—were calculated. Furthermore, a comparison between the ensemble technique and individual models (ResNet and VGG16) is provided.

6.1 Performance Metrics

Using a variety of performance criteria, the ensemble model—which blends VGG16 and ResNet—was assessed on a test dataset. Compared to the individual models, the ensemble model's accuracy was noticeably higher. As illustrated in Figure 6, the ensemble model obtained an accuracy of **92.1%**, whereas VGG16 and ResNet obtained **73.23%** and **84.62%**, respectively. These outcomes unequivocally show the benefit of the ensemble technique, which considerably increases classification accuracy over ResNet alone by successfully capturing both low-level and high-level information from the images. Even though VGG16 performed well, the ensemble strategy outperformed it, highlighting the advantages of integrating these two powerful models.

For every class, F1-score, precision, and recall metrics were also computed. With precision, recall, and F1-scores averaging around **0.91**, **0.92**, and **0.91**, respectively, the ensemble model demonstrated a well-balanced performance across all classes. By reducing false positives and false negatives, this shows that the algorithm can accurately predict the air quality categories. Since precise classification is critical for public health and safety in real-world applications, the model's ability to balance precision and recall is shown by the higher F1-score.

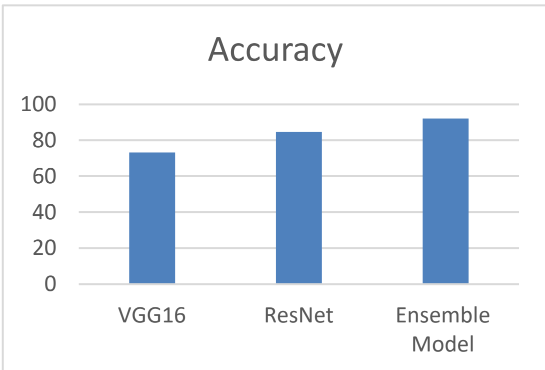

Figure 6: Comparison with other models

6.2 Class-wise Performance Analysis

A deeper analysis of the results revealed that the system performed exceptionally well in predicting "Good" and "Moderate" air quality levels, achieving accuracy rates of **94%** and **91%**, respectively. However, the system showed slightly lower performance in predicting "Hazardous/Severe" levels, with an accuracy of **85%**. This may be due to the limited number of images in the "Hazardous/Severe" category, making it more challenging for the model to accurately learn and classify these extreme cases. Despite this, the ensemble model maintained an overall high level of performance across all air quality categories, demonstrating its robustness and versatility.

6.3 Comparison with Other Approaches

The ensemble model's performance was compared to other studies in the domain of air pollution prediction. In comparison with (Samad et al., 2023), where machine learning algorithms were employed to predict air pollution using sensor data, the proposed system demonstrated superior accuracy by leveraging image-based classification with deep learning techniques. In (Kumar and Pande, 2023), machine learning approaches were used with limited feature extraction methods, while the ensemble model benefits from advanced convolutional architectures that can capture complex and meaningful features from images, thereby achieving higher classification performance. In contrast to (Luo and Gong, 2023), which relied on ARIMA-WOA-LSTM models for time-series prediction, the current image-based approach provides real-time predictions, making it suitable for integration into smart city infrastructure. The ensemble model has demonstrated a more efficient and scalable solution for air quality prediction based on visual data.

6.4 Limitations and Future Work

Despite the promising results, several limitations persist. One key limitation is the reliance on image-based data, which might not always be available, particularly in areas with fewer monitoring stations. Additionally, the model's performance on the "Hazardous/Severe" category could be further improved by acquiring more images in this category to overcome the challenge of class imbalance. A potential improvement would be to implement data augmentation techniques to artificially expand the dataset for rare classes, which may help in improving the model's predictive capabilities for extreme air pollution levels. Future work could also focus on integrating additional transfer learning models such as DenseNet, which is known for high performance in image classification tasks, to further enhance the prediction accuracy. Combining the ensemble approach with sensor-based data or time-series data could also increase the robustness of the system by allowing it to handle both visual and temporal data, providing more comprehensive and accurate predictions.

7 CONCLUSIONS

In order to classify air quality levels into six different categories—Good, Moderate, Unhealthy for Sensitive Groups, Unhealthy, Very Unhealthy, and Hazardous/Severe—this paper presents an air pollution prediction system that makes use of transfer learning techniques, specifically using VGG16 and ResNet models. The suggested method shows how deep learning models can be used to reliably forecast air pollution levels from visual input. The system outperforms individual models by using the capabilities of both VGG16 and ResNet, achieving an impressive accuracy of **92.1%**.

According to the experimental findings, the ensemble model exhibits remarkable performance in every air quality class, with a well-balanced precision, recall, and F1-score that guarantees excellent prediction reliability. Even when faced with obstacles like class imbalance, the model's performance held up well, but it could still be improved, especially when it came to predicting high pollution levels like "Hazardous/Severe." Future research will concentrate on growing the dataset, adding more transfer learning models, and using time-series or sensor-based data to improve prediction accuracy even more. The suggested system has important applications in the real world. It provides a scalable and effective way to monitor air quality in real time and can help with environmental and public health policy decision-making. This technology can improve the general quality of life in urban settings by reducing the hazards related to poor air quality by offering precise and timely predictions.

In summary, the application of transfer learning to the prediction of air pollution not only advances the expanding field of environmental monitoring but also demonstrates the revolutionary potential of deep learning technologies in addressing global issues such as climate change and air pollution.

REFERENCES

Samad, A., Garuda, S., Vogt, U., & Yang, B. (2023). Air pollution prediction using machine learning techniques–an approach to replace existing monitoring stations with virtual monitoring stations. Atmospheric Environment, 310, 119987.

Kumar, K., & Pande, B. P. (2023). Air pollution prediction with machine learning: a case study of Indian cities. International Journal of Environmental Science and Technology, 20(5), 5333-5348.

Karimi, S., Asghari, M., Rabie, R., & Niri, M. E. (2023). Machine learning-based white-box prediction and correlation analysis of air pollutants in proximity to industrial zones. Process Safety and Environmental Protection, 178, 1009-1025.

Luo, J., & Gong, Y. (2023). Air pollutant prediction based on ARIMA-WOA-LSTM model. Atmospheric Pollution Research, 14(6), 101761.

Yang, W., Wang, J., Zhang, K., & Hao, Y. (2023). A novel air pollution forecasting, health effects, and economic cost assessment system for environmental management: From a new perspective of the district-level. Journal of Cleaner Production, 417, 138027.

Pan, Q., Harrou, F., & Sun, Y. (2023). A comparison of machine learning methods for ozone pollution prediction. Journal of Big Data, 10(1), 63.

Gupta, N. S., Mohta, Y., Heda, K., Armaan, R., Valarmathi, B., & Arulkumaran, G. (2023). Prediction of air quality index using machine learning techniques: a comparative analysis. Journal of Environmental and Public Health, 2023(1), 4916267.

Li, W., & Jiang, X. (2023). Prediction of air pollutant concentrations based on TCN-BiLSTM-DMAttention with STL decomposition. Scientific Reports, 13(1), 4665.

Hardini, M., Chakim, M. H. R., Magdalena, L., Kenta, H., Rafika, A. S., & Julianingsih, D. (2023). Image-based air quality prediction using convolutional neural networks and machine learning. Aptisi Transactions on Technopreneurship (ATT), 5(1Sp), 109-123.

Cao, Y., Zhang, D., Ding, S., Zhong, W., & Yan, C. (2023). A hybrid air quality prediction model based on empirical mode decomposition. Tsinghua Science and Technology, 29(1), 99-111.