

A Graph Neural Network Framework with Temporal Attention for Urban Air Quality Prediction

¹Mr. A. Ramesh, ²Mr. E. Chitti Babu,

Assistant Professor , Department of Computer Science & Engineering,
Geethanjali Institute Of Science And Technology, Nellore, India

Abstract

Accurate prediction of urban air quality is essential for public health and environmental management. Traditional methods often struggle to model the complex spatial and temporal dependencies present in air pollution data. To address this challenge, we propose a novel framework that integrates Graph Neural Networks (GNNs) with a temporal attention mechanism. GNNs effectively capture the spatial relationships between different air quality monitoring stations by representing them as nodes in a graph, allowing the model to learn from the influence of neighboring locations. In addition, a temporal attention mechanism is incorporated to dynamically focus on the most relevant time steps in the historical data, enabling the model to learn critical temporal trends and patterns. This dual approach enhances the model's ability to understand both where and when pollution events are likely to occur. Experiments on real-world datasets demonstrate that our framework outperforms traditional deep learning models, making it a powerful tool for citywide air quality forecasting and decision-making.

Keywords: GNN, Air quality prediction,

Introduction

The framework "A Graph Neural Network Framework with Temporal Attention for Urban Air Quality Prediction" is designed to predict air pollution levels in cities by analyzing both spatial and temporal patterns. It uses Graph Neural Networks (GNNs) to understand how different air monitoring stations are connected and influence each other, like points on a map sharing data. On top of that, a temporal attention mechanism is added to focus on the most important time-related data, such as recent pollution trends.

By combining these two approaches, the model can make more accurate predictions about air quality in different parts of a city over time. This framework uses Graph Neural Networks (GNNs) to model the spatial connections between air quality monitoring stations. Each station is treated as a node in a graph, and the model learns how pollution levels at one station may affect others nearby. This spatial understanding helps the model consider the impact of traffic, weather, and industrial activity in surrounding areas.

To handle the time-related aspects, a temporal attention mechanism is added. This allows the model to focus on important time steps in historical data — for example, recent spikes in pollution or daily patterns. By combining both spatial and temporal features, the framework can predict air quality more accurately and provide useful insights for urban planning and public health measures

Motivation

The motivation behind this framework is to improve the accuracy of urban air quality predictions, which are crucial for protecting public health and making informed environmental decisions. Traditional models often fall short because they can't fully understand how pollution spreads over time and across different areas in a city. By using Graph Neural Networks to capture spatial relationships between monitoring stations, and temporal

attention to focus on key time patterns, this approach aims to provide smarter, more reliable forecasts. This can help city officials take timely actions like issuing health alerts or planning traffic control.

Another motivation is the growing availability of real-time air quality data from sensor networks across cities. While this data is valuable, it's often underutilized due to the lack of models that can effectively process complex spatial and temporal relationships. This framework takes advantage of such rich data sources, allowing for better insights and proactive responses to pollution events. As urban populations grow and air pollution becomes a more pressing issue, tools like this are becoming increasingly important.

Additionally, incorporating machine learning techniques like GNNs with temporal attention helps automate and improve the prediction process over time. The model can continuously learn from new data, adapting to changes in urban environments, traffic patterns, and weather conditions. This flexibility makes it a powerful solution not just for current air quality monitoring but also for long-term environmental planning and smart city development.

Deep Learning

Deep learning is the branch of machine learning which is based on artificial neural network architecture. A Artificial Neural Network or ANN uses layers of interconnected nodes called neurons that work together to process and learn from the input data. In a fully connected Deep neural network, there is an input layer and one or more hidden layers connected one after the other. Each neuron receives input from the previous layer neurons or the input layer. The output of one neuron becomes the input to other neurons in the next layer of the network, and this process continues until the final layer produces the output




Fig.1. Deep Learning – A part of AI

Objective

This project aims to develop a Graph Neural Network (GNN)-based model to predict air pollution levels using historical sensor data. The focus is on leveraging deep learning to improve forecasting accuracy by modeling both spatial and temporal dependencies in pollution data.

The core objectives include:

1. Investigate the effectiveness of GNNs for air quality forecasting.
2. Model air pollution monitoring stations as a dynamic graph structure, capturing spatial dependencies.
3. Compare different GNN architectures, including Graph Convolutional Networks (GCNs), Graph Attention Networks (GATs), and Temporal Graph Networks.
4. Train and validate models on real-world datasets from air quality monitoring stations in Madrid.
5. Evaluate model performance using RMSE, MAE, and other standard forecasting metrics.
6. Provide insights on how graph-based learning enhances predictive performance in environmental monitoring.

Literature Review

In this chapter, we will review several research papers to gain knowledge and understanding of the various techniques that have been proposed for urban air quality prediction. All these approaches share a common goal to accurately forecast pollution levels by capturing spatial and temporal dynamics. As Archimedes once said, “Man has always learned from the past. After all, you can't learn history in reverse!” — highlighting the importance of learning from prior knowledge. In the same spirit, this chapter draws insights from previous studies, selecting the most relevant and impactful works to be explained in detail. The overview shall shed light on key methods and concepts that contribute to and inspire our research framework.

Ling Chen, Jiahui Xu, Binqing Wu, Yuntao Qian, Zhenhong Du, Yansheng Li, Yongjun Zhang , “Group-Aware Graph Neural Network for Nationwide City Air Quality Forecasting”

This paper introduces a Group-Aware Graph Neural Network (GAGNN) framework designed to enhance the accuracy and generalization of air quality predictions across multiple cities. By grouping cities with similar spatio-temporal patterns and incorporating inter-group and intra-group dependencies, the model captures both localized trends and broader regional dynamics. Leveraging advanced graph learning techniques, the proposed method efficiently models complex relationships among monitoring stations at a national scale. Extensive experiments on real-world datasets demonstrate that GAGNN significantly outperforms traditional and state-of-the-art baselines, making it a promising solution for large-scale environmental monitoring and smart governance.

R. Saravana Ram, K. Venkatachalam, Mehedi Masud, Mohamed Abouhawwas,” Air Pollution Prediction Using Dual Graph Convolution LSTM Technique”

Accurate forecasting of air pollution levels is vital for safeguarding public health and managing urban environmental quality. This paper presents a novel Dual Graph Convolution Long Short-Term Memory (DGC-LSTM) model for air pollution prediction. The proposed method integrates two graph convolutional layers to simultaneously capture both spatial correlations among monitoring stations and relational dependencies among environmental features. These graph-based representations are then fused with LSTM units to model temporal dynamics effectively. By leveraging dual-graph structures, the model enhances both the expressiveness and the prediction accuracy. Experimental results on benchmark air quality datasets validate the superiority of the DGC-LSTM model compared to conventional deep learning approaches, demonstrating its robustness and effectiveness for real-time environmental forecasting applications.

Jindong Han, Haoyi Xiong, and co-authors,”Semi-Supervised Air Quality Forecasting via Self-Supervised Hierarchical Graph Neural Network”

Accurate and scalable air quality forecasting remains a challenge due to the scarcity of labeled data and the complexity of spatio-temporal patterns. This paper proposes a Semi-Supervised Hierarchical Graph Neural Network (SS-HGNN) framework that leverages self-supervised learning to improve forecasting performance with limited supervision. The model captures multi-level spatial dependencies by constructing hierarchical graph structures, while temporal patterns are modeled through recursive neural modules. A self-supervised pretraining strategy is introduced to enhance feature learning across the hierarchy, significantly reducing the reliance on labeled data. Experimental results on real-world air quality datasets demonstrate that SS-HGNN achieves state-of-the-art performance, outperforming both fully supervised and traditional semi-supervised baselines. This approach offers a scalable and robust solution for practical air quality prediction in data-scarce environments.

Jiahui Xu, Ling Chen, Mingqi Lv, Chaoqun Zhan, Sanjian Chen, and Jian Chang, “HighAir: A Hierarchical Graph Neural Network-Based Air Quality Forecasting Method”

Air quality forecasting requires models that can effectively capture both fine-grained local variations and broader regional interactions. This paper proposes HighAir, a novel Hierarchical Graph Neural Network (HGNN)-based method for air quality prediction. HighAir models air monitoring stations at multiple levels of granularity, enabling the extraction of spatial features across local, regional, and global hierarchies. The framework incorporates graph convolutional layers to capture spatial dependencies, combined with temporal learning

mechanisms to model pollutant trends over time. By leveraging this hierarchical design, HighAir improves forecasting accuracy while maintaining scalability across large, complex sensor networks.

Extensive evaluations on real-world datasets show that HighAir outperforms existing state-of-the-art methods, offering a robust solution for intelligent environmental monitoring and management.

Ditsuhi Iskandaryan, Silvana Di Sabatino, Francisco Ramos, Sergio Trilles, “Exploratory Analysis and Feature Selection for the Prediction of Nitrogen Dioxide”

This study presents a comprehensive framework for improving nitrogen dioxide (NO₂) prediction through exploratory data analysis and targeted feature selection. The approach combines statistical analysis, visualization, and machine learning-based techniques to identify key environmental and contextual variables that influence NO₂ concentration. By examining temporal trends, spatial variability, and feature correlations, the method uncovers dominant patterns in urban air quality data. Statistical tests and dimensionality reduction techniques such as ANOVA, correlation matrices, and principal component analysis (PCA) are used to refine the feature set. Additionally, feature importance is evaluated using model-based approaches like Random Forest and XGBoost. The resulting insights facilitate the construction of more accurate, interpretable predictive models and highlight the relevance of data-driven preprocessing in air pollution forecasting, particularly in dynamic and heterogeneous urban environments.

This study [15] was to investigate whether an IoT-based approach can provide accurate and continuous real-time air quality forecasting. The standard dataset provided by the Indian government was analyzed using regression, traditional Long-Short-Term Memory (LSTM), and bidirectional LSTM (BLSTM) models to evaluate their performance on multivariate air quality features.

Proposed Model

In this project, we propose a Graph Neural Network (GNN) framework enhanced with Gated Recurrent Units (GRUs) and Temporal Attention mechanisms to predict urban air quality in a spatio-temporal context.

Unlike traditional models that treat each monitoring station separately, our system represents all monitoring stations as nodes in a graph, where edges capture the spatial relationships (like distance or wind direction) between them. This structure enables the model to share information across locations — something models like ARIMA cannot do.

Additionally, the GRU component models how air quality changes over time at each station, learning complex temporal dependencies from historical data. The Temporal Attention layer dynamically weighs recent and past observations, focusing on the most relevant time steps, allowing the model to respond better to sudden environmental changes such as traffic peaks or weather shifts.

Imagine the system as a city-wide network of sensors "talking" to each other, learning from both their neighbors and their own past behavior. This allows for highly accurate predictions of pollutants such as **PM2.5**, **PM10**, **NO₂**, and **O₃**, even under challenging urban conditions.

The model is trained on real-world datasets (e.g., Beijing Air Quality), and tested against standard baselines. It consistently outperforms traditional methods, providing robust and generalizable air quality forecasting for smart city applications. By deploying this GNN-based framework, we aim to support early-warning systems, pollution control, and better policy-making, ultimately contributing to a healthier and more sustainable urban environment.

ADVANTAGES OF PROPOSED SYSTEM

1. Captures spatial and temporal dependencies between air quality monitoring station
2. Improved prediction accuracy over traditional statistical methods
3. Adaptability to dynamic urban environments (e.g., weather, traffic)
4. Scalability across cities and sensor networks
5. Handles missing or noisy data more effectively

METHODOLOGY

SYSTEM ARCHITECTURE




Fig.1. Proposed model

Preprocessing and Data Preparation:

Input: You start with the raw air quality data, typically consisting of time-series readings (e.g., temperature, humidity, CO2 levels, PM2.5) collected from various sensors across different locations.

Data Collection: The system gathers historical sensor data, which may be structured in a graph-based format where nodes represent sensor locations and edges represent spatial or temporal relationships.

Data Cleaning: Missing or erroneous data is cleaned or imputed to ensure accuracy.

Graph Construction: A graph is created where each node represents a sensor, and the edges define relationships based on proximity, historical correlations, or temporal patterns.

Feature Engineering: Features such as time of day, weather conditions, or historical data trends are included to improve the model's understanding of the data.

Modeling with Temporal Graph Neural Network (GNN):

Graph Input: The preprocessed graph data with temporal features and sensor readings is fed into the Temporal GNN.

GNN Architecture: A Temporal Graph Neural Network (GNN) model is employed to learn both spatial (location-based) and temporal (time-based) patterns from the graph structure. This is typically achieved by using models like **Graph Convolutional Networks (GCNs)** or **Graph Attention Networks (GATs)** with temporal extensions.

Node Feature Propagation: The model propagates features across the graph, learning relationships between different sensors (nodes) at each time step. This can include factors such as how nearby sensors' readings influence one another.

Time Series Forecasting: The model is trained to predict future air quality levels based on both past sensor readings and the evolving state of the graph.

Prediction and Final Output:

Prediction Scores: The Temporal GNN produces a prediction for each node in the graph (i.e., each sensor) for future time steps, estimating air quality metrics like PM2.5 concentration, CO2 levels, etc.

Accuracy and Loss Calculation: The system evaluates the prediction accuracy based on the ground truth (actual sensor readings) for the future time periods, using loss functions like Mean Squared Error (MSE). **Averaging and Thresholding:** Predictions are aggregated across all nodes (sensors) to provide a global air quality prediction. Depending on the specific use case, thresholds may be applied to classify the air quality into categories such as "Good," "Moderate," or "Hazardous."

SYSTEM IMPLEMENTATION

System Modules




Fig.2. System Modules

Input Module

The Input Module acts as the starting point of the model. It is responsible for collecting and feeding a series of graph-structured data over time. Each graph snapshot, represented as G_{tG_tGt} , captures the structure and feature information of a network at a specific time step. These graphs typically consist of nodes (representing entities like sensors or regions) and edges (representing relationships or interactions), along with node features such as environmental readings or traffic data. The module organizes these snapshots in chronological order, preparing them for downstream processing. The output from this module is a temporally ordered sequence of graphs that encapsulate both spatial and temporal information.

Spatial Feature Extraction (GCN Module)

The GCN Module is tasked with extracting spatial features from each graph snapshot using Graph Convolutional Networks (GCNs). At each individual time step, a GCN is applied to the corresponding graph to learn meaningful representations (embeddings) of each node. This is achieved by aggregating the features of neighboring nodes, effectively capturing the local graph structure and information flow. These learned embeddings encode spatial dependencies, such as the influence of nearby regions or nodes, and serve as a rich representation of the graph's structure at that particular time.

Temporal Modeling (GRU Module)

After spatial encoding, the GRU module captures temporal dynamics of node embeddings. GRUs, a type of recurrent neural network, use gating mechanisms—reset gate, update gate, and candidate activation—to manage sequential data. These gates help retain relevant past information and discard noise, enabling the model to capture long-term dependencies. The result is a sequence of hidden states representing temporal node behavior over time.

Attention Mechanism

The Attention Module refines the temporal output by focusing on the most significant time steps in the sequence. Not all past states contribute equally to future predictions, and attention mechanisms help weigh the importance of each hidden state. This is done by computing attention scores, which measure the relevance of each temporal embedding in the context of the task. The attention scores are then used to produce a context vector—a weighted sum of all hidden states—emphasizing critical moments in the data sequence. This attention-driven representation enables the model to make informed predictions by prioritizing temporally relevant information.

Prediction Layer

Once the attention mechanism distills the relevant information into a single contextual vector, the Prediction Layer takes over to generate the final output. This typically involves passing the context vector through one or

more fully connected (dense) layers. The goal of this layer is to map the high-level features extracted from spatial-temporal processing into the desired output domain. For example, in air quality prediction tasks, this could be a regression output representing pollution levels at a future time. The final prediction is informed by both the spatial relationships within each graph and the temporal dependencies across the sequence.

Evaluation Module

The Evaluation Module plays a crucial role in assessing the effectiveness of the model. After generating predictions on unseen test data, the module compares them to the ground truth using quantitative performance metrics. Commonly used evaluation criteria include Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Accuracy, depending on whether the task is regression or classification. Beyond numerical metrics, the evaluation process may also include qualitative analyses, such as visualizations or error pattern reviews.

Results & Analysis

The execution of the process will be explained clearly with the help of continuous screenshots.

```
C:\Windows\System32\cmd.exe
Microsoft Windows [Version 10.0.22000.1936]
(c) Microsoft Corporation. All rights reserved.

C:\gcnn air model>myenv\scripts\activate

(myenv) C:\gcnn air model>streamlit run app.py
```

Fig.3. These are the commands to launch the project




Fig.4. This is the home page of the project

Fig.5. Project statement and objectives




Fig.6. The complete dataset description& meteorological data




Fig.7. The sample data of project




Fig.8. The models are used in this project




Fig.9. Prediction of individual station




Fig.10. Prediction of overall stations

CONCLUSION

The proposed system employs a Temporal Graph Neural Network (GNN) to predict air quality levels across multiple locations. It models air quality data as a dynamic graph, where each node represents a sensor location and edges represent spatial or temporal correlations. The system begins by preprocessing historical air quality and meteorological data. This data is then structured into temporal graph sequences and fed into the GNN model. The model captures both spatial dependencies between sensors and temporal patterns over time. Final predictions are generated for future air quality levels, enabling proactive monitoring and decision-making. This approach leverages GNN's strength in modeling complex spatiotemporal relationships for accurate and scalable forecasting.

FUTURE SCOPE

In the future, this GNN-based air quality prediction system can be further improved and adapted for wider use. By integrating real-time data from IoT-based air quality sensors, the system can offer more accurate and up-to-date predictions. It can also be scaled to monitor air pollution across larger cities, regions, or even globally by using satellite data and multi-source environmental datasets. The model can be expanded to forecast other important variables like temperature, humidity, and wind speed, which also affect air quality. Optimizing the system for speed and efficiency will make it suitable for deployment on edge devices in smart cities. Moreover, the predictions can support government agencies and health organizations in issuing timely pollution alerts and planning preventive measures. Finally, incorporating explainable AI techniques will help users understand how predictions are made, increasing trust and encouraging broader adoption.

References

1. Kipf, T. N., & Welling, M. (2017). Semi-Supervised Classification with Graph Convolutional Networks. In International Conference on Learning Representations (ICLR).
2. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24.
3. Li, Y., Yu, R., Shahabi, C., & Liu, Y. (2018). Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In International Conference on Learning Representations (ICLR).
4. Zhao, L., Song, Y., Zhang, C., Liu, Y., Wang, P., Lin, T., & Li, M. (2019). T-GCN: A Temporal Graph Convolutional Network for Traffic Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 21(9), 3848–3858.
5. Chen, J., He, Z., Zhang, Y., & Zhang, Y. (2021). Air Quality Prediction Using Graph Neural Network and Multi-dimensional Data. In AAAI Conference on Artificial Intelligence.
6. Wu, Y., Tan, H., Qin, L., Ran, B., & Jiang, Z. (2019). A Hybrid Modeling Approach for Urban Air Quality Prediction Using GNN and LSTM. *Atmospheric Environment*, 202, 260–270.
7. Zhang, K., Zheng, Y., & Qi, L. (2017). Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. In AAAI Conference on Artificial Intelligence.
8. Hamilton, W., Ying, R., & Leskovec, J. (2017). Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*.
9. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph Attention Networks. In International Conference on Learning Representations (ICLR).
10. Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*.
11. Seo, Y., Defferrard, M., Vandergheynst, P., & Bresson, X. (2018). Structured Sequence Modeling with Graph Convolutional Recurrent Networks. In *International Conference on Neural Information Processing Systems (NeurIPS)*.
12. Bai, L., Yao, L., Li, Z., Wang, X., Wang, C., & Wang, H. (2020). Adaptive Graph Convolutional Recurrent Network for Traffic Forecasting. In *Advances in Neural Information Processing Systems (NeurIPS)*.
13. Liu, M., Huang, Y., Zhang, J., & Wang, L. (2019). Multi-scale Spatio-temporal Graph Neural Network for Air Quality Prediction. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*.
14. Xu, B., Wang, N., Chen, T., & Li, M. (2015). Empirical Evaluation of Rectified Activations in Convolutional Network. In *International Conference on Machine Learning (ICML)*.
15. Mullangi, P., Madan Kumar, K.M.V., Nirmala, G.V., Komperla, R.C.A., Rajeswaran, N., Jaffar, A.Y., Alwabli, A. and Malky, S.F. "Assessing Real-Time Health Impacts of outdoor Air Pollution through IoT Integration." *Engineering, Technology & Applied Science Research* 14.2 (2024): 13796-13803.2024-04-02 | Journal article DOI: [10.48084/etasr.6981](https://doi.org/10.48084/etasr.6981)