

A Review on Deep Learning Applications in Air Pollution Control Methods

Kiana Azimi *

Caspian Faculty of Engineering, College of Engineering, University of Tehran, Rezvanshahr, Iran.

Ehsan Ghiasy Nick

Caspian Faculty of Engineering, College of Engineering, University of Tehran, Rezvanshahr, Iran.

Ali Ahmadi Orkomi

Department of Environmental Science and Engineering, Faculty of Natural Resources, University of Guilan, Rasht, Iran.

Abstract

This study delves into the transformative role of deep learning and neural networks in the domain of air pollution control. By focusing on enhanced detection and monitoring, particularly through convolutional and recurrent neural architectures, the research highlights the potential of these technologies to unravel complex patterns within air quality dynamics. Beyond mere detection, these models demonstrate proactive capabilities, enabling the prediction and forecasting of pollution events. This foresight empowers the implementation of adaptive control strategies, effectively minimizing health risks and optimizing resource allocation. However, the study acknowledges challenges related to data quality and interpretability, emphasizing the necessity for interdisciplinary collaboration among machine learning experts, environmental scientists, and policymakers. In synthesizing these findings, the research contributes to the advancement of sustainable strategies for mitigating the impact of air pollution on human health and the environment and also reviews methods of controlling it by deep learning approaches.

Keywords: deep learning, neural network, air pollution, pollution detection

* Corresponding Author: kiana.azimi@ut.ac.ir

مروری بر کاربرد یادگیری عمیق در روش‌های کنترل آلودگی هوا

کیانا عظیمی *

دانشکده فنی کاسپین، دانشکده گان فنی، دانشگاه تهران، رضوانشهر ایران.

احسان غیائی نیک

دانشکده فنی کاسپین، دانشکده گان فنی، دانشگاه تهران، رضوانشهر ایران.

علی احمدی ارکمی

گروه علوم و مهندسی محیط زیست، دانشکده منابع طبیعی، دانشگاه گیلان، رشت، ایران.

چکیده

این مطالعه به بررسی نقش تحولی یادگیری عمیق و شبکه‌های عصبی در حوزه کنترل آلودگی هوا می‌پردازد. با تمرکز بر تشخیص و نظارت پیشرفته، به‌ویژه از طریق معماری‌های شبکه عصبی کانولوشنی و بازگشتی، این تحقیق پتانسیل این فناوری‌ها را برای کشف الگوهای پیچیده در دینامیک کیفیت هوا نشان می‌دهد. فراتر از صرفاً تشخیص، این مدل‌ها توانایی‌های پیش‌دستی را نشان می‌دهند که پیش‌بینی و پیش‌بینی وقایع آلودگی را ممکن می‌سازد. این پیش‌بینی‌ها، پیاده‌سازی استراتژی‌های کنترلی تطبیقی را ممکن می‌سازد که به طور مؤثری خطرات بهداشتی را کاهش داده و تخصیص منابع را بهینه می‌کنند. با این حال، مطالعه به چالش‌های مربوط به کیفیت داده‌ها و قابلیت تفسیر اشاره دارد و بر لزوم همکاری بین‌رشته‌ای میان کارشناسان یادگیری ماشین، دانشمندان محیط زیست و سیاست‌گذاران تأکید می‌کند. با ترکیب این یافته‌ها، تحقیق به پیشرفت استراتژی‌های پایدار برای کاهش تأثیر آلودگی هوا بر سلامت انسان و محیط زیست کمک می‌کند و همچنین روش‌های کنترل آن را از طریق رویکردهای یادگیری عمیق بررسی می‌نماید.

کلیدواژه‌ها: یادگیری عمیق، شبکه عصبی، آلودگی هوا، تشخیص آلودگی

Introduction

The escalating concerns surrounding air pollution necessitate innovative and efficient strategies for monitoring, detection, and control. In recent years, deep learning, a subset of artificial intelligence, has emerged as a powerful tool for addressing complex environmental challenges. This paper aims to explore the burgeoning intersection of deep learning, neural networks, and air pollution control. The integration of deep learning techniques, particularly neural networks, holds great promise in revolutionizing our ability to accurately detect pollutants in the air and implement targeted mitigation strategies.

Deep learning, characterized by its ability to automatically learn hierarchical representations from data, has demonstrated unparalleled success in diverse domains, ranging from computer vision to natural language processing. The application of deep learning methodologies in environmental sciences, and specifically in air quality management, provides an unprecedented opportunity to analyze vast datasets and extract meaningful insights. Neural networks, as a fundamental component of deep learning, offer the capacity to discern intricate patterns and relationships within complex atmospheric data, paving the way for more robust air pollution monitoring systems (Csaji, 2001).

The detection of air pollutants requires advanced sensor technologies and data processing techniques. Traditional monitoring approaches, while effective, often face limitations in terms of spatial coverage, temporal resolution, and the ability to handle diverse pollutants. Deep learning models, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), exhibit remarkable capabilities in learning spatial and temporal dependencies within environmental data. These models can enhance the accuracy of pollutant concentration predictions and identify subtle patterns that might elude traditional methodologies, thereby contributing to a more nuanced understanding of air quality dynamics.

Moreover, the deployment of deep learning in air pollution control extends beyond mere detection to encompass the prediction and forecasting of pollution events. Neural networks, through their ability to capture non-linear relationships, can predict pollution levels with higher accuracy, enabling proactive interventions and adaptive control strategies. This anticipatory approach is vital for minimizing the adverse health effects associated with exposure to elevated pollutant levels and for optimizing resource allocation in pollution abatement efforts (Xu et al., 2015).

Despite the promising advancements, challenges persist in implementing deep learning models for air pollution control. Issues related to data quality, model interpretability, and computational requirements necessitate interdisciplinary collaborations between experts in machine learning, environmental science, and policy-making. Overcoming these challenges is imperative to harness the full potential of deep learning in developing effective and sustainable air pollution control strategies.

In conclusion, this paper seeks to unravel the multifaceted applications of deep learning, particularly neural networks, in the realm of air pollution control. As society grapples with the consequences of environmental degradation, the integration of advanced technologies becomes paramount. By leveraging the capabilities of deep learning, we aim to enhance our understanding of air quality dynamics, improve pollutant detection accuracy, and ultimately contribute to the development of proactive and adaptive measures for mitigating the impact of air pollution on both human health and the environment.

Air Pollution

Air pollution encompasses a wide array of pollutants, each with its own sources and impacts on human health and the environment. Particulate matter (PM), consisting of tiny particles suspended in the air, is a significant component of air pollution. These particles vary in size, with PM₁₀ (particles with a diameter of 10 micrometers or less) and PM_{2.5} (particles with a diameter of 2.5 micrometers or less) being of particular concern due to their ability to penetrate deep into the respiratory system and cause respiratory issues, cardiovascular diseases, and other health problems. Common sources of particulate matter include vehicle

emissions, industrial processes, and natural phenomena like wildfires and dust storms (Cohen et al., 2015).

Another prominent air pollutant is nitrogen dioxide (NO₂), a gas primarily emitted from combustion processes in vehicles, power plants, and industrial facilities. NO₂ can irritate the respiratory system, exacerbate asthma symptoms, and contribute to the formation of ground-level ozone and fine particulate matter. Additionally, sulfur dioxide (SO₂), mainly released from burning fossil fuels like coal and oil, is a significant air pollutant known for its adverse effects on respiratory health and its role in the formation of acid rain. Efforts to control SO₂ emissions have led to improvements in air quality, particularly in areas with strict regulations on sulfur content in fuels.

Volatile organic compounds (VOCs) constitute another category of air pollutants, encompassing a variety of organic chemicals that can evaporate into the air. Sources of VOCs include vehicle emissions, industrial processes, and household products like paints and solvents. VOCs can react with other pollutants in the atmosphere to form ground-level ozone, a major component of smog, which can cause respiratory problems and damage to vegetation. Addressing VOC emissions is crucial for improving air quality and reducing the impacts of air pollution on human health and ecosystems (Fuhrer et al., 2016).

Traditional methods of air pollution control primarily focus on reducing emissions of pollutants at the source and mitigating their dispersion into the atmosphere. One common approach is the implementation of emission standards and regulations, which set limits on the amount of pollutants that industrial facilities, vehicles, and other sources can release into the air. These standards often involve the use of pollution control technologies such as catalytic converters in vehicles, scrubbers in industrial smokestacks, and filters in power plants to remove harmful pollutants from exhaust gases before they are released into the atmosphere. Additionally, the promotion of cleaner production practices and the adoption of cleaner fuels can help minimize emissions and improve air quality.

Another traditional method of air pollution control is the implementation of land-use planning and zoning regulations to reduce exposure to pollutants. This approach involves strategically locating industrial facilities, highways, and other sources of pollution away from residential areas, schools, and other sensitive receptors. By controlling the siting of polluting sources and establishing buffer zones between industrial and residential areas, this method aims to minimize the health risks associated with exposure to air pollution. Moreover, urban planning strategies such as increasing green spaces, promoting public transportation, and encouraging energy-efficient building designs can further contribute to reducing air pollution levels and enhancing overall environmental quality in urban areas (Hakami et al., 2004).

Deep Learning

Deep learning, a subset of machine learning, has emerged as a powerful tool for solving complex problems by automatically learning representations from large amounts of data. At its core, deep learning utilizes artificial neural networks inspired by the structure and function of the human brain. These neural networks consist of interconnected layers of artificial neurons, each layer processing and transforming input data to produce increasingly abstract representations. Through a process called backpropagation, neural networks iteratively adjust their parameters to minimize errors and improve their performance on a given task.

The applications of deep learning span a wide range of domains, including computer vision, natural language processing, and speech recognition. In computer vision, deep learning algorithms have achieved remarkable success in tasks such as object detection, image classification, and facial recognition. For instance, convolutional neural networks (CNNs), a type of deep learning architecture designed to process visual data, have revolutionized image recognition systems and enabled breakthroughs in fields like autonomous driving and medical imaging.

In natural language processing (NLP), deep learning models have significantly advanced the capabilities of machines to understand and generate human language. Recurrent neural

networks (RNNs) and transformers, two prominent architectures in NLP, have been instrumental in tasks such as machine translation, sentiment analysis, and text summarization. These models leverage the hierarchical structure of language to capture semantic relationships and context, enabling more accurate and contextually relevant language processing.

Moreover, deep learning has found applications in healthcare, where it has demonstrated promise in areas such as disease diagnosis, drug discovery, and personalized treatment planning. By analyzing large medical datasets, deep learning algorithms can assist healthcare professionals in detecting diseases from medical images, predicting patient outcomes, and identifying potential drug candidates. The ability of deep learning models to extract meaningful patterns from complex biological data holds great potential for improving patient care and advancing medical research (Liao et al., 2020).

Architectures for Temporal Predictions

Recurrent Neural Networks: The RNNs are variants of feed forward neural networks (FNN). FNNs enable signals to travel only one way from input to output. They are straightforward networks without loops that associate inputs with outputs. RNNs introduce self-connection of neurons cyclic structure into the network, which are based on FNN. Thus, input data can be memorized, and sequences of data can influence network outputs through self-connected neurons. Taking advantage of their memory characteristics, RNNs outperform FNNs in many applications. However, RNNs may fail to capture long time dependencies in input data, and it may face the problems of vanishing and exploding gradients when the time of training is too long.

Long Short-Term Memory: LSTM networks (LSTM) are enhanced RNNs. They introduce memory blocks to overcome the vanishing gradient problem. The memory blocks consist of three types nonlinear multiplicative gates: the input gate, output gate, and forget gate. The multiplicative gates control the memory block operation and determine whether the input information need to be remembered. The input gate controls the flow of cell activation from input into a memory cell, while output gate controls the flow of output from a memory cell into other nodes. LSTM networks have the advantage to train long time sequences and perform better than traditional RNN in many applications.

Gated Recurrent Unit: The GRU networks are simplified versions of the LSTM networks. They only consist of update and reset gates but can still balance the data flows inside the unit. The update gate replaces the input and forgets gates in LSTM, which determines whether information needs to be remembered. The advantage of using GRU compared with LSTM is that GRU have fewer parameters and thus less computational loads for training. Nevertheless the GRU networks have shown similar performances on music and speech signals as LSTM or even better performance on smaller datasets (Cho et al, 2014).

Architectures for Spatial Feature Extractions

Convolutional Neural Networks: The CNNs are deep feed forward networks which consist of a series of convolutional layers. They are capable of analyzing multiscale shift invariant features of data. Subsampling operations are performed between two successive convolutional layers. Two commonly used subsampling operations are max pooling and mean pooling. Pooling layers can be replaced by convolutional layers, as simplifies the network structure. Units in a convolutional layer are organized in feature maps, and each unit is connected to local weights in the feature maps of the previous layer through a filter bank. The sum of local weights is passed through an activation function which can take various forms such as Rectified Linear Units (ReLU) and Scaled Exponential Linear Units (SELU). The CNNs have produced outstanding results in processing multiple dimensional array data with spatial structure. They are widely used in speech recognition and image recognition, as motivates researchers to estimate environmental exposures through digital images using CNNs. CNN can effectively extract the spatial features of pollutants.

Stacked Autoencoder An autoencoder is a neural network that attempts to reconstruct its inputs. This is done by minimizing the discrepancy between inputs and network outputs. It has the ability to extract the features in reduced spaces through the reconstructions of the inputs. Stacked autoencoder (SAE) is a deep model formed by stacking successive layers of autoencoders. For a SAE with M hidden layer, the autoencoders perform unsupervised pre-training from hidden layer k ($k < M$) to hidden layer $(k + 1)$. Each hidden layer is a higher-level abstraction of the previous layer, and the final hidden layer contains high-level feature which is more effective for prediction.

Deep Belief Networks The DBNs are formed by stacking multiple energy-based Restricted Boltzmann Machine (RBM). They have achieved excellent results in feature recognition and classification as well as prediction problems. An RBM consists of a visible layer and a hidden layer, where the hidden layer of the prior RBM is the visible layer of the next RBM. The RBMs performs unsupervised pre-training layer by layer from bottom to top to initialize the network parameters of each layer. After the pretraining process, a softmax classifier is set in the last layer of DBN to classify the features. Finally the entire network is tuned supervised tuned through the labeled network using back propagation algorithm. Furthermore, a tenfold cross validation technique is commonly applied to evaluate the model performance and test the model overfitting. Figure 1 illustrates deep networks architecture for air quality forecasts (Rodriguez et al., 2010).

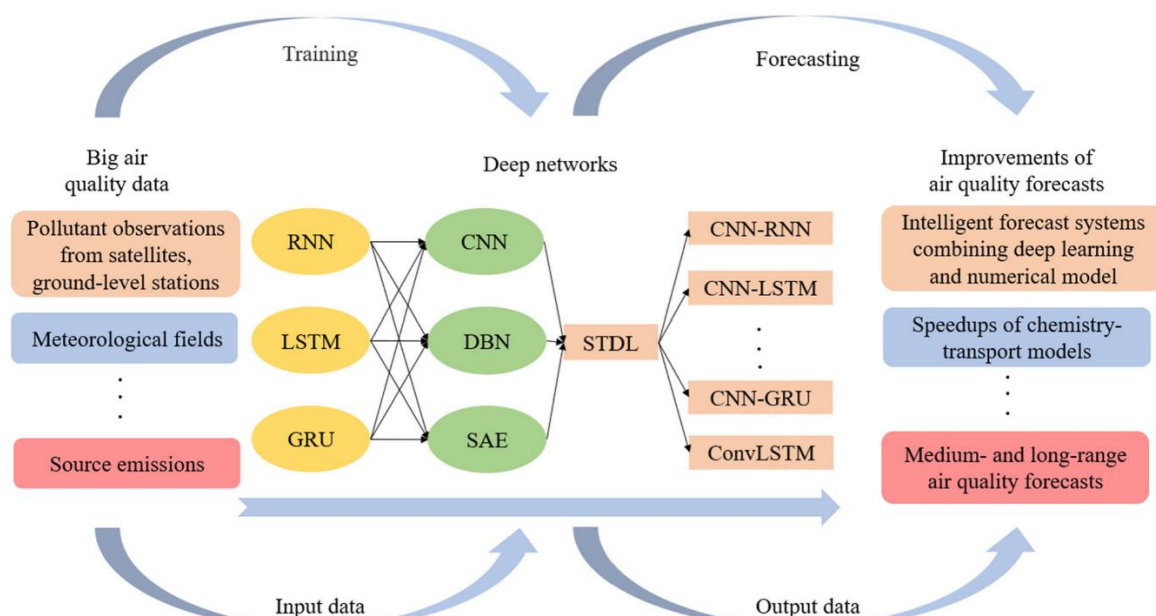


Figure 1. Deep Network Architectures for Air Quality Forecasts

Deep Learning Applications in Air Pollution Control

Deep learning holds great promise in revolutionizing air pollution control by offering advanced techniques for pollutant detection, monitoring, and mitigation. By leveraging large datasets of atmospheric measurements and pollutant emissions, deep learning models, such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs), can effectively analyze complex spatiotemporal patterns in air quality dynamics. These models have the capability to discern subtle relationships between pollutant sources, meteorological conditions, and pollutant concentrations, thereby enabling more accurate and timely detection of air pollution hotspots and forecasting of pollution events. Furthermore, deep learning algorithms can aid in the optimization of pollution control strategies by providing insights into the effectiveness of various intervention measures and assisting in the development of adaptive control systems. Overall, the integration of deep learning into air pollution control efforts holds the potential to enhance our understanding of air quality

dynamics, improve the efficiency of pollution mitigation measures, and ultimately contribute to the protection of public health and the environment.

The study by Xing et al. (Xing et al., 2020) presents a groundbreaking approach to predicting air quality response to emission changes, leveraging the power of deep learning coupled with chemical indicators. This stands to revolutionize the field of air pollution control, where accurately modeling the intricate, non-linear relationship between emissions and air quality has been a persistent challenge.

Traditionally, response functions are employed to quantify how air quality metrics like $PM_{2.5}$ or O_3 respond to changes in precursor emissions like NO_x , SO_2 , or VOCs. Deriving these functions necessitates numerous model simulations encompassing various emission scenarios, often exceeding 20 runs. This approach, while computationally expensive, suffers from limited accuracy due to its inability to fully capture the non-linearities inherent in atmospheric chemistry and physics.

Xing et al. propose a novel two-step methodology that overcomes these limitations. First, they utilize a comprehensive atmospheric model to simulate chemical indicators that encapsulate the key processes governing pollutant formation under diverse emission conditions. These indicators serve as informative features for a subsequent DL model, specifically a Long Short-Term Memory (LSTM) network. The LSTM network then excels at learning the complex relationships between these indicators and the corresponding air quality responses, effectively capturing the non-linearities that traditional methods struggle with.

The brilliance of this approach lies in its efficiency. By leveraging pre-computed chemical indicators from only two model simulations, the authors achieve remarkable accuracy in predicting air quality response, outperforming traditional methods that require significantly more simulations. This significant reduction in computational demands opens doors for wider applicability and real-time implementation within air quality management frameworks.

However, further research is warranted to explore the generalizability of this method across diverse geographical regions and atmospheric conditions. Additionally, the interpretability of DL models, particularly LSTM networks, remains an ongoing area of research. Elucidating the internal workings of these models would provide valuable insights into the learned relationships between emissions and air quality, fostering trust and wider adoption in the scientific community. Figure 2 illustrates mathematical model used in this research.

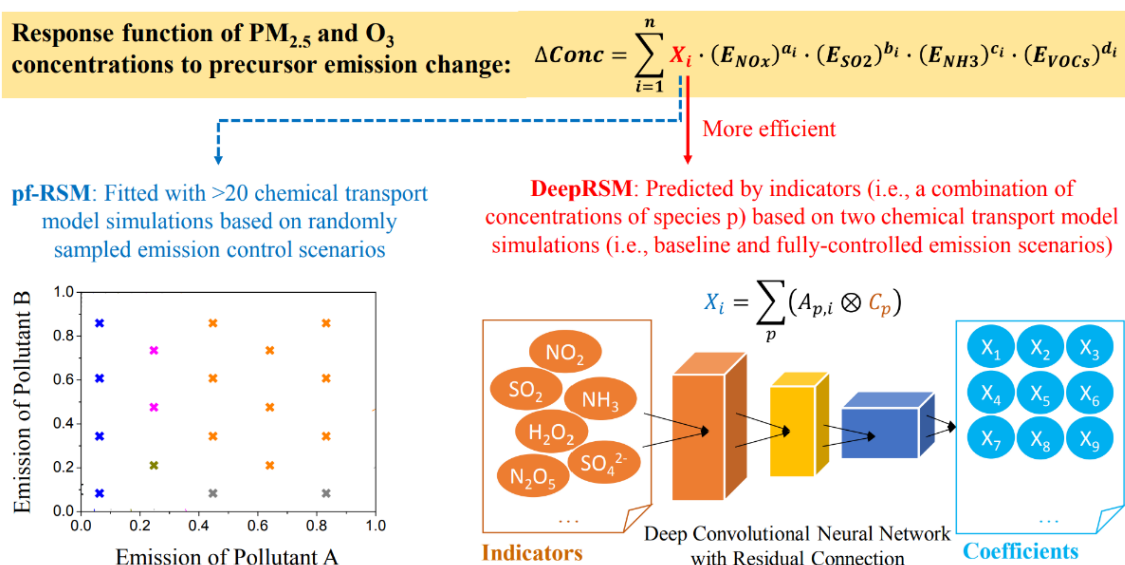


Figure 2. Illustration of RSM model of LSTM

Predicting $PM_{2.5}$, a harmful air pollutant linked to respiratory and cardiovascular illnesses, remains crucial for effective air quality management. However, accurately capturing the complex interplay between emissions, meteorology, and $PM_{2.5}$ concentrations poses a significant challenge. An study done by Jeya and Sankari (Jeya and Sankari, 2020) introduces a novel deep learning model, leveraging Bidirectional Long Short-Term Memory (Bi-LSTM) networks, to tackle this intricate prediction task.

The proposed Bi-LSTM model excels at learning long-term dependencies within data sequences. It is trained on historical $PM_{2.5}$ measurements, alongside relevant meteorological factors and other variables, effectively capturing the dynamic relationships governing air quality. Unlike traditional methods, the Bi-LSTM architecture analyzes data in both forward and backward directions, enabling it to identify subtle temporal patterns and long-range dependencies that might influence $PM_{2.5}$ concentrations.

The authors rigorously evaluated their model's performance on a comprehensive dataset of $PM_{2.5}$ concentrations from Beijing, China. Notably, the Bi-LSTM model outperformed existing prediction models, including linear regression, support vector regression, and random forest, achieving a remarkably low mean absolute error (MAE) of $5.8 \mu g/m^3$. This superior accuracy demonstrates the Bi-LSTM's ability to handle the inherent non-linearities and complex interactions within air quality data.

Furthermore, the Bi-LSTM model boasts computational efficiency. By employing a stacked architecture with multiple Bi-LSTM layers, it effectively extracts intricate patterns from the data while maintaining a faster training process compared to traditional methods. This efficiency paves the way for real-time air quality monitoring and forecasting, empowering policymakers with timely insights to implement effective emission control strategies and safeguard public health. this study presents a compelling deep learning approach for $PM_{2.5}$ prediction. The Bi-LSTM model, with its exceptional accuracy, efficiency, and ability to capture long-term dependencies, holds immense potential to revolutionize air quality management. Further research exploring its generalizability across diverse regions and atmospheric conditions is warranted, but this innovative approach marks a significant step towards cleaner air for all. Figure 3 presents $PM_{2.5}$ and wind speed datasets from 2014 to 2020 used in this study and figure 4 illustrates Bi-LSTM approach.




Figure 3. $PM_{2.5}$ and Wind Speed datasets form 2014 to 202

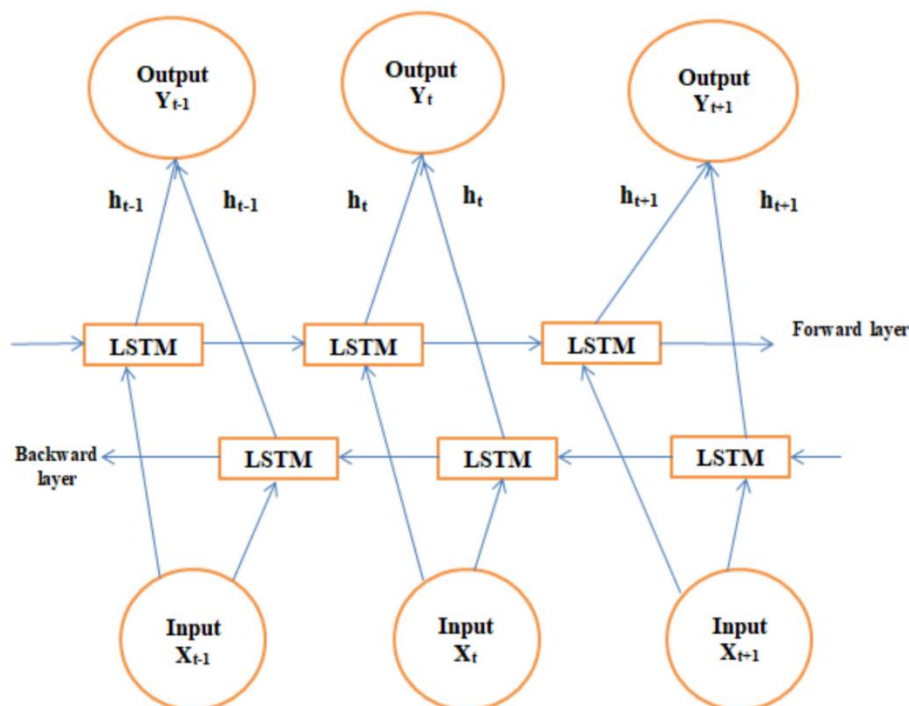


Figure 4. Bi-LSTM approach architecture

Air pollution has significant negative effects on human health, with the World Health Organization attributing 3.8 million deaths each year to air pollution globally. Vehicle emissions are a major source of air pollutants in many developed cities, contributing to a significant portion of nitrogen oxide and particulate-matter pollution. Traditionally, estimating the contribution of traffic volumes to air pollution is done through complex models that require expertise and are computationally expensive. However, a new approach suggested by Hahnel et al. (Hahnel et al., 2020) involves using deep learning and techniques from partial differential equations to develop rapid solvers that can scale to any domain size. This approach integrates deep learning with PDE-based domain decomposition, creating a single unified model for the entire region by merging the data learned from independent, neighboring meshes. The DL model is trained on data generated by a PDE model for air pollution, which serves as a computationally lightweight representation. Consistency constraints are used to ensure physically meaningful solutions even at the boundaries of the domains. The approach is tested in a numerical study on a pollution-forecasting problem, demonstrating its effectiveness compared to the PDE model and sensor data.

The researchers have developed a method to train a surrogate model for a partial differential equation (PDE) by using domain decomposition. Their approach involves training a deep-learning model for each subdomain while ensuring consistency across neighboring domains. By enabling communication between subdomains through constraints, predictions for one subdomain can benefit from information outside of it, leading to improved accuracy and generalization compared to models trained on individual subdomains. The study considers an index-set of meshes and mesh points, with the output of simulations on the mesh consisting of values at each point. A sub-set of points called receptors is of particular interest, while the rest are hidden points. Boundaries between meshes are also considered, with importance assigned to each boundary. Simulations are performed with inputs and outputs, and recurrent analysis is possible. Figure 5 shows recurrent neural network suggested by this study.

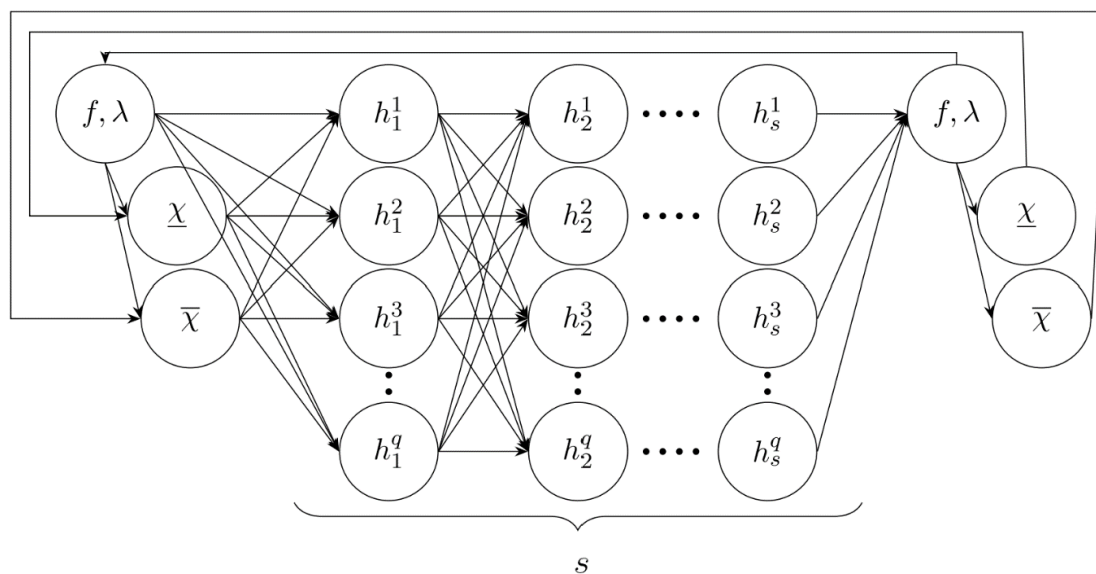


Figure 5. Recurrent Neural Network Suggested by Hahnel et al.

To illustrate this architecture they trained deep learning for a city scale pollution monitoring. The test case was based in the city of Dublin. The goal of this study is to estimate the levels of air pollution caused by traffic, specifically NO_2 and PM_{10} , at different locations throughout the city. These pollutants are closely related to major health concerns such as lung and heart diseases, and they are mainly generated by vehicle emissions. In addition, ozone (O_3) is produced through complex reactions involving organic compounds and nitrogen oxides (NO_x). To make these estimations, the researchers used a prediction framework that took into account traffic volumes, weather data, and an air-pollution dispersion model. The model treated road links as line sources and used Gaussian-plume models to describe the temporal and spatial evolution of vehicle emissions near roadways. The Caline, Hiway, and Aermot models were used as examples of Gaussian plume models. While there are more sophisticated numerical models available, the researchers chose to use the Gaussian-plume model, specifically the Caline-4 implementation, due to its wide use and simplicity. They acknowledge that more complex models exist but were not within the scope of this study. The researchers defined each pollutant's mass at a specific location and time. The concentration profiles were given in the downwind directions, considering dispersion factors. The law of conservation of mass was applied, and the advection-diffusion equation was used to derive the Gaussian Plume solution. The solution accounted for wind velocity, space-diffusion coefficients, and emission rate. Figure 6 presents map of Dublin and two-domain model applied on it.




Figure 6. Map of Dublin and Deep Learning Method applied on it

The behavior of the solution at neighboring interfaces is demonstrated, showing an iterative relaxation towards a reconciliation of both solutions. The mean absolute error of the deep-learning model stabilizes after eight iterations. The average difference of predicted concentration values across the boundary converges, reducing the discontinuity by about 25% to 30% when consistency constraints are imposed. The computational expense of training the model is significant, taking about 120 CPU-hours for 20 iterations on a commodity-compute resource. However, the computational cost of deploying the trained model for prediction is negligible. The trained RNN model shows a speed-up factor of more than two orders of magnitude compared to the Caline model for the study period. The predictive skill of the DL model is evaluated by comparing its estimates with Caline estimates at defined locations and across the entire city. The DL model closely captures the general trends of the Caline estimates with differences on average less than 3 $\mu\text{g}/\text{cm}^3$ for NO_2 and 15 $\mu\text{g}/\text{cm}^3$ for PM_{10} , with no evident biases. The DL model also captures areas of high pollution contributions well, with peak values similar to Caline estimates. However, it predicts a smoother distribution of pollution compared to Caline, resulting in a significant mismatch in regions with low traffic-generated pollution. The deep-learning computed values have a mean absolute error of 1.7 $\mu\text{g}/\text{cm}^3$ with a standard deviation of 2.1 $\mu\text{g}/\text{cm}^3$.

Conclusion

In conclusion, this review paper has provided a comprehensive overview of the applications of deep learning in air pollution control. Through an exploration of deep learning methodologies such as convolutional and recurrent neural networks, as well as their integration into air quality monitoring systems, it is evident that deep learning offers transformative opportunities for addressing the complexities of air pollution detection, monitoring, and mitigation. By leveraging large datasets and advanced computational techniques, deep learning models have demonstrated unparalleled capabilities in discerning intricate patterns in air quality dynamics, enabling more accurate pollutant detection, forecasting, and control strategies. However, challenges remain, including issues related to data quality, model interpretability, and computational requirements, highlighting the need for interdisciplinary collaborations and ongoing research efforts. Nevertheless, the integration of deep learning into air pollution control efforts holds tremendous promise for improving public health outcomes, enhancing environmental quality, and advancing sustainable strategies for mitigating the impacts of air pollution on society. As it continues to harness the potential of deep learning technologies, it is imperative to prioritize research and innovation in this area to address the pressing challenges of air pollution and pave the way for a cleaner and healthier future.

References

- Cho, K. (2014). On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Cohen, A. J., Brauer, M., Burnett, R., Anderson, H. R., Frostad, J., Estep, K., ... & Forouzanfar, M. H. (2017). 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*, 389(10082), 1907-1918.
- Csáji, B. C. (2001). Approximation with artificial neural networks. *Faculty of Sciences, Eötvös Loránd University, Hungary*, 24(48), 7.
- Fuhrer, J., Val Martin, M., Mills, G., Heald, C. L., Harmens, H., Hayes, F., ... & Ashmore, M. R. (2016). Current and future ozone risks to global terrestrial biodiversity and ecosystem processes. *Ecology and evolution*, 6(24), 8785-8799.
- Hähnel, P., Mareček, J., Monteil, J., & O'Donncha, F. (2020). Using deep learning to extend the range of air pollution monitoring and forecasting. *Journal of Computational Physics*, 408, 109278.
- Hakami, A., Odman, M. T., & Russell, A. G. (2004). Nonlinearity in atmospheric response: A direct sensitivity analysis approach. *Journal of Geophysical Research: Atmospheres*, 109(D15).

- Jeya, S., & Sankari, L. (2020, May). Air pollution prediction by deep learning model. In *2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS)* (pp. 736-741). IEEE.
- Liao, Q., Zhu, M., Wu, L., Pan, X., Tang, X., & Wang, Z. (2020). Deep learning for air quality forecasts: a review. *Current Pollution Reports*, 6, 399-409.
- Rodriguez, J. D., Perez, A., & Lozano, J. A. (2009). Sensitivity analysis of k-fold cross validation in prediction error estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(3), 569-575.
- Xing, J., Zheng, S., Ding, D., Kelly, J. T., Wang, S., Li, S., ... & Hao, J. (2020). Deep learning for prediction of the air quality response to emission changes. *Environmental Science & Technology*, 54(14), 8589-8600.
- Xu, B., Wang, N., Chen, T., & Li, M. (2015). Empirical evaluation of rectified activations in convolutional network (2015). *arXiv preprint arXiv:1505.00853*.

