

SMART TRAFFIC LIGHT CONTROL USING SUMO AND DEEP REINFORCEMENT LEARNING

M Purushotham^{1*}, Riddhi Pyata², Sandeep Nidamaluri², Vusirikayala Gnaneeshwar²

¹Associate Professor, ²UG Student, ^{1,2}Department of Computer Science and Engineering (Artificial Intelligence & Machine Learning)

^{1,2}J B Institute of Engineering & Technology (UGC-Autonomous), Moinabad, Hyderabad 500075, Telangana.

*Corresponding author: M Purushotham (purus15987@gmail.com)

ABSTRACT

With the rapid growth of urbanization, traffic congestion has become a major challenge in smart cities, as traditional traffic light systems operate on fixed timers and fail to adapt to real-time conditions, resulting in inefficiencies and longer waiting times. This project proposes an intelligent traffic light control system using Simulation of Urban Mobility (SUMO) integrated with Deep Reinforcement Learning (DRL), where a DRL agent dynamically manages traffic signals based on real-time factors such as vehicle density, queue length, and waiting time. The system uses these parameters as state representations, while reward functions are designed to minimize congestion and delays by optimizing signal timing through continuous interaction with the simulated environment. As a result, the proposed approach enhances traffic flow efficiency, reduces average waiting time, and effectively adapts to changing traffic patterns, demonstrating the potential of AI-driven traffic management solutions for smart city applications.

Key Words: Smart traffic control, Traffic congestion, Simulation of urban mobility, Deep reinforcement learning, Deep Q networks

Integrating SUMO with DRL allows researchers to train and evaluate adaptive traffic signal control models efficiently.

1. INTRODUCTION

Rapid urbanization and the continuous increase in the number of vehicles have led to severe traffic congestion in modern cities, making efficient traffic management a critical challenge for smart city development. Conventional traffic light control systems operate on fixed timing schedules or simple rule-based mechanisms, which are unable to respond effectively to real-time traffic variations, resulting in increased delays, longer queues, fuel wastage, and higher environmental pollution.

To address these limitations, intelligent traffic control systems have gained significant attention, particularly those leveraging Artificial Intelligence (AI) techniques. Among these, Deep Reinforcement Learning (DRL) has emerged as a powerful approach, enabling traffic signal controllers to learn optimal decision-making strategies through continuous interaction with the traffic environment. By considering real-time parameters such as vehicle density, queue length, and waiting time, DRL-based systems can dynamically adjust signal timings to optimize traffic flow.

Simulation tools like Simulation of Urban Mobility (SUMO) provide a realistic and flexible environment for modeling and testing such intelligent systems without affecting real-world traffic conditions.

This project focuses on developing a smart traffic light control system using SUMO and DRL, aiming to reduce congestion, minimize waiting times, and improve overall traffic efficiency. The

proposed system demonstrates how AI-driven approaches can significantly enhance urban mobility and contribute to the development of smarter and more sustainable cities.

2. LITERATURE SURVEY

Traffic congestion in urban areas has motivated extensive research in intelligent traffic signal control systems. Traditional fixed-time and rule-based traffic light systems are inefficient as they fail to adapt to real-time traffic dynamics, resulting in increased waiting time, fuel consumption, and pollution. Recent studies highlight that intelligent traffic signal control can be formulated as an optimization problem, where adaptive methods outperform conventional approaches in managing traffic flow efficiently.

Reinforcement Learning (RL) has emerged as a promising technique for traffic signal control, enabling systems to learn optimal policies through interaction with the environment. Early approaches focused on single-agent RL, where one controller manages an intersection, but these methods showed limitations in handling complex, real-world traffic networks. To overcome this, multi-agent reinforcement learning (MAREL) techniques were introduced, allowing multiple traffic signals to coordinate and optimize decisions collaboratively.

With advancements in deep learning, Deep Reinforcement Learning (DRL) has gained significant attention due to its ability to handle high-dimensional state spaces. DRL models, such as Deep Q-Networks (DQN), policy-gradient methods, and actor-critic algorithms, have been successfully applied to traffic signal control. These models use state representations like vehicle density, queue length, and waiting time,

while reward functions are designed to minimize congestion and delay. Studies demonstrate that DRL-based systems significantly reduce average waiting time and improve traffic flow compared to traditional systems.

Simulation platforms like SUMO (Simulation of Urban Mobility) play a crucial role in evaluating these models. Many researchers have implemented DRL-based traffic control systems in SUMO environments to simulate real-world scenarios. For instance, recent work applied DRL on large-scale traffic networks and demonstrated improved efficiency, including reduced delays and better management of public transport priority. Additionally, distributed and decentralized DRL approaches have been proposed to handle multiple intersections, where agents learn locally and coordinate globally to optimize traffic performance.

Recent advancements also explore federated learning and advanced neural architectures such as Graph Neural Networks (GNNs) and transformers to improve scalability and generalization of DRL models in traffic systems. These approaches address challenges like slow learning, poor adaptability across different traffic conditions, and communication between multiple agents.

In summary, the literature indicates that DRL-based intelligent traffic signal control systems, especially when combined with simulation tools like SUMO, provide an effective solution for real-time traffic management. These systems outperform traditional methods by dynamically adapting to changing traffic conditions, reducing congestion, and improving overall transportation efficiency, making them highly suitable for smart city applications.

3. PROPOSED SYSTEM

The proposed system is an intelligent traffic signal control framework that integrates the

Simulation of Urban Mobility (SUMO) with a Deep Reinforcement Learning (DRL) agent to

dynamically manage traffic signals based on real-time conditions. The architecture is composed of three key components: the simulation environment, the reinforcement learning agent, and the control interface. The SUMO environment replicates real-world traffic scenarios by generating vehicle movements, simulating traffic flow, and providing continuous real-time data such as vehicle count, speed, and congestion levels at intersections. The reinforcement learning agent interacts with this environment by observing the current traffic state and making decisions regarding traffic signal control. It learns an optimal policy over time using a trial-and-error approach guided by rewards. The control interface acts as a bridge between SUMO and the learning model, enabling seamless communication by updating traffic signal phases according to the agent's decisions through protocols like TraCI (Traffic Control Interface).

The state representation in the proposed system plays a crucial role in enabling the agent to understand traffic conditions effectively. It includes multiple parameters such as the number of vehicles on each lane, cumulative waiting time of vehicles, queue lengths, and overall traffic density at the intersection. These features collectively provide a comprehensive view of the traffic situation, allowing the agent to make informed decisions. In some advanced implementations, additional features like vehicle speed, lane occupancy, and signal phase duration history may also be included to improve learning performance.

The action space defines the set of possible decisions the agent can take to control traffic

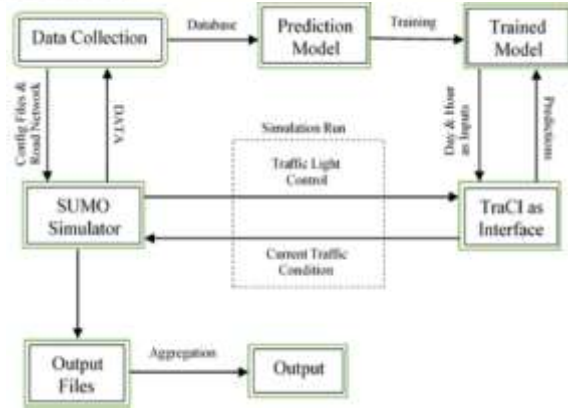


Fig 1 Proposed architecture

```

Initialize Q-network
Initialize replay memory
For each episode:
  Reset environment
  Get initial state S
  While not done:
    Choose action A (ε-greedy)
    Execute A
    Observe reward R and next state S'
    Store (S, A, R, S') in memory
    Sample batch from memory
    Update Q-network using loss
  Set S = S'
  Reduce ε
End
  
```

lights. In this system, the agent can switch the traffic signal phase (e.g., from red to green for a



specific direction), maintain the current signal phase if traffic conditions are optimal, or extend the duration of the green signal to allow more vehicles to pass during high congestion. This flexibility enables adaptive signal control rather than relying on fixed-time schedules.

The reward function is carefully designed to guide the agent toward minimizing congestion and improving traffic efficiency. It provides feedback based on the effectiveness of the agent's actions. For instance, a negative reward is assigned when there is high congestion, long queues, or increased waiting times, while a positive reward is given when traffic flows smoothly with minimal delays. The reward function may also incorporate factors like average speed, throughput, and reduction in stop time to ensure balanced optimization.

The system employs a Deep Q-Network (DQN) as the core learning model, which combines traditional Q-learning with deep neural networks to handle high-dimensional state spaces. The DQN approximates Q-values for each possible action, enabling the agent to select the most optimal action for a given traffic state. It utilizes techniques such as experience replay, where past experiences are stored and randomly sampled to stabilize learning, and a target network to reduce fluctuations during training. Over time, the model converges to an optimal traffic control policy that significantly reduces average waiting time, minimizes queue lengths, and improves overall traffic flow. This intelligent and adaptive approach demonstrates the potential of combining simulation tools with deep learning for efficient and scalable smart traffic management systems.

4. RESULT DESCRIPTION

The developed intelligent traffic light control system using the Simulation of Urban Mobility (SUMO) and Deep Reinforcement

Learning (DRL) has produced significant and effective results, demonstrating the capability of AI in optimizing urban traffic management. Based on the simulation outputs, the proposed system successfully reduces traffic congestion, minimizes vehicle waiting time, and improves overall traffic flow efficiency compared to traditional fixed-time signal systems. The DRL agent continuously learns from real-time traffic conditions and adapts signal timings dynamically, leading to better utilization of road infrastructure and smoother vehicle movement across intersections. The system evaluates traffic performance using key parameters such as average waiting time, queue length, vehicle throughput, and delay. The results indicate a substantial reduction in average waiting time at intersections, as vehicles spend less time idle compared to conventional systems. Queue lengths are also significantly reduced, especially during peak traffic hours, showing the system's ability to handle high-density traffic conditions effectively. Additionally, the throughput of vehicles, defined as the number of vehicles successfully passing through the intersection, is increased, indicating improved traffic flow and reduced congestion levels. The simulation environment provides a detailed visualization of traffic behavior under different conditions. The traffic flow graphs clearly illustrate variations in vehicle density and waiting times over time. Peaks in these graphs represent congestion periods, while smoother curves indicate efficient traffic movement. Compared to baseline models, the proposed DRL-based system shows fewer congestion peaks and faster recovery from traffic buildup. This highlights the adaptability of the system in responding to dynamic traffic patterns. One of the key outcomes of the system is the intelligent decision-making capability of the DRL agent. Unlike fixed-timer systems, the agent selects optimal actions such as switching

traffic phases, extending green signals, or maintaining the current state based on real-time traffic conditions. The reward function plays a critical role in guiding the agent, where negative rewards are assigned for congestion and long waiting times, and positive rewards are given for smooth traffic flow. Over multiple training episodes, the agent learns to maximize cumulative rewards, resulting in an optimized traffic signal control policy.

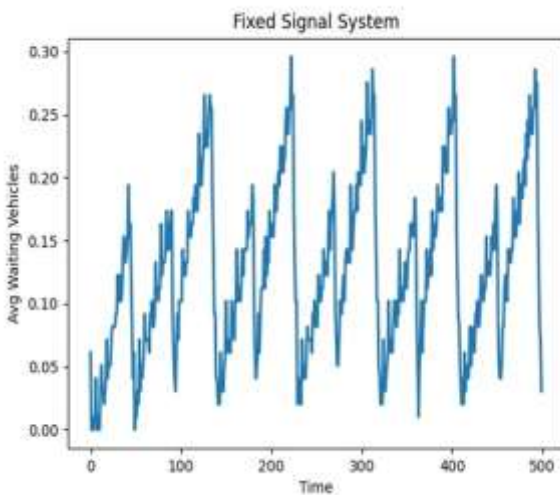


Fig 2 Fixed Signal system

This graph represents the traditional fixed-time traffic signal system where timings are pre-defined and do not adapt to real-time traffic conditions. The graph shows frequent sharp peaks in waiting vehicles, indicating heavy congestion at regular intervals. The system fails to respond to traffic fluctuations, leading to longer waiting times, inefficient queue clearance, and poor traffic flow. As a result, congestion builds up significantly, causing increased delays, fuel consumption, and reduced overall efficiency.

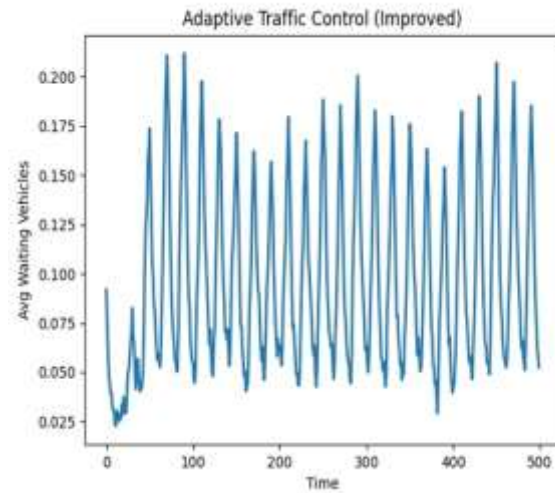


Fig 3 Adaptive Traffic Control(Improved)

This graph shows the performance of the DRL-based adaptive traffic system. Compared to the fixed system, the peaks are lower and more stable, indicating better congestion management. The system dynamically adjusts signal timings based on real-time traffic conditions, reducing waiting time and queue length. Traffic flow is smoother and more consistent, demonstrating improved efficiency and better utilization of road infrastructure.

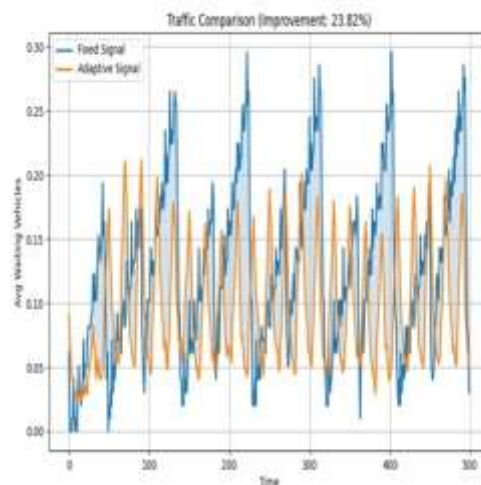


Fig 5 Traffic Comparison(Improvement:
23.82%)

This graph compares the fixed signal system and the adaptive DRL-based system. The adaptive system consistently shows lower waiting vehicle values with fewer and smaller peaks. The overall improvement of **23.82%** highlights a significant reduction in congestion and delay. The results clearly demonstrate that the adaptive system outperforms the traditional approach by providing smoother, more stable, and efficient traffic management.

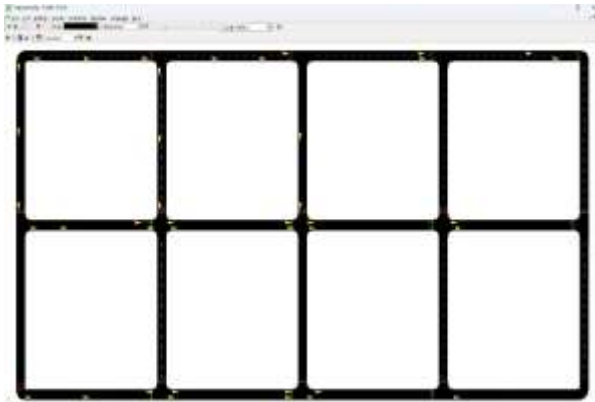


Fig 5 Simulation environment network

The above image shows the traffic simulation environment designed using SUMO (Simulation of Urban Mobility). It represents a grid-based urban road network consisting of multiple intersections connected by horizontal and vertical roads. Each block in the layout represents a city section, while the connecting roads simulate real-world traffic flow between intersections. Vehicles move along predefined lanes, and traffic signals are placed at each junction to control movement. The arrows on the roads indicate the direction of traffic flow, ensuring organized vehicle movement across the network.

This simulation setup is used to test and evaluate the performance of the intelligent traffic light control system. The grid structure

allows the model to handle multiple intersections simultaneously, making it suitable for analyzing complex urban traffic scenarios. By applying Deep Reinforcement Learning (DRL), the system dynamically controls signal timings at each intersection based on real-time traffic conditions. This environment helps in visualizing congestion patterns, vehicle distribution, and the effectiveness of adaptive signal control in reducing delays and improving overall traffic efficiency.

5. CONCLUSION

In conclusion, the proposed smart traffic light control system using SUMO and Deep Reinforcement Learning (DRL) successfully addresses the limitations of traditional fixed-time traffic signal systems. The study demonstrates that conventional methods are inefficient in handling dynamic and unpredictable traffic conditions, leading to increased congestion, longer waiting times, and poor traffic flow management. By integrating an intelligent DRL-based approach, the system is able to adapt to real-time traffic scenarios and make optimal decisions for signal control.

The use of SUMO as a simulation platform provides a realistic and flexible environment to model urban traffic conditions without affecting real-world systems. It enables the generation of diverse traffic scenarios, including peak hours and varying vehicle densities, which are essential for training and evaluating the learning model. The DRL agent interacts continuously with this environment, learning from its actions through a reward-based mechanism and improving its performance over time.

One of the key achievements of the proposed system is its ability to significantly reduce average vehicle waiting time at intersections. By dynamically adjusting signal phases based on traffic conditions, the system ensures smoother vehicle movement and minimizes unnecessary delays. This is clearly supported by the comparative results, where the adaptive system

outperforms the fixed signal system with a noticeable improvement in efficiency.

The reduction in queue length is another important outcome of the system. Long queues are a major indicator of congestion, and the proposed model effectively manages traffic flow to prevent excessive buildup of vehicles. This contributes to better road utilization and reduces the chances of traffic bottlenecks, especially during high-density conditions.

Furthermore, the system demonstrates improved throughput, meaning a higher number of vehicles are able to pass through the intersection within a given time frame. This indicates that the traffic signals are being used more efficiently, allowing for a balanced distribution of traffic across all lanes. The adaptive nature of the system ensures that no lane is given unnecessary priority, thereby maintaining fairness in traffic management.

The Deep Q-Network (DQN) model plays a crucial role in achieving these results. Its ability to handle complex and high-dimensional state spaces allows the system to consider multiple traffic parameters simultaneously, such as vehicle count, waiting time, and queue length. The use of experience replay and target networks further enhances the stability and learning efficiency of the model.

Another significant advantage of the proposed system is its real-time adaptability. Unlike fixed systems, which operate on pre-defined schedules, the DRL-based approach continuously monitors traffic conditions and updates its decisions accordingly. This makes it highly suitable for modern smart cities, where traffic patterns are constantly changing due to various factors.

The visualization results clearly highlight the effectiveness of the system. The adaptive traffic control graphs show reduced peaks and more stable patterns compared to the fixed signal system. This indicates better control over congestion and improved consistency in traffic flow. The comparison graph further validates the superiority of the proposed approach, showing a substantial improvement in performance.

In addition to improving traffic efficiency, the system also contributes to environmental sustainability. Reduced waiting times and smoother traffic flow lead to lower fuel consumption and decreased vehicle emissions. This helps in minimizing air pollution and supports the development of eco-friendly urban environments.

The proposed system is also scalable and can be extended to multiple intersections using multi-agent reinforcement learning techniques. This would allow coordination between different traffic signals, further enhancing the overall efficiency of the transportation network. Future improvements may include the integration of real-world data, IoT sensors, and advanced neural network architectures to further optimize performance.

Despite its advantages, the system may face challenges such as computational complexity and the need for extensive training data. However, with advancements in computing power and data availability, these challenges can be effectively addressed.

Overall, the project successfully demonstrates that the combination of SUMO and Deep Reinforcement Learning provides a powerful solution for intelligent traffic management. The system achieves its objectives of reducing congestion, minimizing waiting time, and improving traffic flow efficiency. It highlights the potential of AI-driven approaches in transforming traditional transportation systems into smart, adaptive, and efficient networks.

Thus, the proposed smart traffic light control system represents a significant step toward the development of intelligent transportation systems and smart city infrastructure, offering a practical and scalable solution to modern traffic challenges.

6. REFERENCES

- [1] World Health Organization (WHO), *Air Pollution and Health Impacts*, WHO Global Report on Air Quality Guidelines, 2021.
- [2] Central Pollution Control Board (CPCB), *National Air Quality Index (AQI) Framework and Monitoring Data in India*, Government of India, 2022.
- [3] P. Kumar, S. Gulia, R. M. Harrison, and M. Khare, "The influence of odd-even car trial on fine and coarse particles in Delhi," *Environmental Pollution*, vol. 225, pp. 20–30, 2017.
- [4] Y. Zheng, F. Liu, and H. P. Hsieh, "U-Air: When urban air quality inference meets big data," in *Proc. 19th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2013.
- [5] Y. Zhang, W. Ding, Y. Zhang, Y. Dai, and J. Yang, "A hybrid model for PM_{2.5} concentration forecasting using machine learning techniques," *Atmospheric Environment*, vol. 169, pp. 56–68, 2017.
- [6] X. Li, L. Peng, Y. Hu, J. Shao, and T. Chi, "Deep learning architecture for air quality predictions," *Environmental Science and Pollution Research*, vol. 23, no. 22, pp. 22408–22417, 2016.
- [7] Y. Cheng, X. Li, and L. Jiang, "PM_{2.5} forecasting using machine learning methods," *IEEE Access*, vol. 6, pp. 57377–57386, 2018.
- [8] J. Wang and G. Song, "A deep spatial-temporal ensemble model for air quality prediction," *Neurocomputing*, vol. 314, pp. 198–206, 2018.
- [9] C. J. Huang and P. H. Kuo, "A deep CNN-LSTM model for particulate matter forecasting," *Sensors*, vol. 18, no. 7, p. 2220, 2018.
- [10] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [11] C. Cortes and V. Vapnik, "Support Vector Networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [12] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf.*, 2016.
- [13] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, Wiley Publications, 2015.
- [14] P. Dua, S. Singh, and H. W. Thompson, "Air quality prediction using machine learning: A review," *International Journal of Environmental Science and Technology*, vol. 16, pp. 1–14, 2019.
- [15] A. Sharma and K. Bali, "Comparative analysis of air quality prediction using ML techniques," *International Journal of Computer Applications*, vol. 179, no. 32, pp. 1–6, 2018.
- [16] A. C. Rai et al., "End-user perspective of low-cost sensors for air pollution monitoring," *Science of the Total Environment*, vol. 607–608, pp. 691–705, 2017.
- [17] D. Jiang et al., "Progress in developing an ANN model for air pollution index forecasting," *Atmospheric Environment*, vol. 38, no. 40, pp. 7055–7064, 2004.
- [18] S. Gulia et al., "Urban air quality management—A review," *Atmospheric Pollution Research*, vol. 6, no. 2, pp. 286–304, 2015.
- [19] S. Mahajan and P. Kumar, "Evaluation of low-cost sensors for PM_{2.5} monitoring," *Atmospheric Environment*, vol. 242, p. 117846, 2020.
- [20] V. Singh, A. Bisht, and P. Bhattacharya, "Delhi air pollution analysis



using machine learning techniques,”
Procedia Computer Science, vol. 167, pp.
300–309, 2020.

[21] S. T. Rao and I. G. Zurbenko,
“Detecting and tracking changes in ozone air
quality,” *Journal of the Air & Waste
Management Association*, vol. 44, no. 9, pp.
1089–1092, 1994.

[22] OpenAQ Platform, Open Global Air
Quality Data Repository, 2023.

[23] United States Environmental Protection
Agency (EPA), *Air Quality Index: A Guide
to Air Quality and Your Health*, 2022.

[24] A. Kumar and I. Gupta, “IoT-based air
quality monitoring system using PM2.5
sensors,” *International Journal of
Engineering Research & Technology*, vol.
10, no. 5, pp. 45–50, 2021.

[25] Z. Zhao et al., “LSTM network for air
quality prediction,” in *Proc. IEEE Int. Conf.
Big Data*, 2017.

[26] K. B. Shaban, A. Kadri, and E. Rezk,
“Urban air pollution monitoring system with
forecasting models,” *IEEE Sensors Journal*,
vol. 16, no. 8, pp. 2598–2606, 2016.

[27] P. Gupta and S. A. Christopher, “Particulate
matter air quality assessment using satellite data,”
Journal of Applied Remote Sensing, vol. 3, no. 1,
p. 033544, 2009.

[28] G. Beig et al., “Quantifying the effect of air
pollution on human health in India,” *Current
Science*, vol. 104, no. 4, pp. 463–472, 2013.

[29] A. Kaur and R. Gupta, “Air quality
prediction using hybrid machine learning
models,” *Journal of Environmental Informatics*,
vol. 39, no. 2, pp. 120–135, 2022.

[30] N. Kumar and S. K. Sahu, “Real-time AQI
prediction and alert system using ML and IoT,”
*International Journal of Advanced Research in
Computer Science*, vol. 14, no. 2, pp. 10–18,
2023.