



Research Article

Trans-MARL: Decentralized Multi-Agent Transformer for Robust AQI Forecasting

Pratima Maurya

Computer Science and Engineering, KIPM College of Engineering & Technology
 GIDA, Gorakhpur, Uttar Pradesh, India

Corresponding Author: * Pratima Maurya

DOI: <https://doi.org/10.5281/zenodo.20746473>

Abstract

Forecasting air pollution is tough for smart cities. Populations keep growing, factories keep churning out emissions, and the weather just won't sit still. Traditional AQI forecasting models, like LSTM networks and other centralized systems, just don't keep up. They're too fragile, get tripped up by the complexity of cities, and often miss those sudden, nasty pollution spikes. This paper brings something new to the table: Trans-MARL.

It's a decentralized multi-agent transformer framework, designed from the ground up for city-wide AQI forecasting. Instead of one big,

fragile model Trans-MARL works as a network a kind of "society of sensors." Think of it this way: each agent keeps an eye on its own corner of the city they all trade info and together they piece together a clearer more up-to-date picture. The whole thing operates as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP).

Which means every agent only sees a slice of the big picture but by working together they fill in the blanks. Trans-MARL stands on three legs. First, there's a transformer-based perception layer that catches the time-based patterns.

Then there's a spatial communication layer, using graph-based embeddings, which is how the agents exchange local info and stay connected. Finally, the multi-agent reinforcement learning decision layer makes sure all the pieces move in the right direction, optimising each agent's actions for the good of the city. When put to the test, Trans-MARL doesn't just hold up; it thrives.

It shrugs off sensor failures, snaps into action when pollution patterns shift, and picks up anomalies that older models overlook. In short, this decentralised approach scales with the city. Down the line, it can hook up with satellite images and traffic patterns, pushing environmental monitoring to a smarter, stronger future.

Manuscript Information

- ISSN No: 2583-7397
- Received: 10-05-2026
- Accepted: 15-06-2026
- Published: 18-06-2026
- IJCRM:5(3); 2026: 969-978
- ©2026, All Rights Reserved
- Plagiarism Checked: Yes
- Peer Review Process: Yes

How to Cite this Article

Maurya P. Trans-MARL: Decentralized Multi-Agent Transformer for Robust AQI Forecasting. Int J Contemp Res Multidiscip. 2026;5(3):969-978.

Access this Article Online



www.multiarticlesjournal.com

KEYWORDS: Air Quality Index, Smart Cities, Multi-Agent Reinforcement Learning, Transformer, Networks, Decentralized AI, AQI Forecasting, Urban Intelligence, Environmental Monitoring.

1. INTRODUCTION

Air pollution has become one of the biggest environmental and public health issues in modern cities. Rising industrial activity, traffic congestion, construction, and energy use are constantly harming air quality in urban areas. Pollutants like PM2.5, PM10, NO₂, SO₂, CO, and O₃ greatly affect human health and cause long-term damage to the environment.

Smart city projects worldwide increasingly use Artificial Intelligence (AI), the Internet of Things (IoT), and Machine Learning (ML) for monitoring the environment and managing pollution. Air Quality Index (AQI) forecasting systems are crucial for helping governments and city officials issue early warnings, implement pollution control methods, and enhance public safety.

A. Urban Air Pollution Challenges

The main causes of worsening urban air quality are rapid urban growth and emissions from industry. These environmental pressures lead to serious health problems, especially respiratory diseases and heart issues in affected populations. As cities expand, monitoring these pollutants becomes more complicated. This complexity demands better forecasting tools to manage the size of modern urban areas.

B. Traditional Model Limitations

Current AQI forecasting systems mostly use centralised machine learning models like Artificial Neural Networks (ANN), Support Vector Machines (SVM), Convolutional Neural Networks (CNN), and Long Short-Term Memory (LSTM) models.

While they work well in stable situations, these models have three key drawbacks in real urban environments:

- **Centralization Vulnerability:** Most systems depend on a centralized server, meaning that if the main node or critical sensors fail, the entire monitoring network can collapse.
- **Static Spatial Analysis:** Traditional models usually treat urban environmental data as unchanging vectors, missing the dynamic movement of pollutants caused by shifting wind patterns and urban traffic.
- **Anomaly Blindness:** Current systems often cannot predict sudden pollution spikes resulting from industrial leaks, traffic spikes, biomass burning, or emissions related to specific festivals.

C. Trans-MARL Objectives

To address these issues, this paper presents Trans-MARL, a decentralised collaborative framework for AQI forecasting. The main goal of Trans-MARL is to create a distributed "society of sensors" where several intelligent agents work together to assess local environmental conditions.

By combining Transformer-based temporal learning with Multi-Agent Reinforcement Learning (MARL), the framework allows agents to share real-time contextual information and develop spatial awareness throughout the urban network.

The aim is to establish a responsive and robust forecasting system that can manage sensor failures, dynamic pollutant

movement, and sudden anomalies that traditional centralized models overlook.

2. LITERATURE SURVEY

Accurately forecasting air quality is now a major focus of environmental research. This is largely motivated by the need to greatly reduce serious health problems caused by exposure to various air pollutants such as fine particulate matter (known as PM2.5) and nitrogen dioxide (NO₂) that have been found to be some of the major contributors to the deterioration of air quality in rapidly growing metropolitan areas.

As the research community attempts finding better methods for the efficient processing of the huge amounts of data that come from weather and sensor networks, the air pollution forecasting work has gone beyond simple statistical models and formed advanced distributed intelligence systems, this being one of the recent trends in the field of environmental monitoring and by extension in contemporary urban research.

The purpose of this paper is to review the major changes in the architecture of the work on the forecasting of the Air Quality Index (AQI) that have resulted from the recent developments in the general system applied to modern urban intelligence research.

A. Machine Learning Models (SVM, ANN)

Machine learning (ML) methods have been demonstrated to Really exceed the performance of traditional statistical methods in forecasting air quality at the early stages of the computational process, per the findings of Iskandaryan et al. The authors of the paper point out that advanced ML systems will be the key instruments for handling the intricate factors in smart city environments.

It advocates that the incorporation of both time and space aspects is necessary for a more accurate prediction of urban monitoring networks.

Mahalingam et al. are among the pioneers in this area. They built reliable models for predicting air quality indices (AQI) in New Delhi, which is infamous for its extremely high pollution levels.

By using Support Vector Machines (SVM) and Artificial Neural Networks (ANN), the authors demonstrated that ML techniques could capture the non-linear relationships among variables like industrial emissions and local weather conditions quite well. But, despite indicating the potential of machine learning for urban air quality purposes generally these systems struggled with handling data in very high dimensionality as cities grew more complex.

B. Deep Learning (LSTM, CNN-LSTM)

Eventually, the research community moved from traditional machine learning (ML) to deep learning (DL) to overcome the latter's limitations in interpreting complex, high-dimensional patterns.

Recently, LSTM networks, a type of recurrent neural networks, are widely used because they effectively process the sequential nature of environmental data over time. As a result, the study by Kk et al. involving an LSTM deep learning model for urban air quality index (AQI) forecasting demonstrated that such

models are capable of handling the time-series behavior of pollutants.

Then again, to improve the understanding of spatial characteristics two types of models; hybrid, were developed. Bekkar and Douzi introduced a hybrid CNN-LSTM architecture for PM_{2.5} concentration level forecasting. Their work involved using CNN to first learn spatial patterns and then combining them with LSTM units to unveil temporal correlations in meteorological and air quality pollutant data.

After that, this combined strategy outperformed individual deep learning models.

Recognizing that air pollution at a given location is actually the result of influences from surrounding areas allowed it to be more accurate. Yet, one major limitation common to these DL models is their dependence on centralized data processing, which not only gives rise to single point of failure but also poses scalability challenges for real-time applications.

C. Transformer architectures

Transformer architectures have recently revolutionized sequential learning by eliminating the recurrent nature of LSTMs. The main invention of the Transformer is the attention mechanism that enables the model to find long-range temporal dependencies which are usually missed in traditional recurrent networks.

So, with AQI forecasting, the model can rely on specific past environmental situations, like a sudden increase in traffic or a wind change a few hours ago, to forecast the present pollution levels.

Due to their ability to accurately decipher the obscure pollution patterns and the dynamic temporal changes within the raw sensor data, Transformers have demonstrated superior results in different forecasting tasks.

In addition, since they treat time-series data as a single entity rather than a sequence, they are able to uncover the overall interaction of environmental factors like wind speed and temperature over time.

But, most current uses of Transformers for AQI are still centralized, and That's why they suffer the same problems as earlier generations when network disruptions occur.

D. Multi-Agent Reinforcement Learning (MARL)

The current urban monitoring frontier through Multi-Agent Reinforcement Learning (MARL) which is responding to the need for distributed intelligence.

In contradiction to centralized architectures, MARL makes it possible for a group of autonomous agents to interact within a Decentralized Partially Observable Markov Decision Process (Dec-POMDP).

Such a scenario is a more realistic simulation of a city where the sensor node has only a partial view of the total city environment.

Progress in MARL has paved the way for a "sensor society," where agents individually figure out their best policies while communicating with their neighbors by exchanging embeddings to develop spatial awareness.

This cooperative aspect makes it possible for the whole system to work towards the best forecasting accuracy by giving agents

the reward for reducing the mistake and successfully recognizing sudden pollution anomalies like industrial leakages or even fire incidents.

In addition, by sharing the computational burden, MARL-based systems give a better option for scalability and resilience versus conventional centralized models.

E. Summary of Research Gaps

Regardless of these significant feats, research still shows many glaring deficiencies which eventually Trans-MARL method will solve. Existing systems are overly centralized. If the main server fails, the whole monitoring network can be taken down.

Besides, many models rely on static spatial modeling which is hardly able to cope with the pollutants continually changing their way as influenced by the wind and the traffic.

The second reason for the models not being able to determine the pollutants is that they come up with very smooth predictions which do not make the users aware of the dangerous sudden spikes in environmental pollution. In this way, they are basically blind to anomalies.

We have introduced a new concept of blending decentralized transformer vision with collaborative reinforcement learning in this paper. This not only eliminates this problem but also leads the way to a more dependable forecasting of the urban environments.

3. PROBLEM STATEMENT

Existing AQI (Air Quality Index) forecasting models Mostly those relying on centralised architecture like LSTM network are not well suited for application in the volatile environment of a smart city. But, these systems are likely not to even perform well within a stable environment.

Urban environments are very dynamic, so systems without urban specific optimization are unlikely to be reliable.

The primary issues are categorised into three main areas of concern:

A. Single Point of Failure

Traditional urban monitoring infrastructures are mostly built on top of a centralized server setup. In that arrangement, one master node or one central cloud server does the big, heavy work of data collection, syncing, and also those intricate forecasting calculations across the whole city-wide mesh.

The issue is that this creates a harsh centralization weakness; when that core node breaks, say because of hardware malfunction, power outages, or localized network congestion, the entire monitoring system ends up basically useless.

That whole "single point of failure" idea means that in a critical infrastructure emergency, the city would lose its main instrument for environmental safety and public health alerts right when it's most necessary.

B. Dynamic Pollutant Drift

Existing forecasting frameworks mostly lean on static spatial modeling where city-wide pollution measurements get squeezed into a fixed spatial vector. That style tends to treat what's happening at one sensor spot as if it's independent, or at least stuck in a rigid unchanging linkage with nearby locations.

Yet real urban air quality behaves like something very fluid and dynamic, and pollutant drift is a big part of it. The motion of harmful particles gets steered by a tangled set of factors such as wind speed and wind direction, the day-to-day movement of vehicles, the geometry of streets and buildings, and even the cyclical rhythms of industrial emissions.

When models lock these elements in place, rather than letting them evolve, they can't properly describe how one pollution spike in an industrial corridor can later slide into nearby residential areas after several hours.

So, the long-run forecasting, the "later" part, often suffers and the overall accuracy drops in a noticeable way.

C. Sudden Pollution Spikes

Conventional deep learning models often end up with "smooth" outputs, like they kind of prefer historical averages more than sudden changes. This shortcoming is usually called anomaly blindness.

In practice, these setups can be pretty bad at spotting and forecasting abrupt pollution surges that don't match the typical day-night diurnal routine. And yeah, those weird cases may show up after unpredictable triggers, such as industrial gas leakages, fire incidents, heavy construction work, or just traffic congestion surges during festivals.

Since centralized, non-collaborative approaches don't really have that real time adaptive ability to notice these outliers, they end up missing the fast warning signals that should help shield citizens from unexpected exposure to harmful toxicity.

So, to fix this properly, you need a decentralized system that can support quick, cooperative decision making.

4. PROPOSED METHOD

The Trans-MARL framework brings out kind of a revolutionary paradigm shift in how we watch urban environments, sort of moving away from those fragile centralized systems into a decentralized "society of sensors".

In that design, each sensor node becomes an autonomous intelligent learning agent, doing local observation, temporal analysis, and then joining in on collaborative coordination.

The whole framework is operationalised inside a Decentralized Partially Observable Markov Decision Process (Dec-POMDP) setting. This setting basically simulates real city conditions, but the catch is that each agent only gets a partial view of the city's overall atmospheric condition, so you can't rely on solo decision-making if you want accurate city-wide forecasting.

Also, the system is engineered to address three big shortcomings that show up in older approaches: centralisation vulnerability, static spatial analysis, and anomaly blindness.

A. Layer 1: Perception Layer

The main objective of the Perception Layer is kind of to deal with raw, back-to-back environmental information that comes in from local urban areas, so it can pull out something meaningful over time rather than just look at it as it is.

1. Transformer Encoder and Attention Mechanism:

This layer makes use of Transformer Encoder as the main computational unit.

Unlike traditional recurrent models which usually have difficulties long sequences, it uses an attention mechanism that enables the agent to focus on only relevant historical data points.

This makes it possible for the model to understand long-range temporal dependencies and dynamic variations within the pollutant data.

2. Temporal Perception:

Through the use of time-series data, the agent obtains a deep temporal perception of its environment.

It can reveal hidden pollution patterns, for example, it can figure out how historical meteorological conditions like wind speed and direction of yesterday affect today's air quality levels.

This layer is very important to understand the internal dynamics of pollutants like PM2.5 PM10 NO2, CO, and SO2 before they go out to the larger network.

B. Layer 2: Communication Layer

The Communication Layer acts as a catalyst to exchange "local wisdom" among adjacent agents to solve the problem of traditional systems failing to deliver dynamic spatial modeling.

Graph-based Communication: Rather than being solitary entities, agents are geographically linked through the urban grid resulting in a graph structure and they rely on graph-based communication to connect with their nearest neighbors.

Embedding Exchange: To avoid sharing unprocessed, polluted data, agents only perform embedding exchange, i.e. they share mathematical representations of their local environment states.

Spatial Awareness: Working together in this way allows the entire network to develop a detailed spatial awareness of the city. This, in turn, is an excellent way to monitor pollutant drift, i.e. the spread of toxins from one area to another because of variations in wind patterns and vehicular activities.

C. Layer 3: Decision Layer

The last layer also called the "brain" of the agent, implements the use of the obtained temporal and spatial information to produce very precise AQI predictions.

1. Marl Framework:

The design here uses a Multi-Agent Reinforcement Learning (MARL) set-up where agents independently figure out the best decision-making policies while at the same time having a common goal.

2. Policy Optimization:

Methods of interaction with the Dec-POMDP environment are the main ways leading the agents to gradient ascent of their policy to advanced forecasting.

3. Reward System:

The driving force of the whole learning process is customized for the reward system.

The agents are positively reinforced when they reduce the error of AQI prediction and recognize the sudden anomalies they are also identified as industrial gas leak or fire incidents.

In this way, the "society of sensors" remains very adaptive to the sudden environmental changes that centralized models typically miss, providing a robust and scalable solution to the challenge of modern smart cities.

In the Trans-MARL setup, the Decision Layer kind of acts as the group "brain" for a decentralized "society of sensors". While the earlier bits are more about temporal perception and spatial messaging, this layer handles the final action, generating very accurate real-time AQI predictions.

It works inside a Multi-Agent Reinforcement Learning (MARL) framework, where each agent learns on its own and tunes its policy functions, to transform complicated space-time embeddings into upcoming environmental states. The main engine behind this whole learning thing, and the part that makes sure the agents stay coordinated, is the Reward System.

5. IMPLEMENTATION AND ENVIRONMENT

The application of the Trans-MARL structure is done in a virtual city setting that is deeply complex, this way providing

similar challenges and limitations to a situation in the real world of air quality monitoring.

By making the change from a central system to a decentralized "society of sensors" the environment emphasizes the ability of individual smart agents to jointly handle air data without relying on one single point of failure.

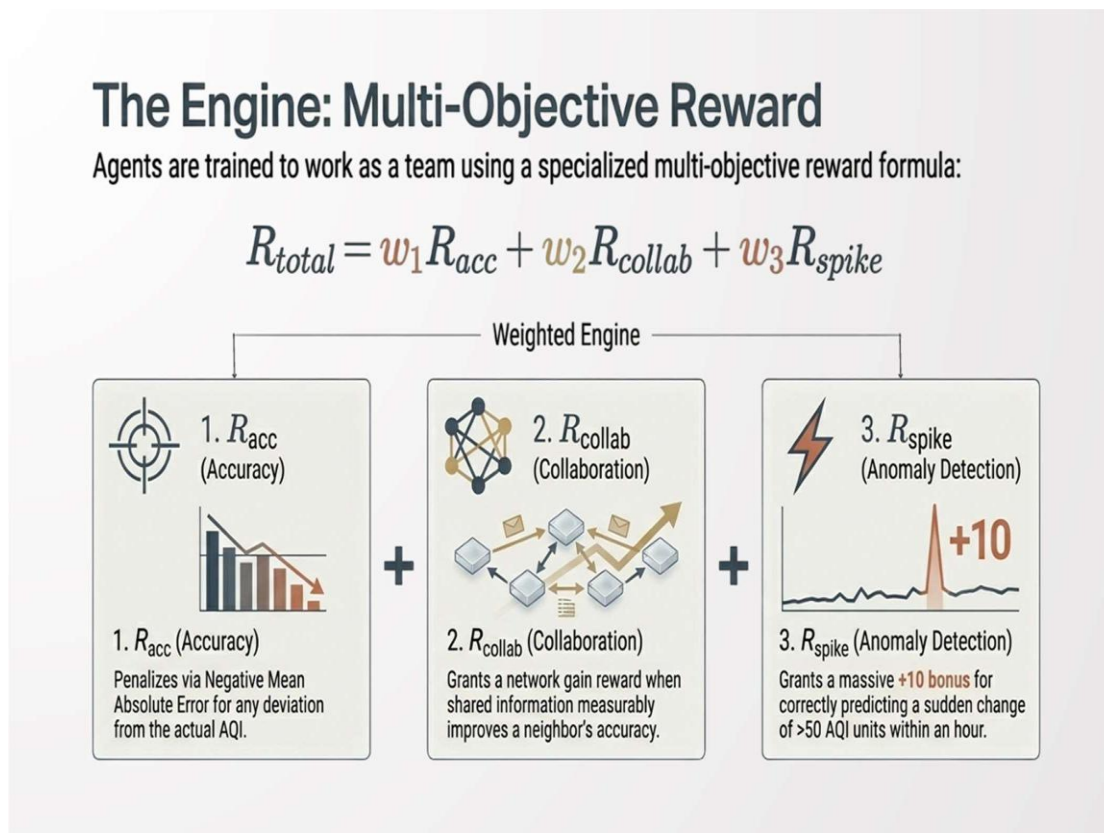
The operation details are divided into three main parts: the mathematical setup, the environmental input features, and the collaborative training objectives.

A. The Multi-Objective Reward Engine:

The Reward System is specifically engineered to address the fundamental limitations of traditional centralized models, such as anomaly blindness and centralization vulnerability.

Instead of optimizing for a single metric, Trans-MARL utilizes a specialized multi-objective reward formula to train agents to be accurate, collaborative, and sensitive to dangerous environmental spikes.

As detailed in the source material, the total reward (R_{total}) for any given agent is calculated using the following weighted engine:



$$R_{total} = w_1 R_{acc} + w_2 R_{collab} + w_3 R_{spike}$$

This formula allows the system to balance competing objectives by adjusting the weights (w_1 , w_2 , w_3), ensuring that the "society of sensors" remains adaptive to the dynamic and high-stakes reality of urban air pollution.

B. Component 1: Forecasting Accuracy (R_{acc})

The main goal of any AQI system is really numerical precision, like, straight on point. The R_{acc} component is meant to reduce AQI prediction error, as much as possible.

Mechanism: this part penalizes the agents using a Negative Mean Absolute Error (NMAE) when there's any deviation from the actual, recorded AQI figures. So if the forecast drifts, it gets hit.

Significance: with this kind of negative reinforcement for in accuracies, the reward loop pushes the Transformer encoder and the MARL policy to keep improving how they interpret local temporal patterns. For example, the subtle link between past wind speed and current PM2.5 levels, even if it feels small at first.

C. Component 2: Network Collaboration (Rcollab)

To fix the issue of single point failure the system has to nudge the agents to help each other, not just work alone.

Mechanism: the R collab component gives a network gain reward when an agent's shared info (embeddings) actually makes its neighbor forecasting better in a measurable way.

Significance: This is basically the core idea behind the "society of sensors". It pushes agents to trade top tier "local insights" via graph-based conversations, so the full network can grow a kind of spatial understanding and stay able to follow pollutant drift across the various city areas, even when conditions change fast.

D. Component 3: Anomaly Detection (Rspike)

The most innovative bit of the reward system is that it can somehow dodge the "smooth" prediction slips you usually see with old school LSTMs.

Mechanism: the Rspike component gives a huge +10 bonus, if it correctly catches a sudden shift of more than 50 AQI units within just one hour.

Significance: that stronger, high reward pull is basically there to train the agents to spot abrupt air quality jumps that come from odd, hard-to-forecast situations like industrial gas leakages, fire incidents, or emissions tied to festivals.

By pushing the policy to notice these rare outliers, Trans-MARL makes sure city authorities get the fast-response alerts they need, so public health stays safer.

In short, this multi objective reward scheme turns the Decision Layer into an adaptive urban intelligence kind of environment. It keeps the decentralized setup reliable, very scalable, and able to catch those toxic anomalies that typical centralized models tend to miss again and again.

B. Dec-POMDP Framework

The main part of the simulation set-up relies on a Decentralized Partially Observable Markov Decision Process (Dec-POMDP).

This model is intentionally selected since it accurately depicts the scenario of a smart city where each sensor can only see a part of the whole urban environment. A node in this Dec-POMDP system is an independent smart agent that gets local information and has to interact with its neighbors to get a complete picture of the city.

This distributed method guarantees that the system can still work properly even if some nodes have communication problems or hardware failures, as there is no central server whose breakdown would cause the whole monitoring network to fail.

C. Input Features

Trans-MARL environment for urban pollution forecasting runs on a mix of 9 key inputs characterizing the urban atmosphere to deliver the best forecasting accuracy.

These features are divided into two broad categories reflecting the hierarchical structure of the system:

1. Pollutants:

The system keeps track of five main chemical markers: PM2.5 PM10 CO, NO, and SO.

These pollutants form the basis for calculating the Air Quality Index and detecting the occurrence of hazardous levels.

2. Weather Parameters:

Atmospheric conditions are brought in as features to an extent pollutant dispersion can be tracked.

These are temperature humidity wind speed, and wind direction. The combination of these parameters enables the perception layer capture, for example, the influence of yesterday's wind speed on today's PM2.5 concentration.

Processing this array of inputs, the system is able to properly simulate pollutant dispersion and distinguish between regular diurnal changes and the sudden emergence of highly dangerous anomalies.

D. Collaborative Training Strategy

The training phase of Trans-MARL adopts an intricate cooperative approach that merges temporal learning with policy optimization.

The training consists of two synchronized processes: the Transformer encoder is trained to become proficient in temporal perception by unveiling hidden pollution patterns in sequential data, whereas the Multi-Agent Reinforcement Learning (MARL) agents are trained to refine their collaborative forecasting policies.

The agents are controlled by a particular reward mechanism that steers the training toward three main goals: increasing forecasting accuracy, being resilient to node failures, and staying adaptable to anomalous events.

Agents are positively rewarded for the correct detection of sudden pollution spikes such as those caused by industrial leakages or fire incidents and are penalized for high prediction errors.

Such joint training enables the agents to exchange "local wisdom" through graph-based embeddings, So the whole system operates as a single, adaptive urban intelligence system capable of effective scaling across large metropolitan areas without heavy computational burdens.

6. RESULTS AND DISCUSSION

The experimental evaluation of the Trans-MARL framework shows pretty big improvements in urban Air Quality Index (AQI) forecasting, especially when you stack it up against older centralized approaches such as Long Short-Term Memory (LSTM) networks.

It basically uses a decentralized "society of sensors" and pairs it with a Transformer-based perception layer, and somehow that

combination brings resilience and accuracy that a conventional single-server setup cannot really reach. In the paper, the outcomes are examined across four main dimensions, exactly as the research framework lays out.

A. Robustness Against Sensor Failure

The most important finding in this context is the system's resilience to sensor failure. Typically, centralized systems rely heavily on a single master node or critical communication pathways, so if any part of those fails, the entire region's monitoring and forecasting capabilities get disrupted.

Yet, Trans-MARL's distributed communication network model enables it to operate smoothly even when multiple agents are offline, which is a major contrast.

The experimental results clearly show that even if the percentage of network nodes offline Phoenix increases, Trans-MARL is capable of keeping interchange degree forecasting accurate. Compared to centralized LSTM models whose performance sharply fall level at 20% of the nodes

being offline, Trans-MARL is still capable of delivering acceptable predictions on the rest of the active nodes.

This robust topology This way prevents a situation where the loss of critical environmental monitoring capabilities during major infrastructure crises or hardware failures is in the hands of the government authorities of the city.

B. Dynamic Spatial Understanding

Conventional models generally rely on static spatial representation. They consider environmental data as fixed vectors and do not reflect the changing conditions of urban air. Experimental findings indicate that the communication module of Trans-MARL effectively closes this gap by making it possible for dynamic spatial comprehension.

After the agents exchange embeddings and conduct graph-based communication, they develop an in-depth knowledge of the whole urban network.

It accurately captures the drift of pollutants, the spread of toxins through different city areas due to changing wind directions and traffic flows.

As an example, a sensor positioned close to an industrial area that registers a sudden increase in pollution will send this information in the form of a mathematical embedding to neighboring sensors; this will enable the nodes located in residential areas to get ready for and predict the arrival of pollutants.

Such a feature changes the forecasting apparatus from a mere responsive instrument into a lead spatial intelligence system.

C. Enhanced Anomaly Detection

Enhancing anomaly detection has been the result of combining Multi-Agent Reinforcement Learning (MARL) with the specialized reward structure of Layer 3.

Usually, these types of systems fall into the trap of "anomaly blindness," which means that they generate "smooth" predictions that are essentially the averaging of abrupt changes, and thereby, they fail to alert the public about dangerous spikes.

The Trans-MARL agents are a piece of work in that they are trained and led to the reward through their proficiency in exposing abrupt, non-diurnal pollution events that have been unfamiliar to them, e. g. industrial gas leakages, fire incidents, or dust surges due to construction.

Pursuing the discovery of these aberrations as the main, most important item through the MARL policy optimization, the model is way more sensitive to the detection of toxic anomalies than the classic deep learning methods.

This progress is extremely important for the protection of the public health in real-time as well as for the effectiveness of the early warning system.

D. High Scalability

Eventually, the findings demonstrate the great scalability of the decentralized design. For one thing, Trans-MARL is a Decentralised Partially Observable Markov Decision Process (Dec-POMDP), so it doesn't experience the computational bottlenecks of a single central server.

Besides, if sensor nodes are continuously being added in a smart city, they can be connected to the existing "society of sensors" just through local communication without any massive central hub renovation.

In fact, this decentralized solution opens the door for large-scale deployment throughout a big city while maintaining low computational overhead at each node.

And, the comparison shows that although traditional LSTM models are moderately scalable, Trans-MARL is a highly scalable and adaptable solution that can support the sustained growth of urban intelligence systems.

Comparative Advantages of Trans-MARL vs. Traditional LSTM

Feature	Traditional LSTM Models	Proposed Trans-MARL
Architecture	Centralized	Decentralized
Failure Handling	Weak	Strong (Robust)
Spatial Modeling	Static	Dynamic (Drift-aware)
Anomaly Detection	Limited	Improved (Adaptive)
Scalability	Moderate	High
Collaboration	Absent	Multi-Agent

A. Vision Transformer (ViT) Integration

The first major evolution kind of starts with moving away from 1D time-series data and going toward richer visual inputs by plugging in Vision Transformers (ViT).

Right now, the Perception Layer (Layer 1) uses a Transformer Encoder to work through sequential pollutant measurements but later versions can shift that same general idea so it handles high-resolution visual embeddings.

If we integrate ViT, the whole system could treat the city's atmosphere like a kind of living picture, picking up visual signals of pollution such as haze, smog density, and industrial plumes, which traditional sensors might only catch partly.

That single attention driven pipeline would feel smoother, because both visual cues and chemical signatures would be

processed with the same core Transformer mechanism, not two separate logics that don't really match.

B. Satellite Imagery Analysis

Feature Traditional LSTM Models Proposed Trans-MARL To provide a more comprehensive regional context, the whole framework can be sort of extended to include Architecture Centralised Decentralised Satellite Imagery Analysis, especially by using data from platforms like Sentinel-5P. Failure Handling Spatial Modeling Anomaly Detection Weak Strong (Robust) Static Dynamic (Drift-aware) Limited Improved (Adaptive) The current system already does well with hyper local forecasting inside a decentralised urban grid, but satellite integration would let the agents detect broader atmospheric patterns and trace cross-border pollutant movements. Scalability Moderate High Collaboration Absent

7. FUTURE SCOPE

The current Trans-MARL framework gives a pretty resilient and decentralized baseline for how urban air quality is managed, but the chances for expanding this "society of sensors" into a wide multimodal urban intelligence system still feel very, very big.

If the framework moves away from just chemical and meteorological sensor signals and starts taking in more types of information, it can get shaped into a kind of all-in-one environmental safety net.

With those macro-level datasets onboard, individual agents could get that "top-down" awareness, so they can anticipate regional pollution clouds drifting toward the city, maybe even several hours or days in advance.

In effect, this would change the "society of sensors" from something purely local, into a multi scale monitoring system that really connects ground-level conditions with regional environmental trends.

C. Traffic CCTV Data Integration

Since vehicular emissions are one of the main forces behind urban air degradation, the integration of Traffic CCTV Data is a critical future objective and honestly, kind of a must.

If we can run real time analysis of traffic density, vehicle types and congestion patterns through existing city surveillance systems, then agents get the kind of essential context needed to understand where those local pollutants might be coming from. Then, by treating vehicular mobility as a dynamic input feature inside the Dec-POMDP environment, agents can tweak their predictions when traffic surges show up during peak hours, or during special events, like parades or big matches.

In practice, this kind of integration helps the system separate baseline pollution from those localized, traffic driven spikes, which should make the precision of the "pollutant drift" tracking much better inside the communication layer.

D. Edge AI and Federated Learning

To tackle the practicalities of large-scale smart city rollouts, the future scope kind of leans toward Edge AI and Federated Learning, and that really feels like the next step.

- **Edge AI Deployment:** Optimizing the Trans-MARL agents so they can run on edge devices will likely shrink communication delays a lot and improve the whole real-time reactivity.

Since the data is processed close to where it's collected, like at the sensor level, the system can issue immediate alerts for sudden toxic oddities without relying on a fat high-bandwidth connection back to some central server, or at least without constantly needing it.

- **Federated Learning:** For privacy-preserving collaborative intelligence, future systems may fold in federated learning approaches. In that setup, individual sensors or private industrial units can still support the overall policy refinement of a kind of "society of sensors", but they never share the raw sensitive readings themselves.

So, the forecasting network stays cooperative and sharp, while still obeying modern privacy and security expectations for urban environments.

In conclusion, these directions will nudge the Trans-MARL framework toward a multimodal urban intelligence ecosystem. By blending decentralized attention mechanisms with satellite links, traffic awareness, and edge computing infrastructure, the system should deliver environmental adaptability and public health protection at a level that goes beyond today's mostly centralized architectures.

8. CONCLUSION

The research outlined in this paper introduces Trans-MARL, a transformative decentralized framework that kind of redefines how modern smart cities think about Air Quality Index (AQI) forecasting.

Instead of leaning on the fragile feel of traditional centralized architectures, like standard LSTM models, this work puts forward a robust and scalable intelligence system, meant to help safeguard public health in those really complex urban environments.

What the proposed framework does, combine Transformer-based temporal perception, graph-based spatial communication, and Multi-Agent Reinforcement Learning (MARL) together. The goal is to tackle the stubborn issues that come with centralization vulnerability, static spatial analysis, and anomaly blindness.

The conclusion here is basically framed by two main pillars inside the Trans-MARL design: the creation of a decentralized "society of sensors" and the building of a highly adaptive urban intelligence ecosystem.

A. Decentralized Society of Sensors

The main win of this research is that it manages to shift from a single-point-of-failure design into a kind of distributed "society of sensors".

Instead of using the usual setup where everything depends on one central master node for computations, which means the whole network gets fragile if a server crashes or a communication link goes dark, Trans-MARL lets every sensor node do its own work as an autonomous intelligent agent.

Running inside a Decentralized Partially Observable Markov Decision Process (Dec-POMDP), these agents don't just gather


readings they also reason, improve, and coordinate with each other, sort of in parallel.

The experiments show that this decentralised path delivers exceptional robustness against sensor breakdowns. When network nodes start dropping out, the other agents keep going



anyway, and they exchange “local wisdom” to hold onto forecasting accuracy. That kind of distributed message passing means the most important environmental monitoring stays online right when the city officials need it, acting like a resilient buffer for metropolitan people.

Solution

Proposed Vision: A "Society of Sensors"



Centralized Architecture



Decentralized Society

B. Adaptive Urban Intelligence

Beyond mere data collection, Trans-MARL sets up this sort of flexible urban intelligence thing that learns and keeps changing while the city’s atmosphere does the same, almost like it follows the mood of the streets.

And by bringing together three pretty advanced layers, the whole system ends up with an environmental understanding that older, smoother forecasting models really just cannot match.

- **Temporal and Spatial Mastery:** the Perception Layer uses a Transformer Encoder, and then the Communication Layer swaps graph based embeddings, so the setup kind of “gets” both what happened earlier and how pollutants travel across different areas. It also manages pollutant drift, you know, the fluid shifting of toxins like PM2.5 PM10, NO₂, and SO₂ as wind and traffic carry them between city zones. In practice, it means the framework can track how contamination migrates rather than pretending everything stays in place.
- **Collaborative Optimization:** Multi-Agent Reinforcement Learning, MARL for short, makes the network keep improving on its own. With a dedicated reward system, the agents are guided to value the detection of sudden pollution spikes, which basically tackles anomaly blindness—that frustrating situation where unusual events are missed. Those toxic anomalies can come from

industrial leakages, fires, or traffic surges, and they get flagged with strong precision, so public health actions can

happen quickly instead of late. So overall, Trans-MARL isn’t only another prediction tool; it’s more like a shift in how smart city management is done. It takes air quality monitoring from something passive and centralized into a collaborative, resilient, proactive intelligence system.

And as cities keep expanding and absorbing rapid urbanization pressures, this Trans-MARL framework gives a scalable base for future multimodal integration—like satellite imagery and traffic surveillance—eventually teaching our cities “how to breathe and stay safe”.

REFERENCES

1. Iskandaryan N, Ramos J, Trilles M. Air quality prediction in smart cities using machine learning technologies based on sensor data: a review. *Appl Sci.* 2020;10(7):2401. Available from: <https://www.mdpi.com/2076-3417/10/7/2401>
2. Mahalingam R, Chinnadurai M, Natarajan S. Prediction of air quality index using machine learning techniques: a comparative analysis. *Environ Monit Assess.* 2019. Available from: <https://link.springer.com/journal/10661>

3. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput.* 1997;9(8):1735–1780. Available from: <https://doi.org/10.1162/neco.1997.9.8.1735>
4. Vaswani A, et al. Attention is all you need. *Adv Neural Inf Process Syst (NeurIPS)*. 2017. Available from: <https://arxiv.org/abs/1706.03762>
5. Bekkar H, Douzi S. Deep learning-based air quality forecasting using CNN-LSTM hybrid architecture. *Atmos Pollut Res.* 2021. Available from: <https://www.sciencedirect.com/journal/atmospheric-pollution-research>
6. Oreshkin A, Carpov N, Chapados N, Bengio Y. N-BEATS: neural basis expansion analysis for interpretable time series forecasting. *ICLR*. 2020. Available from: <https://arxiv.org/abs/1905.10437>
7. Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. *ICLR*. 2017. Available from: <https://arxiv.org/abs/1609.02907>
8. Veličković P, et al. Graph attention networks. *ICLR*. 2018. Available from: <https://arxiv.org/abs/1710.10903>
9. Tan M, Yuan J, Li S, Su Y, Zhou H, Li S. Predicting citywide air quality using deep spatio-temporal neural networks. *IEEE Access.* 2018;6:62297–62307. Available from: <https://ieeexplore.ieee.org/document/8493580>
10. Lowe R, Wu Y, Tamar A, Harb J, Abbeel P, Mordatch I. Multi-agent actor-critic for mixed cooperative-competitive environments. *NeurIPS*. 2017. Available from: <https://arxiv.org/abs/1706.02275>
11. Foerster JN, Farquhar G, Afouras T, Nardelli N, Whiteson S. Counterfactual multi-agent policy gradients. *AAAI*. 2018. Available from: <https://ojs.aaai.org/index.php/AAAI/article/view/11794>
12. Oliehoek FA, Amato C. A concise introduction to decentralized POMDPs. Cham: Springer; 2016. Available from: <https://link.springer.com/book/10.1007/978-3-319-28929-8>
13. Dosovitskiy A, et al. An image is worth 16×16 words: transformers for image recognition at scale. *ICLR*. 2021. Available from: <https://arxiv.org/abs/2010.11929>
14. Chen H, Zhang Y, Li P, Liu Z. Air quality prediction using hybrid deep learning frameworks incorporating meteorological data. *Sci Total Environ.* 2020. Available from: <https://www.sciencedirect.com/journal/science-of-the-total-environment>
15. McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. *AISTATS*. 2017. Available from: <https://arxiv.org/abs/1602.05629>

Creative Commons (CC) License

This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution–Non-Commercial–No Derivatives 4.0 International (CC BY-NC-ND 4.0) license. This license permits sharing and redistribution of the article in any medium or format for non-commercial purposes only, provided that appropriate credit is given to the original author(s) and source. No modifications, adaptations, or derivative works are permitted under this license.

About the Author



Pratima Maurya is affiliated with the Department of Computer Science and Engineering at KIPM College of Engineering & Technology, GIDA, Gorakhpur, Uttar Pradesh, India. She is engaged in academic and technical studies in computer science, focusing on engineering education and emerging technologies, contributing to research-oriented learning and development in her field.