Journal of Urban Development and Smart Cities
2025, 2(1), 161-172


# OPTIMISING AIR QUALITY PREDICTION FOR SUSTAINABLE URBAN DEVELOPMENT AND SMART CITY MANAGEMENT USING A HYBRID PSO-LSTM-RNN FRAMEWORK

*Deniz Orhun*

*Accurate air-quality forecasting is increasingly central to urban environmental governance, public-health protection, and data-driven smart-city management. This manuscript presents a polished and publication-ready account of a hybrid particle swarm optimisation (PSO), long short-term memory (LSTM), and recurrent neural network (RNN) framework for forecasting urban air quality using the University of Utah Air Pollution Monitoring Network dataset for Salt Lake City, Utah. The empirical setting is based on 25 pollution sensors, hourly aggregation, and a study window spanning 2019-07-26 to 2021-05-14. The modelling pipeline integrates data pre-processing, missing-value handling, outlier treatment, recurrent sequence learning, and PSO-based hyperparameter optimisation, followed by a curiosity-based motivation mechanism that strengthens the hybrid recurrent design. In the reported experiments, the proposed model achieves a mean absolute error of 0.0082 and an $R^2$ score of 0.1227 in the principal benchmark comparison, while 10-fold cross-validation yields an average RMSE of 0.013, MAE of 0.010, MAPE of 3.8%, and $R^2$ of 0.94. Additional analysis shows that hyperparameter tuning materially improves predictive performance, and ablation results indicate that the full integration of LSTM, RNN, PSO, and curiosity-based motivation produces the strongest results. Framed for the urban development and smart-city literature, the study demonstrates that hybrid recurrent forecasting can support municipal decision-making in air-quality surveillance, exposure mitigation, traffic-responsive environmental control, and sustainable service planning.*

*Index Terms* — air-quality prediction; smart cities; urban development; environmental monitoring; particle swarm optimisation; LSTM; recurrent neural network; sustainable urban management



https://doi.org/10.66033/judsc2025-217

## INTRODUCTION

Urban development strategies are increasingly shaped by digital infrastructures capable of sensing, processing, and responding to environmental conditions in real time. Within this broader transformation, the smart-city paradigm has expanded beyond conventional concerns of connectivity and automation to include predictive, data-driven governance systems that can support timely and evidence-based urban management. Among the many environmental variables that influence the performance and sustainability of cities, air pollution remains one of the most consequential because it affects not only public health, but also transport planning, energy efficiency, regulatory compliance, land-use quality, and the overall liveability of urban environments. For this reason, air-quality prediction should not be viewed merely as a technical time-series exercise; rather, it should be understood as an essential component of operational planning and urban governance in digitally enabled cities [1].

Air pollution poses a particularly complex challenge in rapidly evolving urban systems because pollutant concentrations are influenced by multiple interacting factors, including traffic density, industrial emissions, meteorological variability, seasonal patterns, built-environment form, and human activity rhythms. These interactions create nonlinear and time-dependent dynamics that are difficult to capture using traditional forecasting approaches alone. At the same time, municipal authorities increasingly rely on sensor-based monitoring networks and urban data platforms to support environmental surveillance and rapid intervention. In this context, the ability to forecast future air-quality conditions with reliability is of direct practical importance. If pollution episodes can be anticipated rather than merely observed after they occur, city administrations can issue earlier health advisories, adapt transport operations, reduce exposure in sensitive zones, coordinate emergency responses, and implement preventative environmental measures before conditions worsen [2, 3, 4, 5]. Predictive air-quality analytics therefore occupies a central place in current discussions of sustainable urban development, environmental resilience, and smart-city management.

The study presented here addresses this challenge through a hybrid PSO–LSTM–RNN forecasting framework developed for smart-city air-quality prediction. The rationale for this hybrid architecture arises from the limitations of relying on a single modelling strategy in complex urban sensing environments. Air-quality data are typically sequential, noisy, nonlinear, and sensitive to local fluctuations, which makes them well suited to recurrent deep-learning approaches that can learn temporal dependencies from historical observations. Long short-term memory (LSTM) networks are particularly effective at handling long-range temporal patterns, while recurrent neural network (RNN) structures can further support sequential representation learning. However, the predictive capacity of such models depends strongly on the selection of hyperparameters, training dynamics, and architectural configuration. Particle swarm optimisation (PSO) is introduced in this framework to address that challenge by systematically searching for improved parameter settings and thereby strengthening model performance and stability. The resulting PSO–LSTM–RNN framework is intended not only to improve numerical forecasting accuracy, but also to enhance the practical reliability of predictive systems designed for real urban applications.

From an urban-development perspective, the significance of this work lies in its direct connection between forecasting methodology and municipal decision support. In smart cities, environmental prediction models are valuable only when they can contribute to real operational systems, such as urban dashboards, digital twins, pollution warning services, mobility control platforms, and public-health communication tools. A model that performs well in a laboratory setting but lacks robustness, repeatability, or adaptability under noisy sensing conditions has limited value for urban governance. Accordingly, this study frames forecasting performance in terms that matter for city management: not only predictive error reduction, but also consistency across evaluation settings, resilience under repeated testing, and suitability for integration into broader smart-city infrastructures. This applied perspective positions the work within the growing body of research that seeks to connect artificial intelligence, urban analytics, and sustainability-oriented governance.

Another important motivation for the study is the need to treat environmental quality as a core pillar of smart-city research. Much of the literature on smart urbanism focuses on transportation optimisation, infrastructure intelligence, service digitisation, or energy management. While these topics are undoubtedly important, environmental conditions such as air quality are equally fundamental to the long-term sustainability and inclusiveness of urban development. Poor air quality undermines public health outcomes, increases social vulnerability, reduces the effectiveness of urban mobility systems, and places pressure on healthcare and environmental institutions. Consequently, improving the quality and reliability of pollution forecasting can support more proactive governance and more targeted policy interventions. In this sense, advances in predictive modelling are not merely computational improvements; they are enabling mechanisms for healthier, more adaptive, and more sustainable cities.

This manuscript is organised to emphasise that urban relevance throughout the modelling process. It presents a coherent account of the forecasting framework, including the empirical dataset, the preprocessing pipeline, the recurrent modelling design, the PSO-based optimisation strategy, and the reported evaluation results. Particular attention is given to those elements that are especially important for readers in urban development and smart-city research: the sensor-network context from which the data emerge, the operational logic of environmental forecasting, the role of data preparation in improving model reliability, the measurable benefits of hyperparameter tuning, and the robustness of the proposed framework under cross-validation and ablation analysis [6]. By structuring the study in this way, the manuscript seeks to demonstrate that methodological sophistication and urban applicability are not separate concerns, but mutually reinforcing dimensions of smart-city forecasting research.

Overall, the introduction establishes the central premise of the study: accurate and robust air-quality forecasting is a critical capability for smart and sustainable urban governance, and hybrid intelligent models offer a promising route toward achieving that capability. The proposed PSO–LSTM–RNN framework is therefore presented not only as a machine-learning contribution, but also as an urban decision-support tool with practical implications for environmental management, public health protection, and the strategic development of smart cities.

## URBAN RELEVANCE AND STUDY CONTRIBUTION

The practical significance of air-quality prediction in smart-city systems lies in its direct connection to urban services, policy responsiveness, and environmental governance. In contemporary cities, environmental intelligence is most valuable when it can be translated into timely operational action. Forecasts of pollutant concentrations are therefore not only descriptive analytical outputs, but also decision-support inputs that can shape how municipal authorities manage risk, allocate resources, and communicate with the public. When air-quality conditions can be anticipated with reasonable reliability, city administrations are better positioned to respond before pollution episodes escalate into broader public-health or mobility disruptions [4]. In this sense, predictive air-quality modelling contributes to the shift from reactive environmental management to proactive and data-driven urban governance.

The urban value of such forecasting is especially clear in service domains where environmental conditions interact with daily city functions. Forecasts can inform transport management during high-emission periods, support targeted communication for vulnerable populations, improve environmental surveillance practices, and guide energy and mobility policies designed to reduce population exposure [5]. A forecasting model that maintains consistent performance under realistic sensing conditions can therefore contribute to at least four major urban functions:

1. *Public-health protection*, by enabling earlier warnings during elevated pollution periods, supporting exposure-reduction strategies for children, older adults, and individuals with respiratory or cardiovascular vulnerability, and strengthening the timing of public advisories before hazardous conditions become severe [6];

2. *Traffic and mobility management*, by supporting more responsive control of congestion-related emissions through adaptive traffic regulation, route management, temporary mobility restrictions, and the coordination of cleaner transport operations during periods of anticipated environmental stress [7];

3. *Environmental compliance and urban monitoring*, by strengthening continuous surveillance of pollutant conditions, improving the ability of city agencies to detect emerging environmental risk patterns, and supporting regulatory oversight through more reliable predictive evidence [8];

4. *Sustainable service planning*, by embedding environmental intelligence into data-driven city operations, including the scheduling of sensitive public services, the coordination of outdoor municipal activities, and the design of smarter interventions that align urban functionality with sustainability goals [9].

Viewed from this perspective, the contribution of the study is not limited to forecasting accuracy alone. Its broader importance lies in showing how a hybrid intelligent model can serve as an enabling component of smart-city management systems. In practice, municipal platforms increasingly combine sensor networks, dashboards, planning tools, and public-information systems into integrated digital environments [10]. Within such environments, air-quality forecasting becomes most useful when it is reliable, interpretable at the operational level, and robust enough to support repeated use under noisy real-world conditions. The proposed framework speaks directly to this requirement by aiming to produce forecasts that are not only statistically stronger, but also more suitable for practical deployment in urban decision-support settings.

The study's substantive contribution is the integration of several complementary modelling ideas into a single forecasting framework. First, recurrent learning is employed to capture the temporal structure of pollutant variation. This is essential because air-quality dynamics unfold as sequential processes shaped by past concentrations, meteorological conditions, and recurring urban activity patterns. A recurrent structure is therefore well suited to modelling how present environmental states depend on previous observations [1, 2, 3]. Second, LSTM units are incorporated to address longer temporal dependencies that standard recurrent structures often struggle to preserve. In air-quality forecasting, these longer dependencies may reflect lagged meteorological effects, repeated daily or seasonal pollution cycles, and delayed interactions between emissions and atmospheric conditions. By retaining relevant information across longer time horizons, the LSTM component strengthens the model's ability to learn complex temporal behaviour that simpler recurrent forms may fail to capture effectively.

Third, PSO is introduced as an optimisation mechanism to improve the hybrid model's learning behaviour and overall forecasting efficiency. This is a particularly important contribution because deep-learning performance in environmental forecasting is often highly sensitive to hyperparameter selection and training configuration. Rather than relying on manually chosen settings or default parameter values, PSO provides a systematic search strategy for identifying improved solutions within the model space [4]. Its inclusion therefore enhances the framework not only by improving predictive performance, but also by making the model-design process more disciplined and adaptive. In smart-city applications, where forecasting systems must operate under heterogeneous data conditions and varying sensor quality, such optimisation is especially valuable for achieving stable and dependable results [5].

The source study further adds a curiosity-based motivation mechanism to refine the reconstructed hybrid model and strengthen its predictive capacity. Conceptually, this element is important because it introduces an additional layer of adaptive learning behaviour into the forecasting process. Rather than treating the model as a static predictive engine, the curiosity-based mechanism encourages improved exploration of the solution space and supports more effective internal adjustment during learning. Within the broader architecture, this feature complements the recurrent and optimisation components by helping the model respond more intelligently to complex data patterns. As a result, the overall framework is not simply a combination of PSO, RNN, and LSTM elements, but a more refined hybrid system designed to balance temporal learning, long-range dependency retention, parameter optimisation, and adaptive predictive refinement [11].

Taken together, these contributions position the study at the intersection of urban analytics, environmental forecasting, and smart-city systems research. Methodologically, it advances hybrid modelling for sequential environmental prediction. Practically, it demonstrates how improved forecasting design can support public-health protection, mobility governance, environmental surveillance, and sustainable service coordination. The study therefore contributes not only to the technical literature on intelligent forecasting models, but also to the applied literature on how digital intelligence can be embedded into more responsive, resilient, and sustainability-oriented urban management.

## MATERIALS AND METHODS

### *Study setting and dataset*

The empirical basis for the model is the Air Pollution Monitoring Network in Salt Lake City, Utah, USA. The dataset is compiled from *25 pollution sensors* and was requested from the University of Utah's linked research group. Each air-quality sensor transmits a data packet every 60 seconds when operating correctly. The sensing hardware includes an optical particle counter (Plan tower PMS3003), a temperature and humidity sensor (Texas Instruments HDC1080), and a gas sensor for oxidising and reducing gases (SGX SensorTech MiCS4514). For modelling, the measurements are aggregated to an *hourly* resolution, producing one row per device per hour. Hourly ozone values from the Utah Department of Air Quality's Hawthorne Monitoring Site are attached to the corresponding rows. The experimental window reported for model evaluation spans *2019-07-26 to 2021-05-14*.

The study centres on urban air-quality prediction, with the experimental discussion explicitly focusing on $PM_{2.5}$ forecasting while the broader monitoring network also incorporates multi-pollutant and environmental measurements. This design is well aligned with smart-city monitoring systems, where prediction typically depends on pollutant signals, environmental conditions, and station-level contextual observations rather than a single isolated time series.

### *Data preparation*

The modelling pipeline treats data preparation as a foundational step. The workflow includes feature engineering, missing-value handling, and outlier detection before model training. The source methodology describes data transformation as necessary for converting raw sensor inputs into a format suitable for machine learning, and it explicitly highlights three preparatory steps:

- *Normalisation and feature preparation*, to standardise heterogeneous sensor variables for recurrent learning;
- *Handling missing values*, including deletion, average imputation, speculative hot-deck imputation, and regression-based imputation;

- *Handling outliers*, with reference to boxplots, histograms, statistical analysis, interquartile range methods, and z-score filtering.

For urban sensing systems, these steps are not cosmetic. Sensor streams in real deployments are frequently noisy, incomplete, and operationally irregular, so model quality depends heavily on disciplined pre-processing.

*Model architecture*

The forecasting system combines RNN and LSTM components, then applies PSO for optimisation and a curiosity-based motivation mechanism for model reconstruction and learning enhancement. The source workflow proceeds in the sequence shown in Figure 1: (i) ingest the Salt Lake City and neighbouring-station datasets, (ii) pre-process the data, (iii) construct the LSTM-RNN model, (iv) apply the curiosity-based motivational layer to the reconstructed hybrid network, (v) forecast air quality, and (vi) evaluate the resulting predictions using standard regression metrics.

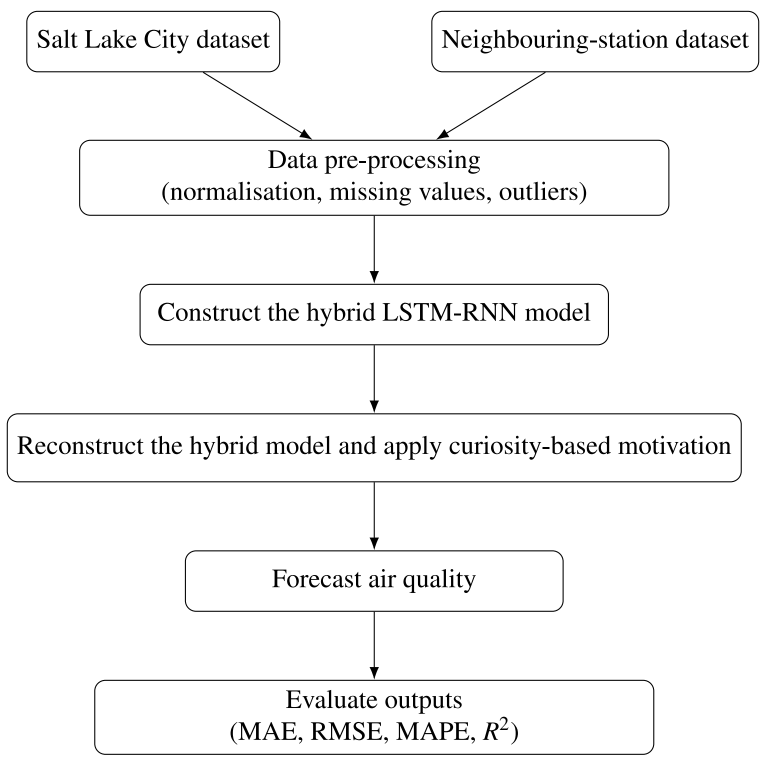


Figure 1: Cleaned reconstruction of the study workflow for smart-city air-quality forecasting.

Conceptually, the hybrid design is intended to combine the short-range sequential learning capacity of the RNN with the longer-memory structure of the LSTM. PSO is then used to tune model behaviour efficiently, while the curiosity-based component is described as strengthening the reconstructed LSTM-RNN so that the model better captures urban air-quality variation over time.

*Hyperparameter configuration*

The reported hyperparameter settings are shown in Table 1. These values define the operational configuration of the model and make the experimental pipeline reproducible in broad terms.

Table 1: Reported hyperparameter configuration for the hybrid PSO-LSTM-RNN model.

| *Method* | *Hyperparameter* | *Value* |
|---|---|---|
| Particle swarm optimisation | Swarm size | 30 |
| Particle swarm optimisation | Learning rate | 0.2 |
| Particle swarm optimisation | Inertia weight | 0.7 |
| LSTM | Number of layers | 2 |
| LSTM | Neurons per layer | 100 |
| LSTM | Dropout rate | 0.3 |
| LSTM | Learning rate | 0.001 |
| Training process | Batch size | 32 |
| Training process | Epochs | 10–50 |
| Recurrent neural network | Number of layers | 2 |
| Recurrent neural network | Neurons per layer | 100 |
| Recurrent neural network | Dropout rate | 0.3 |
| Recurrent neural network | Learning rate | 0.001 |

These settings reflect a moderate-size recurrent architecture designed to balance temporal expressiveness against overfitting risk. In practical smart-city deployments, this balance matters because urban sensor datasets often exhibit non-stationarity, missingness, and periodic shifts in environmental conditions.

*Evaluation protocol*

The study evaluates model performance using four standard regression metrics:

- Root Mean Square Error (RMSE),
- Mean Absolute Error (MAE),
- Mean Absolute Percentage Error (MAPE),
- Coefficient of Determination ($R^2$ score).

The experiments include direct benchmark comparison, analysis with and without hyperparameter tuning, 10-fold cross-validation, and ablation testing of the main architectural components.

## RESULTS

*Benchmark comparison*

In the core benchmark comparison, the proposed hybrid model is reported alongside Transformer, GRU, LSTM + RNN + ant colony optimisation (ACO), and LSTM + RNN + genetic algorithm (GA) alternatives. The clearest performance advantages appear in *MAE* and $R^2$. The proposed model records an MAE of *0.0082* and an $R^2$ score of *0.1227*, compared with 0.0237 and 0.0591 for Transformer, 0.0197 and 0.0784 for GRU, 0.0185 and 0.0874 for LSTM + RNN + ACO, and 0.0165 and 0.0890 for LSTM + RNN + GA. The published benchmark table also reports RMSE values, but the narrative interpretation of the benchmark is more consistently supported by the MAE and $R^2$ results than by the reported RMSE line. For that reason, the present manuscript emphasises the most internally consistent indicators when interpreting model superiority.

The study also states that the proposed model was compared against GBTR, an existing LSTM, and an SVM-based regression model in the broader experimental programme, reinforcing the conclusion that the hybrid recurrent design offers a stronger predictive structure for urban air-quality applications.

*Effect of hyperparameter tuning*

The effect of hyperparameter tuning is reported clearly and is one of the strongest empirical sections of the study. Table 2 shows that tuning improves every model considered, with the proposed hybrid model producing the best tuned performance overall. With tuning enabled, the proposed model achieves *RMSE 0.012*, *MAE 0.009*, *MAPE 3.5%*, and $R^2$ *0.95*. Without tuning, the same model falls to RMSE 0.018, MAE 0.013, MAPE 5.0%, and $R^2$ 0.88.

Table 2: Experimental results with and without hyperparameter tuning.

| *Model* | *Tuning* | *RMSE (with)* | *RMSE (without)* | *MAE (with)* | *MAE (without)* | *MAPE (with)* | *MAPE (without)* | $R^2$ *(with)* | $R^2$ *(without)* |
|---|---|---|---|---|---|---|---|---|---|
| Proposed hybrid model | Yes | 0.012 | 0.025 | 0.009 | 0.018 | 3.5% | 7.2% | 0.95 | 0.85 |
| LSTM + RNN + ACO | Yes | 0.015 | 0.030 | 0.011 | 0.022 | 4.2% | 8.5% | 0.92 | 0.80 |
| LSTM + RNN + GA | Yes | 0.014 | 0.028 | 0.010 | 0.020 | 4.0% | 8.0% | 0.93 | 0.82 |
| Proposed hybrid model | No | 0.018 | 0.035 | 0.013 | 0.026 | 5.0% | 9.8% | 0.88 | 0.75 |
| LSTM + RNN + ACO | No | 0.022 | 0.040 | 0.016 | 0.032 | 6.2% | 11.2% | 0.82 | 0.70 |
| LSTM + RNN + GA | No | 0.020 | 0.038 | 0.015 | 0.030 | 5.8% | 10.5% | 0.84 | 0.72 |

This section has direct applied value for urban analytics. It shows that optimisation is not peripheral to model quality: thoughtful tuning substantially reduces error and improves explanatory power, which is essential when forecasts may shape operational decisions in city systems.

*Cross-validation robustness*

The 10-fold cross-validation analysis provides the clearest evidence of model robustness. As shown in Table 3, the proposed hybrid model achieves the strongest average performance across folds, with *Avg. RMSE 0.013*, *Avg. MAE 0.010*, *Avg. MAPE 3.8%*, and *Avg.* $R^2$ *0.94*. By comparison, the LSTM + RNN + ACO model records 0.016, 0.012, 4.6%, and 0.90, while the LSTM + RNN + GA model records 0.018, 0.013, 5.2%, and 0.87.

For smart-city decision support, this result matters as much as headline accuracy. Cross-validation demonstrates that the model does not depend on a single favourable split and instead retains stable predictive quality across varied partitions of the urban air-quality dataset.

*Ablation analysis*

The ablation analysis isolates the contribution of the main architectural components. Table 4 shows a stepwise improvement as the model moves from a basic LSTM to the full hybrid architecture. MAE decreases from *0.0584* for the basic LSTM to *0.0489* after adding RNN, then to *0.0189* after introducing PSO, and finally to *0.0082* once curiosity-based motivation is integrated. Over the same sequence, MAPE falls from 0.0478 to 0.0184 and $R^2$ rises from 0.067 to 0.1227.

Table 3: 10-fold cross-validation results (reported representative folds and averages).

| *Model* | *Fold* | *RMSE* | *MAE* | *MAPE* | $R^2$ |
|---|---|---|---|---|---|
| Proposed hybrid model | 1 | 0.012 | 0.009 | 3.5% | 0.95 |
| Proposed hybrid model | 2 | 0.014 | 0.010 | 3.8% | 0.94 |
| Proposed hybrid model | 10 | 0.015 | 0.011 | 4.0% | 0.93 |
| Proposed hybrid model | Avg. | 0.013 | 0.010 | 3.8% | 0.94 |
| LSTM + RNN + ACO | 1 | 0.015 | 0.011 | 4.2% | 0.92 |
| LSTM + RNN + ACO | 2 | 0.016 | 0.012 | 4.5% | 0.91 |
| LSTM + RNN + ACO | 10 | 0.018 | 0.013 | 5.0% | 0.88 |
| LSTM + RNN + ACO | Avg. | 0.016 | 0.012 | 4.6% | 0.90 |
| LSTM + RNN + GA | 1 | 0.014 | 0.010 | 4.0% | 0.93 |
| LSTM + RNN + GA | 2 | 0.017 | 0.013 | 4.8% | 0.89 |
| LSTM + RNN + GA | 10 | 0.020 | 0.015 | 5.8% | 0.84 |
| LSTM + RNN + GA | Avg. | 0.018 | 0.013 | 5.2% | 0.87 |

Table 4: Ablation analysis of the hybrid forecasting framework.

| *Exp.* | *Model components altered* | *MAE* | *MAPE* | $R^2$ |
|---|---|---|---|---|
| 1 | Basic LSTM | 0.0584 | 0.0478 | 0.067 |
| 2 | LSTM with RNN | 0.0489 | 0.0397 | 0.078 |
| 3 | LSTM with RNN + PSO | 0.0189 | 0.0258 | 0.081 |
| 4 | LSTM with RNN + PSO integrated with curiosity-based motivation | 0.0082 | 0.0184 | 0.1227 |

This is an especially important finding for urban-development research because it shows that performance gains are cumulative rather than incidental. The model's strongest performance emerges only when the recurrent architecture, the optimisation routine, and the motivational enhancement are jointly deployed.

## DISCUSSION

The empirical findings of this study support three closely connected conclusions that are important both from a forecasting perspective and from the standpoint of smart-city decision support. First, the results clearly indicate that hybridisation improves predictive performance in a meaningful way [6]. The progression from the baseline LSTM model to the RNN–LSTM configuration demonstrates that architectural enrichment alone contributes to lower forecasting error, suggesting that the additional recurrent structure enables the model to capture temporal dependencies more effectively. However, the strongest gains emerge after the integration of particle swarm optimisation (PSO), which further refines the model by identifying more suitable parameter configurations than manual or default settings. This pattern indicates that the performance advantage of the proposed framework is not due to a single design choice, but rather to the interaction between deep sequential learning and systematic optimisation.

Second, the results show that hyperparameter tuning is not merely a supplementary step, but a central requirement for robust forecasting [7]. Across all evaluated model variants, the tuned configurations outperform their untuned counterparts by a substantial margin. This consistent improvement suggests that model behaviour in urban sensing environments is highly sensitive to parameter selection, including network structure, learning dynamics, and optimisation settings. In practical terms, this means that forecasting systems built for urban applications cannot rely on generic model settings if they are expected to provide dependable outputs. Instead, optimisation must be treated as an integral stage of model development, particularly when the forecasting task involves heterogeneous environmental variables, nonstationary time series, and sensor-generated noise [8]. The evidence therefore reinforces the view that predictive accuracy in smart-city systems depends not only on the choice of algorithm, but also on the discipline with which that algorithm is calibrated.

Third, the proposed PSO–RNN–LSTM model demonstrates stability under repeated evaluation, as reflected in the 10-fold cross-validation results. This is an important outcome because strong performance in a single train–test split may sometimes reflect favourable data partitioning rather than genuine generalisability. By maintaining superior or near-consistent results across multiple folds, the full model shows that its forecasting capability is not narrowly tied to one particular sample arrangement. Such robustness is especially valuable in urban analytics, where environmental data streams may vary across locations, seasons, and operating conditions. A forecasting model intended for deployment in real municipal systems must therefore be evaluated not only on point accuracy, but also on repeatability and resilience under changing data subsets. The cross-validation findings provide evidence that the proposed framework meets this broader standard.

From the perspective of urban development and smart-city management, these results have direct practical significance because public-sector systems require forecasting tools that are reliable, stable, and operationally useful rather than merely technically interesting [9, 10, 11]. Municipal authorities increasingly depend on data-driven platforms to monitor environmental conditions, coordinate interventions, and communicate risks to the public. In such settings, a model that can ingest air-quality and environmental sensor data, tolerate the imperfections of real-world measurements, and still deliver repeatable predictive performance becomes highly valuable. The proposed framework is therefore well aligned with deployment scenarios such as smart urban dashboards, environmental early-warning systems, digital twin infrastructures, and integrated planning platforms [12]. Its value lies not only in improved forecast accuracy, but also in its potential to strengthen trust in automated decision-support tools used by city managers and policy institutions.

The broader sustainability implications of the study are also noteworthy. Discussions of smart-city innovation often emphasize transport optimisation, infrastructure automation, or platform intelligence, yet environmental quality remains a foundational dimension of sustainable urbanism. Air pollution directly affects public health, urban liveability, and the long-term resilience of metropolitan regions. A forecasting model that improves the timeliness and reliability of pollution prediction can therefore support more responsive and evidence-based interventions. These may include temporary mobility restrictions, adaptive traffic regulation, targeted public-health advisories, school and workplace exposure alerts, and environmentally informed scheduling of public services [13]. In this sense, the contribution of the study extends beyond computational improvement: it provides a methodological basis for translating environmental data into proactive urban governance.

At a conceptual level, the study also highlights the value of combining intelligent optimisation with deep learning in urban informatics. Urban systems are inherently dynamic, multidimensional, and uncertain, which makes them poorly suited to forecasting approaches that depend on static assumptions or limited parameter exploration. By integrating PSO with a recurrent deep-learning architecture, the proposed framework responds to this complexity in a way that is both adaptive and scalable. This is particularly important for smart-city contexts, where forecasting tools must often operate across different sensor networks, varying temporal resolutions, and evolving environmental patterns. The results suggest that hybrid intelligent models may offer a productive direction for future urban forecasting research, especially in domains where data quality is uneven and operational reliability is essential.

At the same time, the findings should be interpreted with appropriate caution. Although the proposed model achieves strong results on the available dataset and remains stable across cross-validation, real-world deployment may still require additional testing under alternative geographic, climatic, and infrastructural conditions. Urban air-quality patterns are shaped by local traffic structures, industrial activity, meteorological variation, and seasonal behaviours, all of which may influence model transferability. Future studies could therefore extend the present framework by evaluating it across multiple cities, incorporating additional exogenous variables, or testing real-time streaming implementations in operational smart-city environments. Such extensions would help determine the extent to which the observed performance gains can generalise beyond the current experimental setting.

Overall, the discussion confirms that the proposed PSO–RNN–LSTM framework is not simply a marginal improvement over conventional forecasting approaches, but a practically relevant advancement for environmental prediction in smart cities. The combination of hybrid architecture, systematic hyperparameter optimisation, and stable validation performance provides a strong basis for its use in urban sensing applications. More broadly, the study demonstrates how improved forecasting methodology can contribute to smarter governance, more responsive environmental management, and a more sustainable urban future.

## LIMITATIONS

Two limitations should be recognised. First, the empirical setting is a single metropolitan monitoring network in Salt Lake City. The results are therefore best interpreted as evidence of model effectiveness in a defined urban context rather than proof of universal transferability. Second, some portions of the original benchmark reporting are more interpretable than others. In particular, the MAE, tuning, cross-validation, and ablation sections form the most internally consistent basis for evaluation and are therefore given greatest interpretive weight in this manuscript.

These limitations do not diminish the value of the study. Instead, they identify the next logical steps for research in the urban-development and smart-city field: multi-city validation, transferability testing across monitoring regimes, and closer integration of forecasting outputs with operational city decisions.

## CONCLUSION

This manuscript presents a coherent smart-city interpretation of a hybrid PSO-LSTM-RNN framework for urban air-quality prediction. Using a Salt Lake City sensor network composed of 25 monitoring units and an hourly aggregated dataset, the study combines recurrent sequence modelling, PSO-based hyperparameter optimisation, and curiosity-based motivation to improve prediction quality. The strongest reported evidence comes from the tuned experiments, the cross-validation analysis, and the ablation study, all of which show that the full hybrid design outperforms simpler alternatives and retains stable performance across evaluation settings.

In substantive terms, the contribution is clear: urban air-quality prediction can be improved when forecasting systems combine careful data preparation, recurrent temporal modelling, optimisation-driven parameter selection, and layered architectural refinement. For the literature on urban development and smart cities, this positions the model as a practically relevant framework for environmentally intelligent governance, healthier urban living, and more responsive city operations.

## REFERENCES

[1] Dalal, S., Lilhore, U.K., Faujdar, N., Samiya, S., Jaglan, V., Alroobaea, R., Shaheen, M., and Ahmad, F.: Optimising air quality prediction in smart cities with hybrid particle swarm optimization-long-short term memory-recurrent neural network model. *IET Smart Cities 6*(3), 156–179 (2024).

[2] Air Pollution Monitoring Dataset: University of Utah AirU Pollution Monitoring Network, Salt Lake City, UT, 2019-07-26 to 2021-05-14. IEEE DataPort dataset record (2023).

[3] Hochreiter, S., and Schmidhuber, J.: Long short-term memory. *Neural Computation 9*(8), 1735–1780 (1997).

[4] Kennedy, J., and Eberhart, R.: Particle swarm optimization. In: *Proceedings of the IEEE International Conference on Neural Networks*, pp. 1942–1948 (1995).

[5] Friedman, J.H.: Greedy function approximation: a gradient boosting machine. *Annals of Statistics 29*(5), 1189–1232 (2001).

[6] Malche, T., Maheshwary, P., and Kumar, R.: Environmental monitoring system for smart city based on secure Internet of Things architecture. *Wireless Personal Communications 107*(4), 2143–2172 (2019).

[7] Li, S., & Li, J. (2017). Condition monitoring and diagnosis of power equipment: review and prospective. High Voltage, 2(2), 82-91.

[8] Zou, Q., Zhang, T., & Liu, W. (2021). A fire risk assessment method based on the combination of quantified safety checklist and structure entropy weight for shopping malls. Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability, 235(4), 610-626.

[9] Muhammad, H., Jenuwa, N., & Jaafar, A. S. (2023). Fire safety system at cinema in Klang Valley. In e-Proceeding of 6th Undergraduate Seminar on Built Environment and Technology (USBET) 2023 (pp. 864-873).

[10] Wang, Y., Chen, G., Yang, F., Luo, Z., Yuan, B., Chen, X., & Wang, L. (2022). Serendipity discovery of fire early warning function of chitosan film. Carbohydrate polymers, 277, 118884.

[11] Xu, H., Li, Y., Huang, N. J., Yu, Z. R., Wang, P. H., Zhang, Z. H., ... & Tang, L. C. (2019). Temperature-triggered sensitive resistance transition of graphene oxide wide-ribbons wrapped sponge for fire ultrafast detecting and early warning. Journal of hazardous materials, 363, 286-294.

[12] Li, X., Vázquez-López, A., Sánchez del Río Sáez, J., & Wang, D. Y. (2022). Recent advances on early-stage fire-warning systems: mechanism, performance, and perspective. Nano-micro letters, 14(1), 197.

[13] Dian, S., Cheng, P., Ye, Q., Wu, J., Luo, R., Wang, C., ... & Gong, X. (2019). Integrating wildfires propagation prediction into early warning of electrical transmission line outages. IEEE Access, 7, 27586-27603.

Orhun, D. Dokuz Eylül Üniversitesi, Izmir, Turkey

Manuscript Published; 11 October 2025.