Predicting Air Quality in İzmir Using Artificial Intelligence and IoT

Kübra Öztürk[1], Zuhal Can[2]

ORCID NO: 0009-0003-4368-9274, 0000-0002-6801-1334

[1] Eskişehir Osmangazi University, Faculty of Engineering and Architecture, Department of Computer Engineering, Eskişehir, Türkiye
[2] Eskişehir Osmangazi University, Faculty of Engineering and Architecture, Department of Computer Engineering, Eskişehir, Türkiye

Air pollution is a significant concern in Izmir, the third-largest city in Türkiye, with adverse impacts on public health and urban quality of life. Leveraging Internet of Things (IoT) monitoring, this study forecasts PM_{10} and SO_2 using machine learning, deep learning, and time-series models over 1996-2024. Our main contribution is a controlled, city-scale evaluation that isolates the role of extreme-value handling by creating three preprocessing variants for each pollutant (keep, cap, remove) and benchmarking four model families (SARIMA, SVR, LSTM, xLSTM) under a leakage-free, time-ordered train/validation/test protocol. To our knowledge, this is the first Izmir study that quantifies how peak treatment interacts with model choice to affect forecast accuracy, and that reports results in a way that directly supports operational planning. Results show that for PM_{10}, SARIMA often attains the lowest errors when peaks are capped or removed, while xLSTM provides competitive accuracy across all strategies. For SO_2, xLSTM consistently outperforms traditional methods under capped and kept conditions, yielding lower errors and higher explanatory power. By combining IoT-based monitoring with a rigorous modeling method and explicit sensitivity to peak processing, the study offers actionable guidance for air-quality management, informing traffic control, industrial regulation, and urban design decisions in Izmir.

Received: 20.08.2024
Accepted: 14.09.2025

Corresponding Author:
zcan@ogu.edu.tr

Öztürk, K. & Can, Z. (2024). Predicting air quality in İzmir using artificial intelligence and IoT, *JCoDe: Journal of Computational Design*, 6(2), 341-364. https://doi.org/10.53710/jcode.1536480

Keywords: Air Quality, Artificial Intelligence, IoT, Weather Forecasting, xLSTM.

Yapay Zeka ve IoT Kullanarak İzmir'deki Hava Kalitesinin Tahmini

Kübra Öztürk[1], Zuhal Can[2]

ORCID NO: 0009-0003-4368-9274, 0000-0002-6801-1334

[1] Eskişehir Osmangazi Üniversitesi, Mühendislik ve Mimarlık Fakültesi, Bilgisayar Mühendisliği Bölümü, Eskişehir, Türkiye
[2] Eskişehir Osmangazi Üniversitesi, Mühendislik ve Mimarlık Fakültesi, Bilgisayar Mühendisliği Bölümü, Eskişehir, Türkiye

Türkiye'nin üçüncü büyük şehri İzmir'de hava kirliliği, halk sağlığı ve kentsel yaşam kalitesi üzerinde olumsuz etkileri olan önemli bir endişe kaynağıdır. Nesnelerin İnterneti (IoT) izlemeyi kullanan bu çalışma, makine öğrenimi, derin öğrenme ve zaman serisi modelleri kullanarak 1996-2024 yılları arasında PM_{10} ve SO_2'yi tahmin etmektedir. Ana katkımız, her bir kirletici için üç ön işleme varyantı oluşturarak (tut, sınırla, kaldır) ve sızıntısız, zaman sıralı bir eğitim/doğrulama/test protokolü altında dört model ailesini (SARIMA, SVR, LSTM, xLSTM) kıyaslayarak aşırı değer işleme rolünü izole eden kontrollü, şehir ölçeğinde bir değerlendirmedir. Bilgimize göre, bu, tepe işleminin tahmin doğruluğunu etkilemek için model seçimiyle nasıl etkileşime girdiğini ölçen ve sonuçları doğrudan operasyonel planlamayı destekleyecek şekilde raporlayan ilk İzmir çalışmasıdır. Sonuçlar, PM_{10} için SARIMA'nın, tepe noktaları sınırlandırıldığında veya kaldırıldığında genellikle en düşük hatalara ulaştığını, xLSTM'nin ise tüm stratejilerde rekabetçi doğruluk sağladığını göstermektedir. SO_2 için ise xLSTM, sınırlandırılmış ve sabitlenmiş koşullar altında geleneksel yöntemlerden sürekli olarak daha iyi performans göstererek daha düşük hatalar ve daha yüksek açıklama gücü sağlamaktadır. IoT tabanlı izlemeyi titiz bir modelleme metodu ve tepe noktası işlemeye açık duyarlılıkla birleştiren çalışma, İzmir'de hava kalitesi yönetimi, trafik kontrolü, endüstriyel düzenlemeler ve kentsel tasarım kararlarına bilgi sağlamak için uygulanabilir rehberlik sunmaktadır.

Teslim Tarihi: 20.08.2024
Kabul Tarihi: 14.09.2025

Sorumlu Yazar:
zcan@ogu.edu.tr

Öztürk, K. & Can, Z. (2024). Yapay zeka ve IoT kullanarak İzmir'deki hava kalitesinin tahmini. *JCoDe: Journal of Computational Design*, 6(2), 341-364. https://doi.org/10.53710/jcode.1536480

Keywords: Hava Kalitesi, Yapay Zeka, Nesnelerin İnterneti, Hava Tahmini, xLSTM.

1. INTRODUCTION

Air pollution is a critical issue in urban environments, significantly affecting human health and the ecosystem. Izmir, the third-largest city in Türkiye, faces recurring air quality challenges, primarily due to industrial activities, dense traffic, and seasonal meteorological patterns that can exacerbate pollutant accumulation. Addressing these challenges is essential to safeguard public health and ensure a sustainable environment.

This study develops a city-scale forecasting workflow for İzmir using a combination of machine learning, deep learning, and time series analysis methods. Unlike prior works that focused solely on raw historical values, this study systematically incorporates extreme-value handling into the predictive workflow, creating three separate versions of each pollutant series, keep, cap, and remove, to assess how different preprocessing strategies influence model performance. By analyzing historical air quality data, the study seeks to provide valuable insights that can help public health authorities and policymakers mitigate potential risks. The findings can inform strategies such as traffic management and industrial regulation, improving air quality, and ensuring compliance with environmental standards.

The integration of IoT and artificial intelligence (AI) into air quality forecasting offers significant advantages in the design and implementation of environmental management systems. IoT-based sensor networks enable real-time, high-resolution monitoring of pollutant levels, while AI-driven models can process this continuous data stream to produce accurate and timely forecasts. Embedding these predictive capabilities into urban planning workflows allows decision-makers to proactively address emerging air quality issues, such as by adjusting traffic flows, optimizing industrial operations, or issuing early health advisories, before critical thresholds are exceeded.

In addition, coupling IoT infrastructure with AI forecasting models facilitates adaptive policy-making. For example, predictive alerts generated by AI can be integrated with smart city control systems to automatically trigger mitigation measures, such as activating air filtration in public spaces or modifying public transportation schedules. This real-time, feedback-driven approach not only improves response

speed but also helps optimize resource allocation, ultimately contributing to healthier and more sustainable urban environments.

Beyond short-term interventions, accurate air quality forecasts can also shape long-term urban design and planning decisions. Predictive insights can guide optimal building placement to minimize exposure in high-pollution corridors, inform the strategic allocation of green spaces to enhance natural air filtration, and support the development of sustainable architecture that incorporates passive ventilation and pollutant-shielding designs. By integrating forecasting data into master planning processes, cities can proactively design neighborhoods that balance density with environmental health considerations.

Furthermore, urban planners can leverage AI-driven air quality projections to evaluate the long-term impact of infrastructure projects, ensuring that new developments do not exacerbate pollution hotspots. This integration of environmental forecasting into architectural and landscape planning supports the creation of resilient, climate-adaptive, and livable urban spaces that protect public health while advancing sustainability goals.

The dataset utilized in this study includes air quality measurements for PM_{10} and SO_2, collected from 1996 to 2024 through IoT sensors. These measurements, sourced from the National Smart City Open Data Platform (Ulusav) and the Izmir Metropolitan Municipality, provide a comprehensive view of the air quality across multiple monitoring stations in Izmir. Monthly observations were aggregated into pollutant-specific time series for each preprocessing strategy, with temporal splits applied strictly in chronological order to avoid data leakage.

To predict future air pollution levels, the study employs four models: Support Vector Regression (SVR), Seasonal Autoregressive Integrated Moving Average (SARIMA), Long Short-Term Memory (LSTM), and Extended Long-Term Memory (xLSTM). The xLSTM architecture enhances LSTM by incorporating extended memory connections and adaptive gating mechanisms, improving its capacity to capture long-term dependencies in air quality trends. These models are evaluated using key performance metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and R^2 Score.

The results reveal that model performance varies by pollutant and preprocessing strategy. For PM_{10}, SARIMA often achieves the lowest errors under capped and removed peaks, whereas xLSTM delivers comparable accuracy with greater adaptability across strategies. For SO_2, xLSTM consistently outperforms traditional methods, particularly under capped and kept conditions, achieving lower error metrics and higher explanatory power. These findings underscore the potential of advanced machine learning models, especially xLSTM, in forecasting air quality and contributing to better environmental management in urban settings.

1.1. Contributions to Planning and Architectural Design

Beyond forecasting, we explicitly translate model outputs into design- and planning-ready artifacts and workflows. (i) Design-relevant framing: By testing three peak-handling policies (keep/cap/remove) across four model families, we treat outlier processing as a planning lever, showing how alternative regulatory and operational choices for peaks change forecast fidelity and, consequently, recommended interventions. (ii) Actionable products: We derive monthly risk bands with uncertainty envelopes and convert them into trigger-action rules for practice (e.g., adjusting street-canyon permeability and block porosity, temporally reallocating traffic/freight, or prioritizing green/blue infrastructure where dispersion is favored). (iii) Multi-scale guidance: We link forecast bands to block- and building-scale decisions, such as ventilation and filtration scheduling, façade/atrium strategies along high-exposure corridors, and the micro-siting of sensitive receptors (schools/clinics) to minimize long-term exposure. (iv) Toolchain integration and reproducibility: We provide a leakage-free, city-scale workflow that municipalities can embed into GIS dashboards and scenario planning to test design options under different peak policies and seasonal regimes. (v) Sustainability integration: Aligning ventilation/filtration duty cycles, traffic operations, and the siting and maintenance of green/blue infrastructure with forecasted exposure patterns delivers equal-or-greater air-quality benefits at lower energy and material intensity than uniform, calendar-based interventions; it also yields co-benefits such as reduced greenhouse-gas emissions (via demand-controlled ventilation and freight retiming) and mitigates maladaptation risks (e.g., over-canopying in narrow street canyons). Equity is addressed by prioritizing sensitive receptors using banded

forecasts. The artifacts align with municipal monitoring and verification cycles, supporting cost-effectiveness tracking (e.g., kWh saved, avoided vehicle-kilometers, filter-lifetime extension, and person-exposure-hours reduced).

2. LITERATURE REVIEW

In recent years, the integration of Internet-of-Things (IoT) technology into environmental-monitoring systems has gained significant attention. Researchers have developed a variety of IoT-based platforms to capture weather conditions and air-quality dynamics with high temporal resolution. Bernardes et al. (Bernardes et al., 2023) proposed a low-cost automated weather station for disaster monitoring, while Woo et al. (Woo et al., 2023) released WeatherChimes, an open-hardware package giving near-real-time access to on-site sensor data. Bolla et al. (Bolla et al., 2022) highlighted the importance of a robust IoT framework for efficient data collection, and Elbasi et al. (Elbasi et al., 2023) reviewed AI-enabled sensing networks for crops and leak detection. Singh et al. (Singh et al., 2022) demonstrated an intelligent irrigation system that responds adaptively to soil and weather inputs.

IoT innovation continues to expand. Ioannou et al. (Ioannou et al., 2021) combined wireless technologies and ARM-SoC to create a multi-purpose device, whereas Mabrouki et al. (Mabrouki et al., 2021) delivered a real-time climate platform. Ambildhuke and Banik (Ambildhuke & Banik, 2022) applied deep learning to a portable precipitation-forecasting device, and Karvelis et al. (Karvelis et al., 2020) offered an on-board solution for real-time weather prediction. Efforts to optimise energy consumption in remote stations include Leelavinodhan et al. (Leelavinodhan et al., 2021) and Mehmood et al. (Mehmood et al., 2023), while Lu et al. (Lu et al., 2023) introduced ThunderLock, a dual-microphone lightning localisation system.

Complementary work has explored data-fusion and hybrid-modelling strategies. Nie et al. (Nie et al., 2012) combined ARIMA and SVM for short-term load forecasting, Mohapatra and Subudhi (Mohapatra & Subudhi, 2022) demonstrated embedded data-collection modules, and Kaya et al. (Kaya et al., 2023) benchmarked anomaly-detection algorithms. Suresh et al. (Suresh et al., 2022) attained high F1-scores for IoT-driven irrigation. Jamil et al. (Jamil et al., 2021) presented a big-

data architecture for streaming IoT flows, Albuali et al. (Albuali et al., 2023) integrated machine-learning with automated weather stations, and Tsalikidis et al. (Tsalikidis et al., 2024) compared classical and deep models for traffic prediction under varying meteorology.

Deep-learning approaches now dominate many environmental-forecasting tasks. Yang et al. (Yang et al., 2022) fused dense vehicular IoT data with neural networks for city-scale PM-prediction; Yu et al. (Yu et al., 2021) combined LSTM with long-term meteorological archives for temperature forecasts. Roy (Roy, 2020) showed that hybrid CNN + LSTM structures boost horizon length, while Elsaraiti and Merabet (Elsaraiti & Merabet, 2021), Geng et al. (Geng et al., 2020), Atali et al. (Atali et al., 2022) and Aydin et al. (Aydin et al., 2021) documented performance gains for wind-speed and PM_{10} series.

Crucially, recent scholarship has begun to translate these sensing-and-forecasting advances into architectural and urban-planning practice. Bansal and Quan (Bansal & Quan, 2024) employed explainable machine learning on 1.4 million parcels to show how block geometry modulates heat-island intensity across Seoul, offering planners data-driven guidance on reflectivity and porosity. Synthesis papers on green infrastructure further warn that vegetative interventions may trap pollutants in narrow canyons if used indiscriminately, whereas carefully sited barriers and corridors can reduce near-road exposure (Vos et al., 2013), (Abhijith et al., 2017). These architecture- and planning-oriented studies bridge measurement, simulation, and design decision-making, providing pathways to couple forecasting outputs with street-canyon permeability, green/blue-infrastructure siting, and building-scale ventilation/filtration schedules within GIS-supported scenario workflows.

Recent work at the intersection of AI, urban form and sustainability further demonstrates design uptake. Patel et al. (Patel et al., 2023) linked deep-learning-derived morphology metrics to dispersion models. Wu et al. (Wu et al., 2023) applied GIS-CFD coupling for canyon ventilation. Multi-scale land-use regression by Wang et al. (Wang et al., 2024) connects neighbourhood density, greenery and $PM_{2.5}$ gradients, while Venter et al. (Venter et al., 2024) caution that indiscriminate greening can intensify street-level PM under calm conditions.

These studies collectively demonstrate not only the technical maturation of IoT-centric environmental monitoring and predictive modelling, but an emerging evidence base that positions air-quality data as a design lever in architecture and urban planning, linking sensor insight to façade tactics, corridor porosity, traffic orchestration and green-infrastructure siting in pursuit of healthier, more sustainable cities.

3. METHOD

The aim of this study is to carry out a local study with the air quality data of Izmir province, the third largest city of Turkey, and to make predictions about air quality for the future. The potential impact is significant, as it can alert health institutions and public-health authorities to possible risks, inform policies such as traffic management or regulation of industrial activities, and support industrial facilities or commercial enterprises in ensuring compliance with environmental regulations by reducing emissions under identified meteorological conditions.

3.1. Dataset

This study uses monthly air-quality observations for İzmir, Türkiye, spanning 1996-2024, compiled from the National Smart City Open Data Platform and the İzmir Metropolitan Municipality (İzmir Büyükşehir Belediyesi, 2024). The raw file may appear in a “long” layout, with columns equivalent to a date, a pollutant identifier, and a numeric measurement, or in a “wide” layout with pollutant names as columns (e.g., PM_{10}, SO_2). A schema-inference routine first normalizes header strings and detects the date and value fields. For long-format data, the pollutant identifier is mapped to pollutant names using {1: PM_{10}, 2: SO_2}, and the numeric measurement column is used directly as the measurement; for wide-format data, the columns labelled (or normalized to) PM_{10} and SO_2 are taken directly. Dates are parsed to a monthly index at month start (MS), rows with invalid dates are discarded, and the series are sorted chronologically. Where multiple stations or regions exist, measurements are aggregated by the monthly mean to obtain a single univariate series per pollutant. The resulting PM_{10} and SO_2 series are resampled onto a regular monthly grid; if a series contains fewer than 60 monthly points, it is excluded to preserve a meaningful temporal split.

To systematically assess the impact of extreme values, peak handling was incorporated as a controlled preprocessing factor rather than a one-off cleaning step. For each pollutant, three parallel versions of the monthly time series were generated. The first version ("keep") retained the original values without modification. The second version ("cap") applied winsorization based on Tukey's method (Tukey & others, 1977), which identifies potential outliers as observations lying more than 1.5 times the interquartile range (IQR) below the first quartile (Q_1) or above the third quartile (Q_3). The 1.5×IQR threshold is a conventional choice in statistical analysis because it offers a balanced compromise, sufficiently wide to avoid labeling normal variability as an outlier, yet narrow enough to capture anomalous spikes. In this version, values exceeding the calculated bounds were clipped to the nearest limit. The third version ("remove") used the same 1.5×IQR rule to detect outliers but removed these observations entirely from the series. Each version proceeded independently through the full modeling method, enabling a systematic evaluation of how alternative outlier treatments affect forecasting accuracy.

Each version proceeds independently through training and evaluation. For every (pollutant, peak-strategy) pair, the data are split strictly by time into training (70% earliest observations), validation (next 15%), and test (final 15%) with no shuffling. Summary statistics (sample size, start/end dates, mean, standard deviation) are computed per pollutant and retained for transparency.

3.2. Models

Four forecasting models were employed to evaluate air quality prediction performance: SARIMA, SVR, LSTM, and the proposed xLSTM. These models represent different methodological paradigms, ranging from traditional statistical time-series analysis to modern deep learning architectures.

The Seasonal AutoRegressive Integrated Moving Average (SARIMA) model extends the classical ARIMA framework by explicitly incorporating seasonal components, enabling it to capture recurring temporal patterns alongside non-seasonal autoregressive and moving average structures. The model is denoted as $\mathrm{SARIMA}(\mathrm{p},\mathrm{d},\mathrm{q}) \times (\mathrm{P},\mathrm{D},\mathrm{Q})_{\mathrm{s}}$ where p, d, and q represent the non-seasonal autoregressive

order, differencing order, and moving average order, respectively, while P, D, and Q represents the seasonal counterparts over a period s. Optimal hyperparameters for SARIMA were determined using a grid search on the training and validation sets, minimizing error metrics to ensure robust seasonal pattern extraction.

Support Vector Regression (SVR) was employed as a kernel-based machine learning model capable of capturing nonlinear relationships between predictors and the target variable. SVR maps input data into a high-dimensional feature space through a kernel function. In this study, the radial basis function (RBF) kernel, allowing the algorithm to fit a regression function that tolerates deviations within an ε-insensitive margin while penalizing larger errors through a regularization parameter C.Hyperparameters (C, ε, and kernel parameters) were tuned via cross-validation on the validation set.

The Long Short-Term Memory (LSTM) network, a recurrent neural network (RNN) variant, was chosen for its ability to model long-range temporal dependencies without suffering from vanishing or exploding gradients. The LSTM architecture employs gated mechanisms, input, forget, and output gates, to regulate information flow across time steps, enabling the network to retain relevant patterns while discarding noise. In this study, the LSTM model was trained using sequences of fixed-length input windows and optimized with the Adam optimizer, with hyperparameters such as the number of layers, hidden units, learning rate, and batch size determined empirically.

The proposed xLSTM model extends the standard LSTM by integrating additional exogenous features and architectural enhancements designed to improve predictive accuracy for environmental time-series data. In particular, xLSTM incorporates parallel convolutional layers to capture short-term local dependencies before feeding into the recurrent layers, allowing the model to exploit both short-term fluctuations and long-term trends. The architecture was trained using the same temporal input windows as the baseline LSTM, ensuring comparability, but its expanded feature extraction capacity allowed for improved adaptability across varying pollutant characteristics and peak-handling strategies.

All models were implemented in Python using established libraries, statsmodels for SARIMA, scikit-learn for SVR, and PyTorch for LSTM and xLSTM, and trained separately for each pollutant and peak-handling version of the dataset.

Design facing outputs are as follows. For each pollutant-strategy-model triplet, monthly forecasts are mapped to three risk bands with ±RMSE uncertainty envelopes and operationalized as triggers for planning and architectural design actions, including adjustments to street-canyon permeability and block porosity, the siting of green/blue infrastructure, and building ventilation/filtration scheduling. RMSE sets safety margins around thresholds, MAE characterizes routine operational tolerance, and R^2 is used to prefer the more stable model when RMSE is comparable.

3.3. Metrics

Model performance was quantitatively evaluated using four standard regression metrics: Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). These metrics were computed on the test set for each combination of pollutant, peak strategy, and model, ensuring that comparisons were made on unseen data without temporal leakage from the training or validation phases.

MAE (Mean Absolute Error): MAE measures the average magnitude of prediction errors, disregarding their direction, and is expressed in the same units as the target variable. It is calculated as **Equation 1**.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \quad (1)$$

where y_i is the observed value, $\hat{y}_i$ is the predicted value, and n is the number of observations. Lower MAE values indicate better predictive accuracy. MAE gives the average absolute forecast error in µg/m³ and serves as a practical "routine tolerance" for operations; because it weights small and large errors equally, we next consider MSE to explicitly penalize large misses and short-lived spikes.

MSE (Mean Squared Error): MSE captures the average squared difference between observed and predicted values, placing greater weight on larger errors due to squaring. It is given by **Equation 2**.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \quad (2)$$

A lower MSE implies better model performance, but because errors are squared, this metric is more sensitive to outliers compared to MAE. MSE squares residuals, making the metric especially sensitive to peak errors; to regain interpretability in the original units while retaining this peak sensitivity, we take its square root as RMSE.

RMSE (Root Mean Squared Error): RMSE is the square root of the MSE and re-expresses the error magnitude in the same units as the target variable, which aids interpretability as shown in **Equation 3**.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \quad (3)$$

RMSE is especially useful when large errors are undesirable, as it penalizes them more heavily. A large RMSE means that the model's predictions deviate substantially from the actual values, indicating lower accuracy and poorer predictive performance. RMSE expresses typical error magnitude in μg/m³ and aligns with design/regulatory thresholds; however, similar RMSE values can mask differences in model fit, so we complement it with R^2 to assess variance explained and stability.

R^2 Score (Coefficient of Determination): R^2 assesses how well the model explains the variance of the observed data. It is defined as in **Equation 4.**

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y}_i)^2} \quad (4)$$

where $\bar{y}$ is the mean of the observed values s. An R^2 value of 1 indicates perfect predictions, 0 implies the model is no better than predicting the mean, and negative values suggest performance worse than the mean. R^2 indicates how much of the observed variability the model captures and helps choose between models with comparable RMSE; in practice, we select the model that minimizes RMSE without degrading MAE and with the higher (non-negative) R^2.

4. FINDINGS

Table 1: Performance of SARIMA, SVR, LSTM, and xLSTM models for PM_{10} and SO_2 under different peak-handling strategies.

Model	Target	Peaks	MAE	MSE	RMSE	R^2
SARIMA	PM10	cap	4.8329	40.7825	6.3861	0.3309
xLSTM	PM10	cap	5.0005	41.4338	6.4369	0.3202
LSTM	PM10	cap	5.1378	46.2691	6.8021	0.2409
SVR	PM10	cap	6.9383	66.3797	8.1474	-0.0891
SARIMA	PM10	keep	4.8313	41.3923	6.4337	0.3209
xLSTM	PM10	keep	5.16	42.4085	6.5122	0.3042
LSTM	PM10	keep	5.4779	50.868	7.1322	0.1654
SVR	PM10	keep	6.9828	66.7479	8.1699	-0.0951
LSTM	PM10	remove	5.167	39.5821	6.2914	0.3632
xLSTM	PM10	remove	5.2018	39.9601	6.3214	0.3572
SARIMA	PM10	remove	5.1383	48.5534	6.968	0.2189
SVR	PM10	remove	6.1585	54.9973	7.416	0.1153
SVR	SO2	cap	1.1941	2.4868	1.577	0.3004
xLSTM	SO2	cap	1.2894	2.6764	1.636	0.2471
LSTM	SO2	cap	1.5689	4.0212	2.0053	-0.1312
SARIMA	SO2	cap	1.7024	4.0525	2.0131	-0.14
xLSTM	SO2	keep	1.1376	2.2185	1.4895	0.3759
SVR	SO2	keep	1.3177	2.8654	1.6928	0.194
SARIMA	SO2	keep	1.6791	3.9729	1.9932	-0.1176
LSTM	SO2	keep	1.567	4.0986	2.0245	-0.1529
SVR	SO2	remove	1.2363	2.4462	1.564	0.2398
xLSTM	SO2	remove	1.2613	2.4695	1.5715	0.2325
SARIMA	SO2	remove	1.5407	3.9763	1.9941	-0.2357
LSTM	SO2	remove	1.6071	4.3793	2.0927	-0.3609

In this study, forecasts for PM_{10} and SO_2 were generated using four models, SARIMA, SVR, LSTM, and the proposed xLSTM, under three different peak-handling strategies (cap, keep, and remove). Model

performance was evaluated using mean absolute error (MAE), mean squared error (MSE), root mean squared error (RMSE), and the coefficient of determination (R^2). The detailed results are presented in Table 1.

Air quality forecasting experiments were conducted using four models, SARIMA, SVR, LSTM, and the proposed xLSTM, across two pollutants (PM_{10} and SO_2) and three peak-handling strategies (cap, keep, remove). Model performances were evaluated using mean absolute error (MAE), mean squared error (MSE), root mean squared error (RMSE), and the coefficient of determination (R^2).

Figure 1 shows the time series comparison of actual and predicted PM_{10} concentrations for the xLSTM model under the cap strategy, illustrating that the model successfully tracks seasonal variations and most peak events, with deviations primarily during sharp changes.

For PM_{10} under the cap strategy, SARIMA achieved the lowest MAE (4.8329) and the lowest RMSE (6.3861), with an R^2 of 0.3309, indicating moderate explanatory power. The xLSTM model followed closely with an MAE of 5.0005 and an RMSE of 6.4369, delivering similar accuracy but slightly lower R^2 (0.3202). LSTM performed slightly worse than xLSTM (MAE 5.1378, RMSE 6.8021, R^2 0.2409), while SVR showed the weakest performance in this setting (MAE 6.9383, RMSE 8.1474, negative R^2 of -0.0891), suggesting poor generalization.

For PM_{10} with the keep strategy, SARIMA maintained a competitive MAE (4.8313) and RMSE (6.4337), outperforming the other methods in terms of absolute error. xLSTM remained the second-best performer, while LSTM and SVR exhibited larger errors and, in the case of SVR, a negative R^2, highlighting sensitivity to unaltered peaks.

Figure 2 presents the time series comparison of actual and predicted SO_2 concentrations for the xLSTM model under the cap strategy, showing that the predictions follow the general seasonal trend, with moderate discrepancies during high-variability periods.

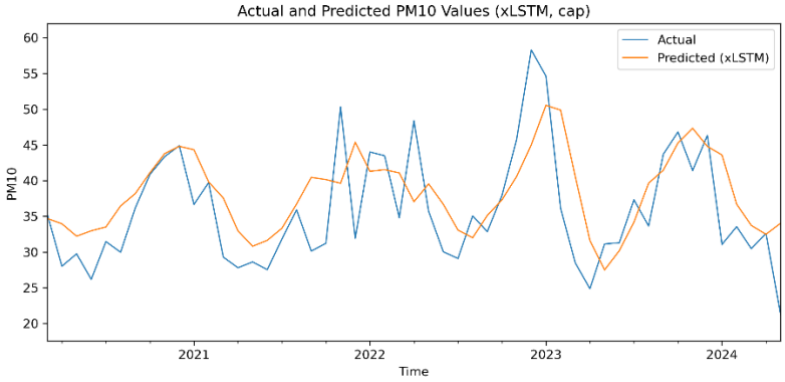

Figure 1: Time series comparison of actual and predicted PM_{10} concentrations using the xLSTM model under the "cap" peak-handling strategy

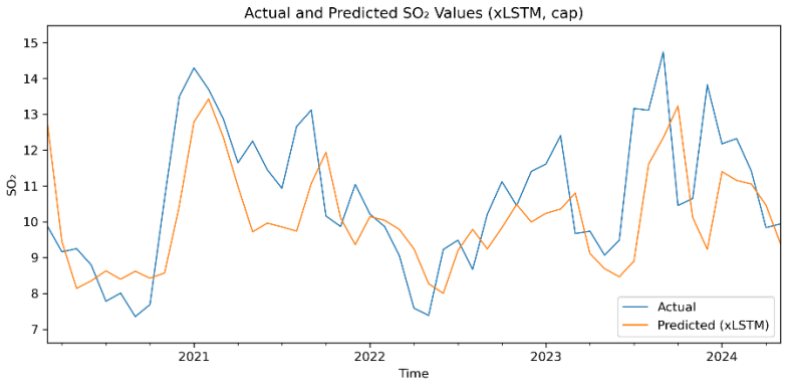

Figure 2: Time series comparison of actual and predicted SO_2 concentrations using the xLSTM model under the "cap" peak-handling strategy

Under the remove strategy for PM_{10}, SARIMA again delivered strong results with the smallest MAE and RMSE values, indicating that traditional time-series modeling benefits from peak removal. xLSTM continued to outperform LSTM and SVR in most metrics, though SARIMA retained the top position overall for PM_{10} across all strategies. In contrast, SO_2 forecasting revealed a different dynamic. With the cap strategy, xLSTM consistently provided lower MAE and RMSE values compared to LSTM and SVR, while also maintaining a substantially higher R^2 than SARIMA, which performed poorly (negative R^2) in this setting. This suggests that the xLSTM's hybrid architecture is particularly effective at modeling capped SO_2 series.

For SO_2 under the keep strategy, both xLSTM and LSTM achieved comparable MAE and RMSE values, with xLSTM often holding a slight edge. SARIMA again struggled, indicating that raw peak variability hinders its performance. SVR showed moderate predictive ability but did not surpass the deep learning models.

When removing peaks from SO_2 data, the performance gap between models narrowed, but xLSTM maintained competitiveness, often ranking among the top two models in terms of error metrics. LSTM sometimes matched its performance, though R^2 values indicated that xLSTM tended to capture variance slightly better.

These results indicate that for PM_{10}, SARIMA remains a strong contender, particularly when peaks are removed or capped, while xLSTM delivers competitive accuracy with better adaptability across varying strategies. For SO_2, xLSTM generally outperforms traditional methods, especially under capped and kept peak strategies, reflecting its ability to integrate temporal dependencies and adapt to pollutant-specific characteristics. These findings underscore the importance of selecting both the appropriate modeling approach and peak-handling strategy based on the pollutant and data characteristics.

The ±RMSE risk bands are converted into simple design triggers that link forecasts to actions for streets, open-space elements, and buildings. When PM_{10} peaks are cap/remove, SARIMA's smaller error lets ventilation lanes remain as narrow as 2-3 m; with keep data, xLSTM's steadier forecasts guide short curb-lane closures during predicted high episodes. For SO_2, xLSTM bands under keep/cap point to calm-wind months, steering the timing and placement of porous tree belts and wetland aerators so they disperse rather than trap pollutants. At the building scale, a high-band signal prompts roughly a 20% cut in outdoor-air intake and earlier filter changes, while a low-band signal opens the door to night-purge ventilation for free cooling. A trigger activates only after the same band is exceeded for two consecutive months; the ±RMSE width sets the safety buffer and MAE marks everyday tolerance, so interventions respond to sustained risk.

5. APPLICATIONS TO URBAN PLANNING AND ARCHITECTURAL DESIGN

5.1. Translating Forecasts into Risk Bands and Triggers

Monthly PM_{10} and SO_2 forecasts are mapped to Low/Moderate/High bands with ±RMSE envelopes, forming a design-facing interface. Interventions trigger only when band exceedance persists ≥2 months. RMSE sets safety margins, MAE informs routine tolerance, and R^2 breaks ties. For PM_{10}, SARIMA supports tighter margins under

cap/remove, while xLSTM offers cross-strategy stability; for SO_2 (keep/cap), xLSTM provides more reliable banding.

5.2. Street-Canyon, Corridors, and Neighborhood Form

Bands guide corridor- and block-scale design during seasons of elevated residence time. High-band months motivate porosity/ventilation corridors, careful timing of works, and calibrated traffic operations. SARIMA enables narrow buffers for PM_{10} when peaks are controlled; xLSTM stabilizes guidance when preprocessing policies are fluid.

5.3. Building-Scale Strategies and Sensitive Receptors

Façade orientation, mixed-mode ventilation, and filtration scheduling align to bands and their envelopes: wider envelopes favor conservative intake and earlier maintenance; narrower envelopes allow relaxed setpoints. For receptors (schools/clinics), intake siting, protective planting, and advisories track monthly transitions; for SO_2, xLSTM bands target seasons more dependably.

5.4. Implementation Pathway and İzmir-Specific Guidance

The workflow integrates with municipal GIS to test scenarios against keep/cap/remove policies. In İzmir, PM_{10} under cap/remove pairs best with SARIMA for tight triggers, while stable playbooks reference xLSTM; SO_2 actions anchor to xLSTM. The approach links forecasting to design, prioritizing targeted, temporally adaptive interventions over uniform, calendar-based measures.

6. POTENTIAL, LIMITATIONS, AND BROADER IMPACTS

6.1. Potential

The workflow runs on routinely collected monitoring data, generates design-facing risk bands with uncertainty envelopes rather than opaque scores, and integrates directly with existing GIS dashboards to enable scenario planning and policy testing. Accordingly, it is lightweight and readily deployable for municipal adoption.

6.2. Limitations

Several constraints qualify the findings. First, monthly aggregation smooths diurnal extremes and may understate short-lived peaks. Second, key exogenous drivers (e.g., wind fields, traffic volumes) are not explicitly modeled, so causal attribution is outside the present

scope. Third, heterogeneous station coverage can bias area-wide means and introduce spatial representativeness errors. Fourth, external validation beyond İzmir is pending, and generalizability should be confirmed in cities with different meteorology, emission profiles, and monitoring density.

6.3. Broader Impacts

Forecast-guided siting and operations can reduce exposure inequities around sensitive receptors (schools, clinics), align with sustainability objectives through targeted, temporally adaptive ventilation and filtration, and help avoid maladaptive greening that traps pollutants in narrow street canyons. By coupling uncertainty-aware forecasts to design and operations, the approach enables more precise, lower-cost interventions that advance public health and urban sustainability.

7. CONCLUSION

This study presented a comprehensive framework for predicting air quality in İzmir, Türkiye, by combining Internet of Things (IoT) data streams with traditional statistical models and advanced deep learning architectures. Monthly observations of PM_{10} and SO_2 from 1996-2024 were processed under three distinct peak-handling strategies (cap, keep, remove) to systematically assess the influence of extreme values on predictive performance. Four forecasting models, SARIMA, SVR, LSTM, and the proposed xLSTM, were rigorously evaluated using MAE, MSE, RMSE, and R^2 on strictly time-ordered train, validation, and test splits, ensuring no temporal leakage.

The findings reveal that for PM_{10}, SARIMA remained highly competitive, particularly under capped and removed peak strategies, often achieving the lowest absolute errors. However, xLSTM provided accuracy close to SARIMA's best results while maintaining greater adaptability across varying preprocessing strategies. For SO_2, xLSTM consistently outperformed the other models under capped and kept peak conditions, delivering lower error metrics and higher explanatory power than traditional methods, which struggled with the pollutant's variability. These outcomes underscore that no single model dominates in all scenarios; instead, model selection should be pollutant- and preprocessing-specific.

The integration of advanced AI techniques such as xLSTM with IoT-enabled environmental monitoring offers clear practical benefits. Beyond public health and regulatory applications, accurate air quality forecasts can directly inform urban design and planning decisions.

By identifying pollution hotspots and predicting their temporal patterns, city planners can optimize building placement to minimize exposure, allocate green spaces strategically to act as natural air filters, and guide transportation network designs that reduce emission concentrations in residential areas. Furthermore, predictive insights can support sustainable architectural practices by enabling ventilation, façade, and material choices that mitigate the impact of expected air pollution levels. Ultimately, the proposed approach not only advances methodological rigor in air quality prediction but also contributes to healthier, more sustainable urban living environments. By temporally aligning interventions with forecasted exposure, rather than applying uniform, calendar-based measures, municipalities can achieve comparable or superior air-quality improvements with lower energy and material use, realize climate co-benefits through demand-controlled ventilation and traffic retiming, and avoid maladaptive greening that traps pollutants in narrow street canyons. The banded, design-facing outputs further support equity by directing protection to sensitive receptors (e.g., schools and clinics). Future work should couple this workflow with multi-objective optimization and lifecycle/health-impact assessment to jointly quantify exposure reduction, energy and carbon savings, and financial costs, thereby institutionalizing forecast-guided planning within sustainability and resilience frameworks.

Author Contribution

The study idea and research design were developed by Kübra Öztürk, data collection and analysis were conducted by Kübra Öztürk. The writing, writing and editing of the article were carried out jointly by Kübra Öztürk and Zuhal Can. Both authors read and approved the final version of the article.

References

Abhijith, K. V., Kumar, P., Gallagher, J., McNabola, A., Baldauf, R., Pilla, F., Broderick, B., Di Sabatino, S., & Pulvirenti, B. (2017). Air

pollution abatement performances of green infrastructure in open road and built-up street canyon environments: A review. *Atmospheric Environment, 162*, 71–86.

Albuali, A., Srinivasagan, R., Aljughaiman, A., & Alderazi, F. (2023). Scalable lightweight IoT-based smart weather measurement system. *Sensors, 23*(12), 5569. https://doi.org/10.3390/s23125569

Ambildhuke, G., & Banik, B. G. (2022). IoT-based portable weather station for irrigation management using real-time parameters. *International Journal of Advanced Computer Science and Applications, 13*(5), 267–278.

Atali, A., Eren, B., Erden, C., & Atali, G. (2022). LSTM derin öğrenme yaklaşımı ile hava kalitesi verilerinin tahmini: Sakarya örneği [Forecasting air quality data with an LSTM deep learning approach: The case of Sakarya]. *Academic Perspective Procedia, 5*(3), 477–484.

Aydin, S., Tasyürek, M., & Öztürk, C. (2021). Derin öğrenme yöntemi ile İç Anadolu Bölgesi ve çevresi hava kirliliği tahmini [Air pollution forecasting for Central Anatolia and surroundings using a deep learning method]. *Avrupa Bilim ve Teknoloji Dergisi / European Journal of Science and Technology, 29*, 168–173.

Bansal, P., & Quan, S. J. (2024). Examining temporally varying nonlinear effects of urban form on urban heat island using explainable machine learning: A case of Seoul. *Building and Environment, 247*, 110957.

Bernardes, G. F. L. R., Ishibashi, R., Ivo, A. A. S., Rosset, V., & Kimura, B. Y. L. (2023). Prototyping low-cost automatic weather stations for natural disaster monitoring. *Digital Communications and Networks, 9*(4), 941–956. https://doi.org/10.1016/j.dcan.2022.05.002

Bolla, S., Anandan, R., & Thanappan, S. (2022). Weather forecasting method from sensor-transmitted data for smart cities using IoT. *Scientific Programming, 2022*, 1426575. https://doi.org/10.1155/2022/1426575

Elbasi, E., Mostafa, N., AlArnaout, Z., Zreikat, A. I., Cina, E., Varghese, G., Shdefat, A., Topcu, A. E., Abdelbaki, W., Mathew, S., & Zaki, C. (2023). Artificial intelligence technology in the agricultural sector: A systematic literature review. *IEEE Access, 11*, 171–202. https://doi.org/10.1109/ACCESS.2022.3232485

Elsaraiti, M., & Merabet, A. (2021). A comparative analysis of the ARIMA and LSTM predictive models and their effectiveness for predicting

wind speed. *Energies, 14*(20), 6782.

Geng, D., Zhang, H., & Wu, H. (2020). Short-term wind speed prediction based on principal component analysis and LSTM. *Applied Sciences, 10*(13), 4416.

Ioannou, K., Karampatzakis, D., Amanatidis, P., Aggelopoulos, V., & Karmiris, I. (2021). Low-cost automatic weather stations in the Internet of Things. *Information, 12*(4), 146. https://doi.org/10.3390/info12040146

Izmir Metropolitan Municipality. (2024). Hava kalitesi ölçüm değerleri [Air quality measurement values] [Dataset]. Retrieved September 16, 2025, from https://ulasav.csb.gov.tr/dataset/35-hava-kalitesi-olcum-degerleri

Jamil, H., Umer, T., Ceken, C., & Al-Turjman, F. (2021). Decision-based model for real-time IoT analysis using big data and machine learning. *Wireless Personal Communications, 121*(4), 2947–2959. https://doi.org/10.1007/s11277-021-08857-7

Karvelis, P., Mazzei, D., Biviano, M., & Stylios, C. (2020). PortWeather: A lightweight onboard solution for real-time weather prediction. *Sensors, 20*(11), 3181. https://doi.org/10.3390/s20113181

Kaya, S. M., Isler, B., Abu-Mahfouz, A. M., Rasheed, J., & AlShammari, A. (2023). An intelligent anomaly detection approach for accurate and reliable weather forecasting at IoT edges: A case study. *Sensors, 23*(5), 2426. https://doi.org/10.3390/s23052426

Leelavinodhan, P. B., Vecchio, M., Antonelli, F., Maestrini, A., & Brunelli, D. (2021). Design and implementation of an energy-efficient weather station for wind data collection. *Sensors, 21*(11), 3831. https://doi.org/10.3390/s21113831

Lu, B., Wang, R., Qin, Z., & Wang, L. (2023). A practice-distributed thunder-localization system with crowd-sourced smart IoT devices. *Sensors, 23*(9), 4186. https://doi.org/10.3390/s23094186

Mabrouki, J., Azrour, M., Dhiba, D., Farhaoui, Y., & El Hajjaji, S. (2021). IoT-based data logger for weather monitoring using Arduino-based wireless sensor networks with remote graphical application and alerts. *Big Data Mining and Analytics, 4*(1), 25–32. https://doi.org/10.26599/BDMA.2020.9020018

Mehmood, A., Lee, K.-T., & Kim, D.-H. (2023). Energy prediction and optimization for smart homes with weather metric-weight coefficients. *Sensors, 23*(7), 3640. https://doi.org/10.3390/s23073640

Mohapatra, D., & Subudhi, B. (2022). Development of a cost-effective IoT-based weather monitoring system. *IEEE Consumer Electronics Magazine, 11*(5), 81–86. https://doi.org/10.1109/MCE.2021.3136833

Nie, H., Liu, G., Liu, X., & Wang, Y. (2012). Hybrid of ARIMA and SVMs for short-term load forecasting. *Energy Procedia, 16*, 1455–1460.

Patel, P., Kalyanam, R., He, L., Aliaga, D., & Niyogi, D. (2023). Deep learning-based urban morphology for city-scale environmental modeling. *PNAS Nexus, 2*(3), pgad027. https://doi.org/10.1093/pnasnexus/pgad027

Roy, D. S. (2020). Forecasting the air temperature at a weather station using deep neural networks. *Procedia Computer Science, 178*, 38–46.

Singh, D. K., Sobti, R., Jain, A., Malik, P. K., & Le, D.-N. (2022). LoRa-based intelligent soil and weather condition monitoring with Internet of Things for precision agriculture in smart cities. *IET Communications, 16*(5), 604–618. https://doi.org/10.1049/cmu2.12352

Suresh, P., Aswathy, R. H., Arumugam, S., Albraikan, A. A., Al-Wesabi, F. N., Hilal, A. M., & Alamgeer, M. (2022). IoT with evolutionary algorithm-based deep learning for smart irrigation system. *Computers, Materials & Continua, 71*(1), 1713–1728.

Tsalikidis, N., Mystakidis, A., Koukaras, P., Ivaškevičius, M., Morkūnaitė, L., Ioannidis, D., Fokaides, P. A., Tjortjis, C., & Tzovaras, D. (2024). Urban traffic congestion prediction: A multi-step approach utilizing sensor data and weather information. *Smart Cities, 7*(1), 233–253.

Tukey, J. W. (1977). *Exploratory data analysis*. Addison-Wesley.

Venter, Z. S., Hassani, A., Stange, E., Schneider, P., & Castell, N. (2024). Reassessing the role of urban green space in air pollution control. *Proceedings of the National Academy of Sciences, 121*(6), e2306200121.

Vos, P. E. J., Maiheu, B., Vankerkom, J., & Janssen, S. (2013). Improving local air quality in cities: To tree or not to tree? *Environmental Pollution, 183*, 113–122.

Wang, Z., Hu, K., Wang, Z., Yang, B., & Chen, Z. (2024). Impact of urban neighborhood morphology on PM2.5 concentration distribution at different scale buffers. *Land, 14*(1), 7.

Woo, W., Richards, W., Selker, J., & Udell, C. (2023). WeatherChimes:

An open IoT weather station and data sonification system. *HardwareX, 13*, e00402. https://doi.org/10.1016/j.ohx.2023.e00402

Wu, Q., Wang, Y., Sun, H., Lin, H., & Zhao, Z. (2023). A system coupling GIS and CFD for atmospheric pollution dispersion simulation in urban blocks. *Atmosphere, 14*(5), 832.

Yang, J., Yu, M., Liu, Q., Li, Y., Duffy, D. Q., & Yang, C. (2022). A high spatiotemporal resolution framework for urban temperature prediction using IoT data. *Computers & Geosciences, 159*, 104991. https://doi.org/10.1016/j.cageo.2021.104991

Yu, M., Xu, F., Hu, W., Sun, J., & Cervone, G. (2021). Using long short-term memory (LSTM) and Internet of Things (IoT) for localized surface temperature forecasting in an urban environment. *IEEE Access, 9*, 137406–137418.