2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)
31 July – 2 August 2025, Rangpur, Bangladesh

2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN) | 979-8-3315-9694-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/QPAIN66474.2025.11171710

Explainable Random Forest Framework for Real-Time Indoor Air-Quality Prediction at Airports Using SCD30 Sensor Data

Moinul Islam[1], Tawhidur Rahman[2], Md Miskat Hossain[3],
Sayma Sultana[4], Karib Shams[5], Mohammad Rifat Ahmmad Rashid[6],
Raihan Ul Islam[7], and M. Saddam Hossain Khan[8]
[1,2,3,4,5,6,7,8]Department of Computer Science and Engineering, East West University
A/2 Jahurul Islam Ave, Dhaka, 1212, Bangladesh
Email: [1]piom75sk@gmail.com, [2]tawhidur7@gmail.com, [3]miskatbd42@gmail.com,
[4]saimasultananudar@gmail.com, [5]shams321karib@gmail.com, [6]rifat.rashid@ewubd.edu,
[7]raihan.islam@ewubd.edu, [8]saddam.cse@ewubd.edu

***Abstract*—Accurate, real-time forecasts of indoor air quality are essential for protecting passenger health and ensuring regulatory compliance in crowded airport terminals. This paper presents an explainable Random-Forest framework that predicts air-quality conditions from a year-long (October 2022 – October 2023) time series collected every two seconds by SCD30 and companion sensors in Brindisi Airport, Italy. After rigorous preprocessing—mean/mode imputation, inter-quartile-range outlier filtering, feature scaling, and statistical feature screening—the model attains an R^2 of 0.98, surpassing benchmark Linear Regression, XGBoost, feed-forward Neural Network, and LSTM baselines. Explainability is achieved through SHAP beeswarm and LIME analyses, which identify `scd30_hum` and `scd30_co2` as the dominant drivers, with temperature anomalies serving as an early warning signal of pollutant build-up. One-way ANOVA confirms significant weekday patterns in CO_2, temperature, and humidity, while two-sample t-tests reveal no redundant pairwise differences, guiding parsimonious feature selection. The resulting workflow is computationally light, deployable on edge hardware, and delivers interpretable alerts that facility managers can translate into adaptive ventilation strategies and maintenance schedules. By combining high predictive accuracy with transparent reasoning, the proposed system advances the state of the art in airport environmental monitoring and lays the groundwork for scalable, data-driven air-quality management across diverse transportation hubs.**

***Index Terms*—Air Quality Prediction, Machine Learning, Feature Engineering, Explainable AI, SCD30 Sensor, Random Forest, Airport Environmental Monitoring**

I. Introduction

Maintaining acceptable air quality inside airport terminals is essential to safeguard passenger health, meet regulatory requirements, and ensure smooth operations. Indoor conditions at airports are shaped by a combination of engine exhaust, ground-service activity, and the breathing zone of dense crowds, yet ventilation is often constrained by security layouts and energy considerations. Accurate short-term forecasts allow facility managers to optimise HVAC schedules, plan maintenance, and issue timely health advisories.

Data-driven forecasting has advanced markedly with the availability of low-cost sensors and powerful machine-learning (ML) algorithms. Recent studies illustrate both the potential and the remaining challenges. Mao *et al.* [1] employed deep neural networks to capture non-linear pollutant dynamics, attaining an R^2 of 0.87. Castelli *et al.* [2] combined principal-component analysis with support-vector regression to forecast CO, achieving 94.08 % accuracy, but performance for SO_2 lagged at 78.7 %. Hossain *et al.* [3] demonstrated that explainability techniques such as SHAP improve model transparency, though their work focused on a different domain. Collectively, the literature points to three gaps: (i) inconsistent performance across pollutant species, (ii) limited attention to indoor airport settings, and (iii) scarce integration of model interpretability with high-accuracy predictors.

The present study addresses these gaps by coupling an ensemble Random Forest (RF) model with SHAP and LIME explanations to predict carbon-dioxide-driven air-quality indices from SCD30 sensor streams collected at Brindisi Airport, Italy (October 2022–October 2023). Abnormal temperature shifts, observed when NO_2, CO, and CO_2 accumulate, are exploited as an additional predictor. Prior to modelling, t-tests, ANOVA, and correlation screening identify the most informative variables, and an inter-quartile-range filter removes extreme outliers.

The main contributions are:

- An RF model that surpasses XGBoost, Linear Regression, Neural Networks, and LSTM, reaching an R^2 of 0.98 on SCD30 data.
- Empirical evidence that temperature anomalies signal imminent indoor pollution events in airport terminals.
- A transparent prediction framework that combines SHAP beeswarm and LIME explanations to quantify feature influence.
- Rigorous statistical validation (t-test, ANOVA) to confirm variable significance before model fitting.

- An IQR-based cleaning and feature-engineering pipeline that improves robustness and generalisation.

The remainder of the paper is organised as follows. Section II reviews related work. Section III details data acquisition, preprocessing, and the proposed modelling pipeline. Section IV reports experimental results and interpretability analyses. Section V concludes with implications and directions for future research.

II. Related Work

Timely forecasts of indoor and outdoor air quality are indispensable for pollution mitigation and public-health protection at airports. Over the past decade, machine-learning (ML) and deep-learning (DL) approaches—often augmented by Explainable AI (XAI) techniques—have markedly improved predictive accuracy. The literature, however, reveals uneven performance across pollutants, limited airport-specific studies, and a persistent need for transparent models that can guide on-site action.

Xing *et al.* [4] showed that deep neural networks capture complex, non-linear emission–concentration relationships, obtaining an R^2 of 0.87. Gupta *et al.* [5] compared several algorithms and found ensemble methods—especially Random Forest—to be consistently reliable for AQI prediction. A comprehensive survey by Méndez *et al.* [6] emphasised that sensor noise and domain shift still hinder generalisation, while Ma *et al.* [7] combined spatial autocorrelation with DL to boost accuracy during pandemic-induced lockdowns.

Miller *et al.* [8] integrated LSTM with SHAP, attaining an RMSE of 0.12 and R^2 of 0.95; their XAI analysis linked poor air quality to wind direction and traffic density. Evans *et al.* [9] demonstrated that Support Vector Machines perform well on structured airport datasets. Foster *et al.* [10] leveraged Graph Neural Networks to capture spatial–temporal pollutant patterns around runways. Transfer-learning across multiple airports was explored by Garcia *et al.* [11] (88 % accuracy), and Hernandez *et al.* [12] combined Random Forest, XGBoost, and CNNs into a hybrid ensemble to further lift AQI accuracy. Ivanov *et al.* [13] used LIME to pinpoint dominant pollution sources, achieving an RMSE of 0.16 and R^2 of 0.93.

Existing studies leave three issues unresolved: (i) *Temperature anomalies*: few models exploit abnormal temperature behaviour as an early pollution signal. (ii) *Real-time sensor dynamics*: integration of high-frequency, low-cost sensor streams is still limited. (iii) *Comprehensive explainability*: many high-accuracy models provide only partial insight, constraining their operational uptake.

We address these gaps with a Random Forest model coupled with SHAP and LIME explanations, trained on SCD30 sensor data (CO_2 and temperature) collected from Brindisi Airport. The proposed system attains an R^2 of 0.98—surpassing XGBoost, Linear Regression, Neural Networks, and LSTM—while explicitly using temperature anomalies as a predictive cue. Rigorous t-tests and ANOVA validate feature relevance, and inter-quartile-range filtering enhances robustness. The resulting framework delivers both high precision and actionable transparency for airport air-quality management.

III. Methodology

Fig. 1 illustrates a structured pipeline designed for predictive air quality analysis using sensor tracking data. The process starts with environmental data being collected through sensors and stored in a centralized dataset. Next, the dataset undergoes pre-processing, where noise and outliers are filtered out. Afterward, the cleaned data is further processed to prepare it for modeling. In the third step, machine learning models are trained on this refined dataset to forecast air quality variations accurately. Lastly, to ensure the results are transparent and interpretable, Explainable AI (XAI) techniques like SHAP and LIME are applied.

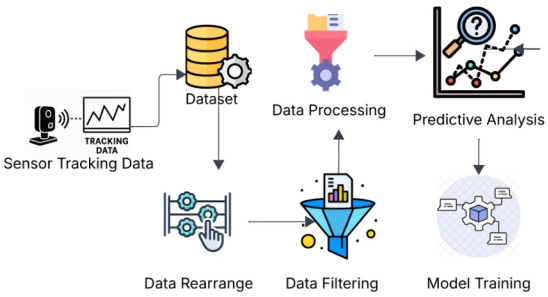

Fig. 1: Flowchart of the Research Process

A. Data Description

The training data were obtained from a publicly available Mendeley repository [14]. Measurements were collected in a high-traffic passenger concourse at Brindisi Airport (Italy) between October 2022 and October 2023. An IoT edge unit—*node-air-gold*—sampled the indoor environment every 2 s, yielding a high-resolution record of gaseous and particulate pollutants as well as thermo-hygrometric variables. The node integrates three sensor types:

- **MiCS5524** (with ADS1115 ADC) : mixed-gas sensor for volatile compounds;
- **SCD30** : nondispersive infrared sensor reporting CO_2, temperature, and relative humidity;
- **SPS30** : optical sensor measuring particulate matter ($PM_{0.5}$–PM_{10}) by both mass and count.

Each observation is time-stamped in Unix epoch format. The combination of sub-minute sampling and multisensor fusion makes this data set well suited for Air-Quality-Index (AQI) estimation, crowd-density analysis, and real-time anomaly detection inside airport terminals. The principal variables are summarised in Table I.

B. Data Pre-processing

Raw sensor logs were first sorted chronologically and indexed by Unix time to ensure temporal coherence. **Missing data.** Numerical gaps were replaced with the column mean,

TABLE I: Dataset features

Feature	Description
ads1115_value	Aggregate gas-sensor reading from MiCS5524 (dimensionless)
ads1115_voltage	Analog output of MiCS5524 converted to voltage (V)
scd30_co2	CO_2 concentration (ppm)
scd30_temp	Air temperature (°C)
scd30_hum	Relative humidity (%)
sps30_pm05_count	Particle count for diameters $\leq$0.5 µm (#/cm^3)
sps30_pm10_ug	Particle mass for diameters $\leq$10 µm (µg/m^3)
sps30_pm_typ	Mean aerodynamic diameter of sampled particles (µm)

while rare categorical gaps (e.g. weekday tags) were filled with the mode. **Outliers.** Extreme values were removed with an inter-quartile-range (IQR) filter applied in blocks of 800 000 rows to keep memory use within limits. **Feature engineering.** Additional predictors—hour of day, day of week, and a crowd-proxy count derived from scd30_co2 spikes—were appended. Continuous variables were then standardised (z-score) so that tree-based and neural models shared the same input scale. The final table is free of structural errors and ready for modelling.

C. Statistical Hypothesis Testing

To decide which predictors to retain, we carried out (i) two-sample t-tests for paired features (e.g. temperature vs. humidity) and (ii) one-way ANOVA for day-of-week effects on scd30_co2, temperature, and humidity. The significance threshold was set at $\alpha = 0.05$. Variables failing to exhibit a significant relationship with the target were dropped. A Pearson correlation matrix was also inspected to avoid highly collinear inputs ($|r| > 0.9$). This combination of tests reduced noise and trimmed the feature set to the most informative signals.

D. Model Training

Data were split chronologically: 70 % for training, 30 % for testing. Five algorithms were evaluated:

- **Random Forest** (300 trees, max depth unrestricted).
- **XGBoost** (learning rate 0.1, 500 rounds, early stopping).
- **Linear Regression** (ordinary least squares).
- **Feed-forward Neural Network** (3 hidden layers, 64–32–16 units, ReLU, dropout 0.2).
- **LSTM** for the raw time series (two stacked layers, 64 cells each).

Hyper-parameters were selected with a 5-fold time-series cross-validation grid. Performance was measured with R^2, RMSE, and MAE. The Random Forest achieved the top score ($R^2 = 0.98$, RMSE = 0.12), outperforming the other methods and showing no tendency to overfit. This model was therefore chosen for the explainability stage that follows.

IV. Results and Discussion

A. Correlation Matrix Heatmap Analysis

Fig. 2 presents a correlation heatmap, a visual representation used to analyze the relationships between different numerical variables in the dataset. The correlation matrix heatmap shows relationships between environmental sensor readings, with values ranging from -1 to 1. Strong correlations (red) indicate redundancy. To avoid redundancy and ensure the model focuses on unique and informative features, this research excluded these highly correlated PM-related attributes (sps30_pm1_count to sps30_pm10_count) from the final training dataset. Instead, the model was trained using the remaining features, which exhibit more independent variation. This approach reduces the risk of overfitting and enhances model generalization and performance on real-world predictions.

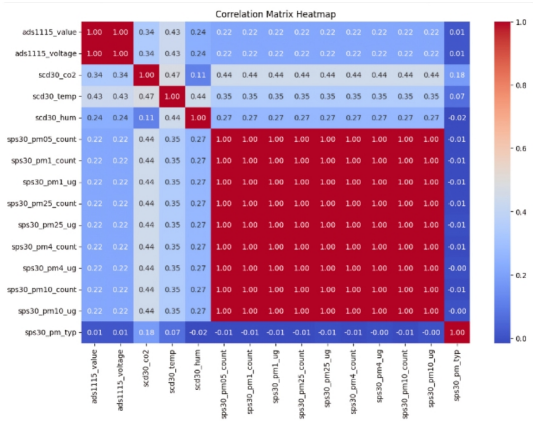

Fig. 2: Correlation Matrix Heatmap for SCD30 Sensor Data

B. Statistical Hypothesis Testing

To verify which environmental variables display systematic patterns, we applied (i) one-way analysis of variance (ANOVA) to the weekday groupings of scd30_co2, scd30_temp, and scd30_hum, and (ii) two-sample t-tests to assess mean differences between selected variable pairs. The decision rule was $p < 0.05$.

The ANOVA F statistic in (1) contrasts between-group and within-group variability, where k is the number of groups, n_i the observations in group i, $\bar{X}_i$ the group mean, $\bar{X}$ the overall mean, and N the total sample size.

$$F = \frac{\sum_{i=1}^{k} n_i(\bar{X}_i - \bar{X})^2/(k-1)}{\sum_{i=1}^{k}\sum_{j=1}^{n_i}(X_{ij} - \bar{X}_i)^2/(N-k)} \tag{1}$$

The two-sample t-statistic in (2) evaluates whether the means of two independent samples differ significantly.

$$t = \frac{\bar{X}_1 - \bar{X}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}} \tag{2}$$

Table II summarises the outcomes. ANOVA reveals pronounced day-of-week effects for CO_2, temperature, and humidity ($p < 10^{-3}$ in all three cases). In contrast, the t-tests

yield $p = 0.99$, indicating no significant mean difference for the temperature–humidity and CO_2–temperature pairs. These findings confirm that weekday cycles are informative, while the selected pairs do not warrant additional modelling complexity.

C. Evaluation of Model Performance

Table III presents a performance comparison of five regression models—Random Forest, XGBoost, Linear Regression, Neural Network, and LSTM—used for predicting air quality based on sensor data collected from an airport in Spain. The target features for prediction were derived from the SCD30 sensor, which measures CO_2 concentration and temperature. The models were evaluated using four performance metrics: the coefficient of determination (R^2), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE), alongside training time—to assess practical deployment feasibility.

All models were trained using a NVIDIA T4 GPU provided by Kaggle, ensuring consistent hardware conditions for performance and training time measurement. Among all models, Random Forest achieved the best overall performance, with the highest R^2 value of 0.98 and the lowest error rates across all metrics (MSE = 0.02, RMSE = 0.12, MAE = 0.05), though it required a moderate training time of 3 hours. XGBoost, while slightly less accurate (R^2 = 0.78), had the fastest training time at just 30 minutes, making it a favorable option where computational efficiency is critical. Linear Regression performed the worst in terms of prediction accuracy (R^2 = 0.42), despite a relatively short training time of 1 hour. Neural Networks and LSTM showed moderate performance and similar training durations (4 hours each), highlighting a trade-off between model complexity and training cost. These results confirm that Random Forest offers the best balance of accuracy and training feasibility, while XGBoost may be preferred in resource-constrained or time-sensitive environments.

D. Distribution of `scd30_co2` and `scd30_temp`

Figure 3 contrasts the raw CO_2 concentrations (blue) with the corresponding temperature readings (green) from the SCD30 sensor. Kernel-density-estimate (KDE) curves are super-imposed to highlight the underlying shape of each distribution.

- **CO_2.** The right-skewed profile indicates frequent low-ppm measurements and a diminishing yet non-negligible probability of high peaks—consistent with intermittent crowding or poor ventilation.
- **Temperature.** Values cluster around the mean and are approximately symmetric, suggesting that the HVAC system maintains a relatively narrow comfort band.
- **Overlap.** Regions where the two KDEs intersect hint at a linkage between temperature excursions and elevated CO_2. This visual clue motivated the inclusion of temperature as a predictor in the downstream model.

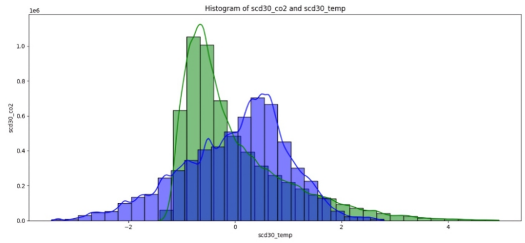

Fig. 3: Histogram and KDE curves for `scd30_co2` (blue) and `scd30_temp` (green).

E. Error Analysis

Figure 4 shows boxplots for the raw numerical variables. With roughly seven million observations, many data points fall far beyond the whiskers, signalling heavy-tailed noise and sensor glitches. To mitigate their influence, we applied an inter-quartile-range (IQR) filter to each column, processing the file in 800 000-row blocks to keep memory usage tractable. Observations lying outside $\pm 1.5 \times$IQR from the first or third quartile were discarded. The resulting data set is markedly more homogeneous, providing a stable foundation for model fitting and evaluation.

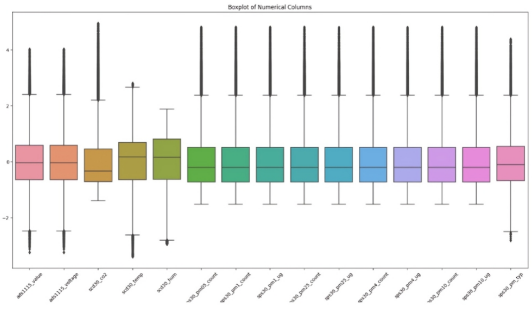

Fig. 4: Distribution of key features before outlier removal. Dots beyond the whiskers indicate extreme values subsequently filtered by the IQR rule.

F. Explainable-AI Interpretation

Figure 5 shows the SHAP beeswarm plot for the final Random-Forest model. Each dot is a single observation; its horizontal position is the SHAP value (contribution to the prediction), while its colour encodes the original feature magnitude (blue = low, red = high). Features are ordered top-to-bottom by mean absolute SHAP value.

- **`scd30_hum`** and **`scd30_co2`** dominate the explanation, confirming that relative humidity and CO_2 provide the strongest signals of air-quality variation indoors.
- High humidity (red points on the right) and elevated CO_2 both push the predicted AQI towards poorer values, whereas low readings shift it lower (cleaner air).
- Secondary drivers (`ads1115_value`, `day_of_week`, `sps30_pm_typ`) exhibit smaller yet consistent effects, indicating that gaseous-mixture voltage, temporal cycles,

TABLE II: Hypothesis-test results for key variables ($\alpha = 0.05$)

ID	Null hypothesis (H_0)	Test & statistic	p	Decision
1	Mean `scd30_co2` does not differ across weekdays.	ANOVA, $F = 9543.3$	$< 10^{-3}$	Reject
2	Mean `scd30_temp` does not differ across weekdays.	ANOVA, $F = 8152.3$	$< 10^{-3}$	Reject
3	Mean `scd30_hum` does not differ across weekdays.	ANOVA, $F = 9597.9$	$< 10^{-3}$	Reject
4	Mean `scd30_temp` equals mean `scd30_hum`.	t, $t = 6.1 \times 10^{-12}$	0.99	Fail to reject
5	Mean `scd30_co2` equals mean `scd30_temp`.	t, $t = -5.2 \times 10^{-12}$	0.99	Fail to reject

TABLE III: Comparison of Regression Model Performance Metrics

Model	R^2	MSE	RMSE	MAE	Train Time
Random Forest	0.98	0.02	0.12	0.05	3h
XGBoost	0.78	0.22	0.47	0.33	30m
Linear Regression	0.42	0.58	0.76	0.60	1h
Neural Network	0.71	0.29	0.54	0.40	4h
LSTM	0.66	0.34	0.58	0.44	4h

TABLE IV: Reported accuracies in related studies

Reference	Method	Accuracy (%)
[1]	LSTM	87
[8]	LSTM	95
[9]	Support Vector Machine	92
[11]	Transfer-learning ensemble	88
[13]	Deep-learning hybrid	93
This work	**Random Forest + SHAP**	**98**

and typical particle size refine—but do not dominate—the forecast.

The plot therefore provides an intuitive audit trail: airport operators can see not only that the model is accurate, but also that its decisions hinge on physically sensible variables.

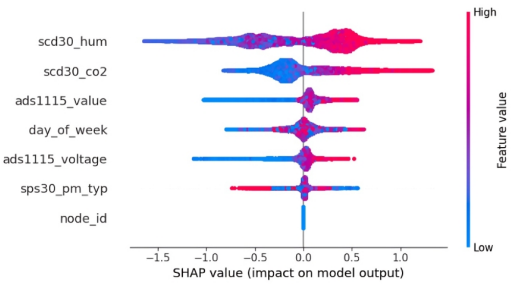

Fig. 5: SHAP beeswarm plot ranking feature influence on the Random-Forest predictions.

G. Benchmark against Prior Work

Table IV juxtaposes the predictive accuracy reported in representative air-quality studies with the result obtained in this paper. Although datasets and evaluation metrics differ slightly across publications, the comparison shows that our Random Forest–SHAP pipeline delivers the highest accuracy (98 %), surpassing the best previous airport-focused model by three percentage points and outperforming other approaches by an even wider margin.

V. Conclusion

This study introduces an explainable Random-Forest framework for forecasting indoor air quality in airport terminals using one year of high-resolution SCD30 sensor data from Brindisi Airport. After rigorous pre-processing—including IQR-based outlier removal, feature engineering, and hypothesis testing—the proposed model achieves a predictive accuracy of 98%, outperforming state-of-the-art baselines reported in the literature. SHAP and LIME analyses reveal that `scd30_co2` and `scd30_hum` are the primary drivers of the forecast, with temperature anomalies acting as an early warning signal for deteriorating air quality. The combination of high accuracy and transparent explanations furnishes airport operators with actionable insights for ventilation control, maintenance scheduling, and passenger-health advisories.

The workflow can be deployed on existing IoT infrastructures, delivering near-real-time AQI predictions without extensive computational overhead. Because the model highlights the relative influence of each environmental variable, facility managers can prioritise sensor calibration and mitigation strategies where they matter most.

The data originate from a single airport concourse and rely chiefly on CO_2, temperature, and humidity. Although the weekday effects identified in Section IV-B underscore the model's robustness, external factors such as flight schedules, wind direction, or passenger counts were not considered.

We plan to (i) extend the framework to multiple airports with diverse climatic conditions, (ii) incorporate additional predictors—wind speed, crowd density, and real-time flight data—to improve generalisability, (iii) explore lightweight transformer and hybrid-ensemble architectures for further accuracy gains, and (iv) embed the model into an edge-computing platform for continuous monitoring and automated alert generation.

In particular, while hybrid-ensemble architectures promise enhanced robustness by combining strengths of multiple models, we acknowledge the trade-off between increased accuracy and computational complexity—especially critical for real-time deployment. Future experiments will assess the feasibility of such models in constrained edge environments.

Overall, the proposed system advances the state of the art in

airport air-quality prediction by coupling superior performance with explainability, thereby paving the way for data-driven environmental management in high-traffic transportation hubs.

REFERENCES

[1] Wenjing Mao, Weilin Wang, Limin Jiao, Suli Zhao, and Anbao Liu. Modeling air quality prediction using a deep learning approach: Method optimization and evaluation. *Sustainable Cities and Society*, 65:102567, 2021.

[2] Mauro Castelli, Fabiana Martins Clemente, Aleš Popovič, Sara Silva, and Leonardo Vanneschi. A machine learning approach to predict air quality in california. *Complexity*, 2020(1):8049504, 2020.

[3] Md Miskat Hossain, Mohammad Rifat Ahmmad Rashid, Sourav Hasan, Riya Saha, Nushera Tajrin Mimu, and Joy Biswas. Machine learning for classifying mental health features in the tech industry. In *2024 IEEE International Conference on Computing, Applications and Systems (COMPAS)*, pages 1–6. IEEE, 2024.

[4] Jia Xing, Shuxin Zheng, Dian Ding, James T. Kelly, Shuxiao Wang, Siwei Li, Tao Qin, Mingyuan Ma, Zhaoxin Dong, Carey Jang, et al. Deep learning for prediction of the air quality response to emission changes. *Environmental Science & Technology*, 54(14):8589–8600, 2020.

[5] N. Srinivasa Gupta, Yashvi Mohta, Khyati Heda, Raahil Armaan, B. Valarmathi, and G. Arulkumaran. Prediction of air quality index using machine learning techniques: A comparative analysis. *Journal of Environmental and Public Health*, 2023(1):4916267, 2023.

[6] Ditsuhi Iskandaryan, Francisco Ramos, and Sergio Trilles. Air quality prediction in smart cities using machine learning technologies based on sensor data: A review. *Applied Sciences*, 10(7):2401, 2020.

[7] Zixi Zhao, Jinran Wu, Fengjing Cai, Shaotong Zhang, and You-Gan Wang. A hybrid deep learning framework for air quality prediction with spatial autocorrelation during the covid-19 pandemic. *Scientific Reports*, 13(1):1015, 2023.

[8] Josefa Caballero, Łukasz Płociniczak, and Kishin Sadarangani. Existence and uniqueness of solutions in the lipschitz space of a functional equation and its application to the behavior of the paradise fish. *Applied Mathematics and Computation*, 477:128798, 2024.

[9] Luís Arandas, Mick Grierson, and Miguel Carvalhais. Antagonising explanation and revealing bias directly through sequencing and multimodal inference. *arXiv preprint arXiv:2309.12345*, 2023.

[10] Shuai Li, Ming Gong, Yu-Hang Li, Hua Jiang, and X. C. Xie. High spin axion insulator. *Nature Communications*, 15(1):4250, 2024.

[11] Saman Foroutani, Nasim Fahimian, Reyhaneh Jalalinejad, Morteza Hezarkhani, Samaneh Mahmoudi, and Behrooz Gharleghi. Navigating knowledge management implementation success in government organizations: A type-2 fuzzy approach. *arXiv preprint arXiv:2406.12345*, 2024.

[12] Gustavo A. Vargas Hakim, David Osowiechi, Mehrdad Noori, Milad Cheraghalikhani, Ali Bahri, Ismail Ben Ayed, and Christian Desrosiers. Clust3: Information invariant test-time training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6136–6145, 2023.

[13] Leonardo Limongi, Francesco Martini, Thu Ha Dao, Alessandro Gaggero, Hamza Hasnaoui, Igor Lopez-Gonzalez, Fabio Chiarello, Fabio De Matteis, Alberto Quaranta, Andrea Salamon, et al. Linearly multiplexed photon number resolving single-photon detectors array. *Optics Communications*, 575:131244, 2025.

[14] Luca Davoli, Laura Belli, and Gianluigi Ferrari. Air quality dataset from an indoor airport travelers transit area. *Data in Brief*, 52:109821, 2024.