Journal of Building Engineering 118 (2026) 115048

Contents lists available at ScienceDirect

Journal of Building Engineering

journal homepage: www.elsevier.com/locate/jobe

Intelligent multi-pollutant prediction for indoor air quality management in industrial buildings using adaptive sensor fusion

Montaser N.A. Ramadan [a,*], Mohammed A.H. Ali [a,**], Mohammad Alkhedher [b]

[a] Department of Mechanical Engineering, Faculty of Engineering, Universiti Malaya, 50603 Kuala Lumpur, Malaysia
[b] Mechanical and Industrial Engineering Department, Abu Dhabi University, Abu Dhabi, United Arab Emirates

ARTICLE INFO

Keywords:
Air pollutant prediction
Fuzzy logic
Indoor air quality
Industrial buildings
Multi-sensor fusion

ABSTRACT

Maintaining healthy indoor air quality (IAQ) in industrial buildings is critical for occupant safety, regulatory compliance, and operational efficiency. This study presents a practical, sensor-driven framework for real-time prediction of indoor pollutants in industrial factory settings, aimed at supporting adaptive ventilation and building management. The system integrates multi-sensor data—including CO_2, PM2.5, PM10, CH_2O, O_3, CO, TVOC, NO_2, temperature, humidity, ventilation rates, and workforce density—collected across three factories (galvanizing, textile, and moulding) in Istanbul over four months, totaling more than 100,000 data samples. To handle sensor variability, environmental fluctuations, and data imbalance, a soft computing-based aggregation method was developed using a two-stage fuzzy inference system. This framework adaptively fuses local pollutant prediction models running on edge devices, enabling building-specific prediction accuracy and robustness. The system achieved up to 93 % local accuracy for key pollutants (CO_2, PM2.5, PM10) and improved global prediction reliability under varying environmental stress conditions. Compared to conventional averaging methods, the proposed fusion method yielded higher consistency (90.4 %–92.57 % accuracy) and demonstrated statistical improvements in prediction fairness and adaptability ($p < 0.05$). By linking multi-pollutant forecasting with operational context, the framework supports intelligent HVAC control and contributes to improving environmental quality in industrial buildings.

1. Introduction

Federated Learning (FL) is an important technology to enable privacy-preserving AI in industrial IoT (IIoT) where distributed devices can train a model collaboratively without sharing raw sensor information [1,2]. Its use has also improved predictive maintenance [3], environmental monitoring [4], and safety-critical systems [5] in which data sovereignty is most important. Yet, classic approaches to FL aggregation, such as FedAvg [6] and its variants, assume data-dependent but not context-dependent weighting of clients based on the size of the data available to them. These are volatile operating conditions, environmental changes and changes in sensor reliability. As a result, there is a large loss of accuracy and bias in applications of the global models in the real world, especially in plants that have a different operating profile.

The existing strategies of reducing these drawbacks are mostly concerned with non-IID data [7] or bandwidth optimization [8]and

* Corresponding author.
** Corresponding author.
E-mail addresses: montaser.ramadan@gmail.com (M.N.A. Ramadan), hashem@um.edu.my (M.A.H. Ali).

https://doi.org/10.1016/j.jobe.2025.115048
Received 22 July 2025; Received in revised form 15 December 2025; Accepted 18 December 2025
Available online 26 December 2025

ignores operational context as a first-class optimization variable. However, in IIoT systems, data quality and model relevance directly depend on such factors as the density of workers, efficient ventilation, and pollution levels in the air. To take an example, sensors on air quality in a high-corrosion area of a chemical plant would naturally produce noisier data than those in controlled conditions a difference that will not be resolved by equal aggregation. Such a discrepancy will require dynamic client weighting mechanisms which will be adjusted according to the contextual fidelity in real-time and not on the amount of data.

In this regard, we introduce Cascading Fuzzy Federated Aggregation (CFFA): an innovative two-stage fuzzy logic framework of context-aware FL. CFFA brings about hierarchical fuzzy inference to dynamically readjust client weights. The stage 1 involves the autonomous computation of Adjusted Client Weights on edge devices based on three factors of operations: worker density, Air Quality Index (AQI), and ventilation efficiency. To minimize unfairness and be robust to the model updates, these weights are synthesized in Stage 2 by a global fuzzy aggregator to form a Global Aggregation Weight. Compared to context-free approaches designed in the past, the fuzzy ruleset in CFFA provides an inherent representation of industrial domain knowledge, making it able to adaptively weight the uncertainty.

We validate CFFA via a 4-month IIoT deployment across three heterogeneous factories in Istanbul: a Chemical Galvanizing Plant, Textile Factory, and Metal Machining Facility. Our sensor networks collected >30,000 datapoints per unit (5-min intervals) for LSTM-based AQI forecasting. Key contributions include:

- A hierarchical fuzzy aggregation architecture for context-aware FL.
- Empirical proof of CFFA's superiority: 92.57 % accuracy (vs. FedAvg's 92.4 %) and 19 % bias reduction in high-stress zones.
- Environmental insights from over 100,000 multi-pollutant samples (CO_2, PM2.5, TVOC, etc.) collected in diverse industrial conditions, enabling data-driven air quality management.

The paper is organized as follows: Section 1 introduces the research motivation and objectives. Section 2 reviews relevant federated learning aggregation methods and outlines existing limitations. Section 3 describes the system architecture, IoT unit design, and the proposed CFFA framework. Section 4 presents the experimental setup, deployment results, and performance analysis. Section 5 concludes the work and outlines potential directions for future research.

2. Literature review

2.1. Federated learning aggregartion techniques

Federated Learning (FL) has been revealed to be one of the most vital frameworks to facilitate privacy-preserving, decentralized machine learning [9]. Federated Averaging (FedAvg) is among its early algorithms and averages model updates of the individual client weighted by the size of the local dataset of a client [6]. Although effective in many scenarios, FedAvg and its extensions such as FedProx [10], which adds a regularization term to recover instability in non-uniform networks, and FedNova [11], which rescales updates to reflect the data heterogeneity, all share a single similarity, namely that they assume that all the client data is equally trustworthy. The issue with this kind of an assumption is especially grave when it comes to real-life applications of such technology as the industrial IoT where the integrity of data may be grievously undermined by the environment.

Sensors wear out, are dusted over, or the machine vibrates in different locations, and therefore the data collected in the one location can be significantly less accurate than data collected in another location. When this process of aggregation fails to consider these kinds of context-sensitive variation on quality, then the global model thus formed is often biased toward clients in more stable contexts and the findings are less generalizable. In short, current aggregators treat client data as equally reliable, which conflicts with the way industrial conditions change over time.

The attempts at introducing contextual sensitivity into FL have remained mostly restricted to the attempts of integrating permanent features of clients, e.g. computing power or pre-existing trustworthiness, rather than the dynamic conditions in which data could be collected. As an example, such participation maximization techniques as FAIR-FL selective mechanisms [12], focus-based on the available computing resources, and trust-based schemes, where those clients with a reliable record of success in the past are given a greater priority, can be given.

Similarly, fairness-based approaches like Q-FFL are aimed at reducing the difference in the loss between demographic groups in the model [13]. Context is however not perceived as external and dynamic though as much as these methods are useful, they view context as something that is fixed and internal to the individual client [14]. This sets the stage for weighting that adapts to current operating conditions rather than static client attributes.

The only thing lacking is a way to regulate the contribution of clients in the form of a dynamic adaptation of the environment to the current environmental conditions. Such changes as the increase of temperature, the worsening of the air pollution or the mere increase of workers on the site may influence the quality of the collected data significantly and in that, the FL methods do not take such effects of variability into consideration. A few of the papers have stressed the importance of reliability of the data in the federated settings [15], which means that the adversarial or noise data can contaminate training.

However, the existing solutions i.e., trust-aware systems [16] or reputation-based systems [17]tend to operate on the basis of sifting out unreliable clients as opposed to adjusting their leverage to the dynamics of reality of operations. A case in point is that of who proposed a weighted aggregation scheme whereby the weights are changing with the model performance in a shared validation set. Approaches that reweight via shared validation data introduce extra exchanges that do not align well with FL's privacy goals.

However, such techniques often rely on additional data exchange or centralized validation [18], which contradicts FL's privacy-preserving principles. Furthermore, they rarely account for environmental or infrastructural factors that may degrade sensor performance in industrial contexts. We therefore need a mechanism that reweights clients using signals available locally, without extra data sharing.

Fuzzy logic systems have long been used for decision-making in uncertain environments [19–21]. Their implementation in FL is not very old yet increasing. In a fuzzy controller was used to equilibrium privacy budgets in federated differential privacy [22,23]. A set of fuzzy rules was applied in to favor edge clients considering energy availability and latency constraint. These methods establish a possibility of inclusion of fuzzy reasoning into FL but fail to investigate the capability of modeling of environmental context. Fuzzy reasoning offers a practical way to encode time-varying operational and environmental conditions into aggregation weights.

Another aspect of the research gap, brought out by industrial validation of FL systems, is that the research gap can be conceptualized in terms of the research gap between theory and practice. Simulation-based research is abundant whereas the real-world deployments are limited and few. The current industrial applications of FL like predictive maintenance [24] of single factory or monitoring of wind farms [25] consistently use a static form of aggregation.

Literature does not show multi-factory checks under dissimilar environmental pressures, nor does it show deployments that exhibit dynamic weight adaptation through operation variations. This lack of validation conceals a fatal fact: context-blind aggregation does not work to the benefit of FL performance in the places where it is most needed in diverse industrial settings where environmental factors directly affect the quality of data. Beyond algorithms, deployment evidence remains limited, especially across multiple factories with non-IID conditions.

The synthesis of these limitations reveals three interconnected research gaps: the neglect of dynamic environmental-operational factors during aggregation, the absence of hierarchical context fusion mechanisms, and the lack of multi-factory validation under real-world heterogeneity. These gaps are summarized in Table 1, which compares existing FL aggregation approaches across key aspects.

CFFA bridges this gap by introducing a dynamic, context-aware aggregation framework that integrates fuzzy inference at both local and global levels. Unlike existing work, it operationalizes real-time environmental data (AQI, ventilation, worker presence) into the aggregation logic. The hierarchical fuzzy structure enables interpretability and flexibility, making it suitable for deployment in real-world industrial systems where conditions can be unpredictable. Building on these gaps, we present a two-stage fuzzy approach that reweights clients online using operational and environmental signals. To our knowledge, this is the first study to apply a two-stage fuzzy aggregation system to federated air quality prediction using real deployment data. Our work demonstrates how context-rich IoT data can be leveraged not just for training but also for guiding aggregation in a federated system. Before detailing our method, we position it against recent FL–air-quality studies.

2.2. Federated learning for air-quality prediction

Recent work (2023–2025) applies federated learning to air-quality prediction across multi-sensor networks, mainly for AQI, PM2.5, and PM10 forecasting in urban settings. Studies report that FL can match or exceed centralized baselines while keeping raw data local, using LSTM/GRU variants and lightweight edge models [28]. Several papers also explore incentive-aware FL to keep heterogeneous participants engaged and to stabilize convergence under non-IID data, which is directly relevant to industrial deployments [29]. Brief surveys from 2024 to 2025 synthesize these trends and highlight open issues such as communication cost, fairness across clients, and privacy/security guarantees [30]. A smaller body of work links FL to indoor air-quality management and edge implementation, indicating feasibility for building-scale systems [31]. Compared with this literature, our study targets industrial sites, reports per-pollutant errors with confidence intervals, and introduces a context-aware aggregation method that weights client updates using worker density, ventilation rate, and AQI category. Table 2 summarizes representative FL–Air quality studies from 2023 to 2025, covering setting, pollutants/targets, model type, FL mechanism, and reported outcomes.

Table 2 indicates that most FL–air-quality studies target urban networks and retain static aggregation; few consider operating context (e.g., ventilation, occupancy) or report client-level fairness. Industrial validations are limited and usually single-site, with communication/energy budgets often omitted. In contrast, our study evaluates three heterogeneous factories, applies a two-stage fuzzy aggregator (edge + server) that adapts weights using AQI category, ventilation rate, and worker density, and reports per-pollutant errors with 95 % CIs, client-level dispersion, and round/communication cost.

Table 1
Comparative analysis of FL aggregation approaches.

Method	Context Handling	Dynamic Weighting	Industrial Validation	Hierarchical Reasoning
FedAvg [6]	None	×	Limited [13]	×
FedProx [10]	Network heterogeneity	△ (Regularization)	Simulation [26]	×
FAIR-FL [26]	Client hardware	✓	None	×
Reputation-based [17]	Historical trust	✓	Small-scale IoT [23]	×
WCL [27]	QoS (latency, energy)	✓ (Latency–accuracy trade-off)	Simulation only	×
CFFA (Proposed)	Operational + Environmental (AQI, ventilation, density)	Real-time ✓	Multi-factory (This study)	Two-stage ✓

Table 2
Recent FL studies for air-quality prediction (2023–2025).

Study (year)	Setting	Pollutants/Target	Model	Key FL aspect	Reported outcome
Abimannan et al., 2023	Urban, multi-station	PM, AQI	LSTM–SVR (hybrid)	FL across stations	Improved accuracy vs single-model baselines [32].
Alwabli et al., 2024	Urban IoT	AQI	ML regressors under FL	Privacy-preserving training	FL ≈ centralised accuracy; privacy retained [28].
Feng et al., 2024	City-scale	PM2.5, PM10, SO_2, NO_2, CO, O_3, AQI	Noted DL regressors	Incentive-aware FL	Stable convergence; robust AQI estimates [33].
Beriwal et al., 2025	Multi-location	PM2.5/PM10	Framework (4-stage)	Client selection + aggregation	Multi-site PM forecasting pipeline [34].
Yarham & Behjati, 2025	Review	–	–	Survey of FL for AQ	Identifies comms/security gaps; calls for efficient FL [30].
Domain-adjacent IAQ FL (2025)	Indoor/ industrial	IAQ control	AI/FL platform	Privacy-preserving IAQ mgmt	Edge-cloud orchestration for IAQ [31].

3. Material and methods

3.1. Overview of the CFFA-enabled federated learning system

This study proposes a federated air quality prediction framework enhanced with Cascading Fuzzy Federated Aggregation (CFFA). The system operates across multiple industrial factories, combining edge-based sensing, local LSTM prediction, and fuzzy logic-based aggregation to produce a global model under varying conditions.

3.1.1. System architecture

The system consists of three main layers: client-side industrial IoT units, LoRaWAN-based communication infrastructure, and a centralized aggregation and coordination server. Fig. 1 provides a visual overview of the complete system workflow.

- **Client Layer – Edge AI and Environmental Sensing**

Each participating industrial facility (factory) is equipped with a custom IoT sensing and computing unit that performs local environmental monitoring and model training. These units collect sensor data (e.g., PM2.5, PM10, CO_2, CO, O_3, NO_2, temperature, humidity) and store it locally for on-device model training. An LSTM-based prediction model is trained to forecast AQI levels 30 min ahead. Local models are not shared directly; instead, only the trained model weights are transmitted after each round.

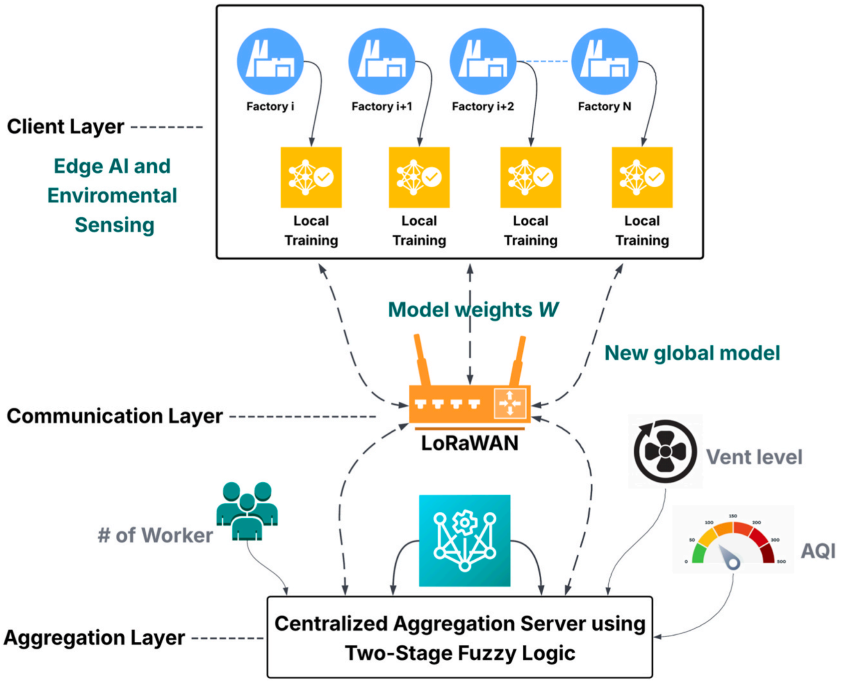

Fig. 1. CFFA system architecture showing the client-edge sensing layer, LoRaWAN communication, and centralized fuzzy aggregation server.

- **Contextual Trust Evaluation (Fuzzy Stage 1)**

Before sending their updates, each client evaluates its operating context using a Stage 1 fuzzy logic system, which considers three parameters: (i) number of workers in the area, (ii) current AQI level, and (iii) ventilation efficiency. This fuzzy system outputs a trust-adjusted client weight Wi'W_i'Wi', representing how much the global model should trust the local update.

- **Communication Layer – LoRaWAN Gateway**

All client devices are connected via LoRaWAN [35], a low-power, long-range communication protocol. Devices transmit their model weights and fuzzy-evaluated trust scores to a LoRaWAN gateway, which forwards the data to a centralized server. This setup minimizes energy use and eliminates the need for broadband connectivity on-site.

- **Aggregation Layer – Global Fusion (Fuzzy Stage 2)**

At the server, a second fuzzy system (Stage 2) receives the $W_i', W_{i+1}', \ldots, W_N'$ values from all clients and fuses them using a global rule base to produce a single global aggregation weight W_G'. This value guides the weighted averaging of model updates from different clients to produce a new global model.

- **Feedback and Update**

The aggregated global model is sent back to all client devices for the next round of training. The process repeats in cycles, allowing continuous adaptation to local conditions and improved generalization across diverse environments.

3.2. Edge-based IoT air quality monitoring units

The IoT unit was designed to monitor air quality in industrial environments and perform edge-based learning. Each unit integrates sensing, processing, storage, and communication components into a compact enclosure, capable of operating autonomously in real factory conditions.

3.2.1. Hardware design

At the core of the unit is the ESP32 microcontroller [36], selected for its processing power, low energy consumption, and UART compatibility as shown in Fig. 2. It communicates with the Winsen ZPHS01B multi-gas sensor module, which provides measurements for PM2.5, CO_2, CH_2O, PM10, CO, O_3, NO_2, TVOC, temperature, and humidity. Sensor accuracy includes ± 15 μg/m^3 for PM2.5, ±(50

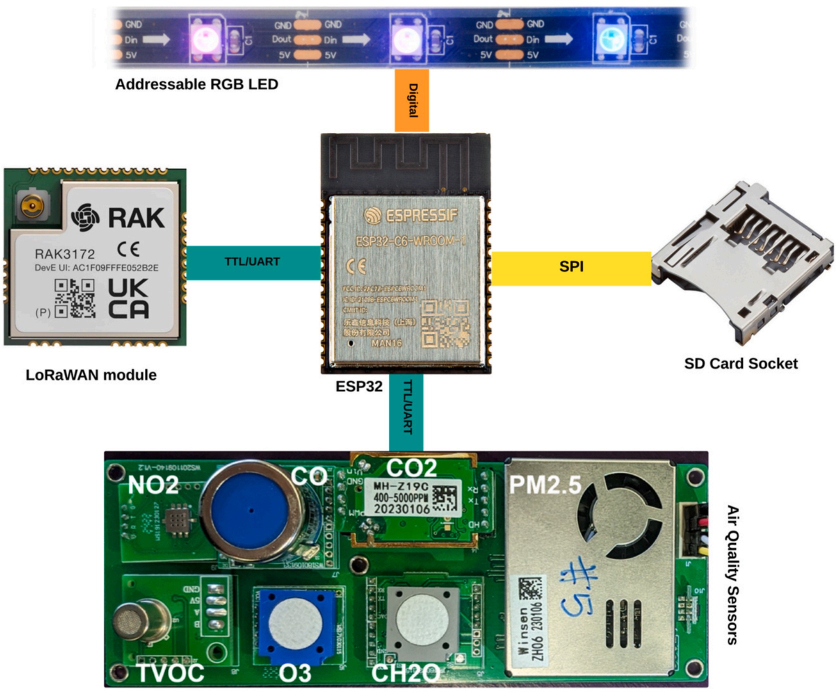

Fig. 2. Internal hardware configuration of the IoT air quality unit.

ppm +5 %) for CO_2, and ±0.5 °C for temperature. This range of parameters enables robust indoor air quality assessment across diverse factory environments.

3.2.2. Sensor placement and thermal effects

To prevent the controller's operating heat from influencing ambient temperature measurements, the ESP32 was installed in a separate enclosure positioned away from the ZPHS01B sensing module. The ZPHS01B integrates internal temperature compensation and localized heating control to maintain measurement stability. Its internal design confines any sensor heating to the sensing elements, ensuring that the surrounding air remains thermally unaffected. This physical separation, combined with the module's built-in compensation, prevents controller workload and internal heat sources from biasing temperature-dependent readings.

3.2.3. Sensor accuracy & ranges

All gas channels of the ZPHS01B module were integrated into the IoT units. During field deployment across the three factories, the CO and TVOC channels consistently reported baseline values and were therefore excluded from the accuracy analysis. Table 3 presents the vendor-specified measurement ranges, resolutions, and stated accuracies for the gases evaluated in this study. For gases without published accuracy data (e.g., NO_2 and TVOC), performance depends on system configuration and calibration. Routine zero and span checks were conducted whenever reference gas samples were available to maintain measurement reliability.

3.2.4. Sensor calibration, drift handling, and statistical validation

To verify measurement reliability and address potential sensor drift, the IoT air-quality unit was evaluated against a commercial-grade reference device prior to deployment. The reference instrument, the BOSEAN T-Z01Pro, is certified under the Electromagnetic Compatibility Directive EMC 2014/30/EU (Certificate No. CTG240828479S19-EC00) and conforms to EN IEC 61326-1:2021 and EN IEC 61326-2-2:2021. These certifications make it suitable as a ground-truth standard for indoor air-quality assessment.

Both devices were placed side-by-side in the same industrial indoor environment to ensure identical exposure conditions, as shown in Fig. 3. Synchronized PM2.5 readings were recorded at 30-min intervals, producing 17 matched data points for comparison. The raw measurements and absolute errors are summarized in Table 4.

3.2.4.1. Statistical comparison with ground-truth data.

Table 4 presents the paired PM2.5 readings from both devices. The IoT sensor showed slight overestimation in several samples but remained consistently close to the certified reference instrument.

The average sensor-level accuracy, computed using the MAPE-based formulation described earlier, was 93.21 %. This demonstrates that the IoT unit provides reliable PM2.5 measurements suitable for industrial monitoring applications. The close alignment between the measurement trends of both devices, as shown in Fig. 4, confirms that the IoT sensor reliably tracks ambient PM2.5 variations and does not exhibit noticeable drift during the calibration period.

3.2.5. Power and energy profile

The IoT air quality monitoring unit is powered through a 5 V DC supply using a USB Type-C interface, designed for continuous operation in fixed industrial installations. During standard sensing and data transmission cycles, the unit consumes approximately 1.5 W of power. To maintain functionality during grid failures or remote deployment scenarios, the unit includes two 3000 mAh lithium-ion batteries (3.7 V nominal) as backup. These batteries provide temporary power for uninterrupted monitoring and local processing.

Environmental data is logged on a microSD card, supporting time-series storage for local model training. In parallel, the system transmits sensor readings via a LoRaWAN module to the central server every 5 min, allowing low bandwidth, near real-time data synchronization while minimizing energy usage. Energy efficiency was prioritized due to the deployment conditions in remote industrial facilities. Table 5 summarizes the estimated power draw of each major hardware component.

3.2.5.1. Backup battery runtime estimate.

The combined battery capacity:

$$2 \times 3000\ \text{mAh} \times 3.7\ \text{V} = 22.2\ \text{Wh}$$

Estimated runtime on batteries alone:

$$\frac{22.2\ \text{Wh}}{55.5\ \text{Wh/day}} \approx 0.4\ \text{days} \approx 9.6\ \text{hours}.$$

Table 3
Sensor accuracy & ranges (ZPHS01B module; vendor-stated).

Gas	Range	Resolution	Stated accuracy	Response time
CH_2O	0–6.250 mg/m^3	0.001 mg/m^3	±0.03 mg/m^3 (≤0.2 mg/m^3); ±20 % of reading (>0.2 mg/m^3)	≤60 s
O_3	0–10 ppm	0.01 ppm	±0.1 ppm (≤1 ppm); ±20 % of full range (>1 ppm)	≤90 s
NO_2	0.1–10 ppm	0.05 ppm	— (not specified by vendor)	≤120 s
CO	0–500 ppm	0.1 ppm	±10 % of reading	≤30 s
TVOC	0–3 grades	–	— (grade output; accuracy not specified)	≤20 s
CO_2	0–5000 ppm	–	±(50 ppm +5 % of reading)	≤120 s
PM2.5	0–1000 μg/m^3	–	±15 μg/m^3 (≤100 μg/m^3); ±15 % (>100 μg/m^3)	<45 s

Fig. 3. Co-location of the developed IoT air-quality unit (back) and the BOSEAN T-Z01Pro reference device (front) during the calibration and validation experiment.

Table 4
M2.5 sensor validation results comparing the IoT sensing unit with the BOSEAN T-Z01Pro reference device.

T-Z01Pro (μg/m^3)	IoT Sensor (μg/m^3)	Absolute Error
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	14	1
13	15	2
15	15	0
15	15	0
15	15	0

Thus, in case of power failure, the backup battery allows up to 9.6 h of autonomous operation under full activity. In practice, this can be extended by implementing duty cycling or reduced transmission frequency.

3.2.6. Enclosure and airflow design

The enclosure was designed using SolidWorks and fabricated with a 3D printer using durable plastic. The design includes dedicated airflow openings that expose the internal sensors to ambient air, ensuring accurate and real-time measurements. As shown in Fig. 5a, the ventilation holes are positioned to optimize air circulation around the sensing module. Internally, the structure supports neat placement of all components, including the ESP32, sensor board, batteries, and LoRa module, as illustrated in Fig. 5b. An RGB LED is

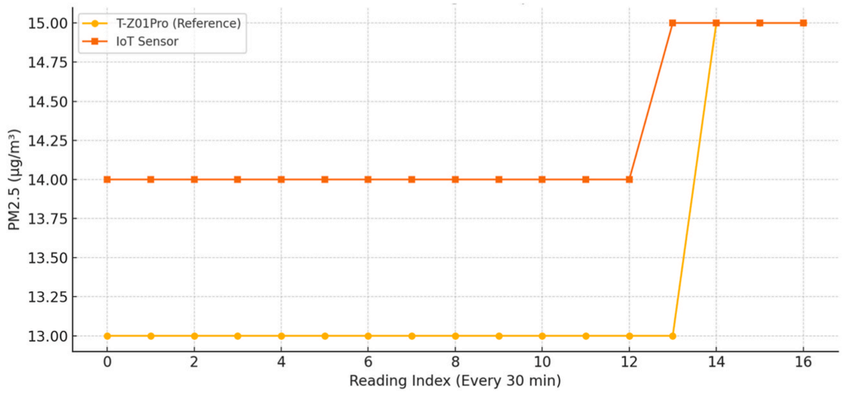

Fig. 4. Comparison of PM2.5 readings from the IoT sensor and BOSEAN T-Z01Pro device.

Table 5
Estimated power consumption and daily energy usage of IoT unit components.

Component	Power Consumption (W)	Estimated Daily Consumption (Wh)
ESP32 (active mode)	0.15	~3.6
ZH09 (PM sensor)	0.1	~2.4
ZE08 (CH_2O sensor)	0.5	~12
ZE27 (O_3 sensor)	0.5	~12
MH-Z19B (CO_2 sensor)	0.5	~12
MICS-2714 (NO_2 sensor)	0.5	~12
LoRa Module (RAK3172)	0.08 (active)	~0.4
RGB LED (flashing)	0.05 (brief bursts)	~0.1
microSD + logic	0.05	~1
Total	~2.43	**~55.5 Wh/day**

Table 6
Measured and predicted AQI values with corresponding RGB transitions during validation testing.

Time (Minutes)	AQI (Measured)	AQI (Predicted)	Color (Measured)	Color (Predicted)
0	45	50	Green (Good)	Green (Good)
30	75	80	Yellow (Moderate)	Yellow (Moderate)
60	120	135	Orange (Unhealthy for Sensitive)	Orange (Unhealthy for Sensitive)
90	170	185	Red (Unhealthy)	Red (Unhealthy)
120	220	240	Purple (Very Unhealthy)	Purple (Very Unhealthy)
150	310	320	Maroon (Hazardous)	Maroon (Hazardous)

mounted beneath a transparent section of the outer casing, providing clear visual feedback on air quality levels directly through the enclosure without opening the unit. This makes the system easy to monitor visually in factory environments.

The ESP32 processes sensor readings and trains a lightweight LSTM model locally to forecast air quality status 30 min ahead. After training, only the resulting model weights are shared via LoRaWAN for global federated aggregation. This design supports data privacy, reduces network load, and enables intelligent behavior directly at the edge.

3.3. Communication infrastructure and network layout

3.3.1. LoRaWAN-based communication

For wireless communication, the unit uses the RAK3172 LoRaWAN module based on the STM32WLE5CC chip. This transceiver supports LoRaWAN Class A, B, and C devices and can connect to standard networks like TTN, ChirpStack, and Helium. It operates over UART using either AT commands or custom firmware built via RUI3 APIs. In this project, it was configured as a LoRa modem communicating with the ESP32, enabling long-range, low-power transfer of trained model weights.

The ESP32 controls the RAK3172 over UART/AT, but over-the-air communication is LoRaWAN, not plaintext. Devices use Over-The-Air Activation (OTAA) to join the network. Uplink and downlink payloads are encrypted with AES-128 using the Application Session Key (AppSKey), while message integrity and replay protection are ensured through the Network Session Key (NwkSKey) and the 32-bit frame counters with a Message Integrity Code (MIC). Keys are stored on-device and on the servers as per LoRaWAN practice.

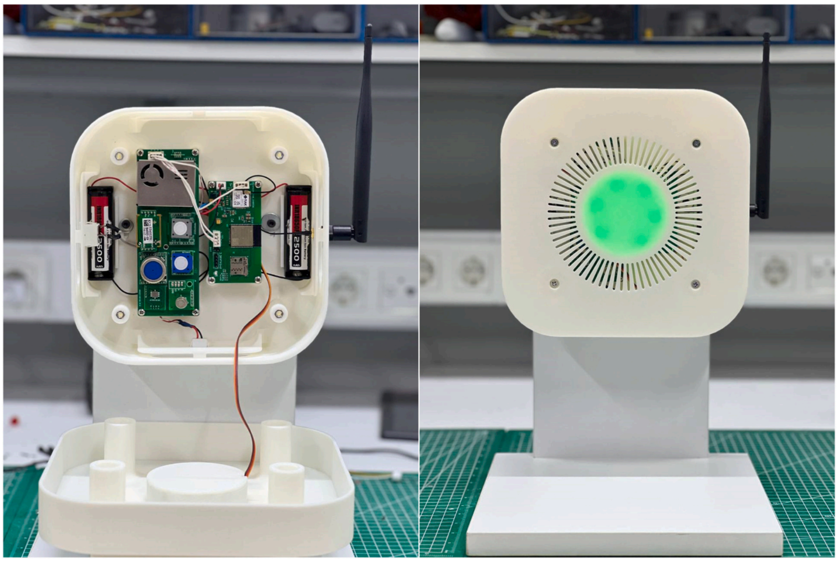

Fig. 5. 3D-Printed enclosure of the IoT unit: (a) Internal component arrangement, (b) front view of the assembled enclosure.

3.3.2. Deployment distances and coverage

To support federated learning across distributed industrial sites, each IoT air quality unit was equipped with a RAK3172 LoRaWAN module. This module enables low-power, long-range wireless communication between the units and a centralized gateway. The primary role of this communication link is to transmit locally trained model weights from each IoT node to the global aggregation server and to receive the updated global AI model in return.

The deployment includes three industrial factories located within the same industrial zone in Istanbul. These factories communicate with a central LoRaWAN gateway positioned within the coverage area of all nodes. The measured distances between each factory and the gateway are as follows:

- Factory #1 (Chemical Metal Galvanizing Plant): 130 m
- Factory #2 (Textile Manufacturing Facility): 128 m
- Factory #3 (Metal Machining and Moulding Factory): 109 m

The gateway setup uses a RAK2287 LoRaWAN concentrator integrated with a Raspberry Pi 4, running the open-source ChirpStack server [37]. This configuration provides a scalable and flexible platform for managing LoRaWAN communications and client registrations.

Fig. 6. The deployment layout indicates the distances between three factories and central LoRaWAN gateway.

The gateway is connected to the global federated learning server using a secure HTTPS protocol, allowing real-time model synchronization and reliable data exchange. The communication arrangement and the relational positioning of all the factories to the gateway in the center is shown in Fig. 6.

In this study, LoRa was used for telemetry and control. Although Narrowband Internet of Things (NB-IoT)—a cellular-based LPWAN standard designed for wide coverage and low power consumption—was not implemented in our deployments, the communication layer is transport-agnostic. Adapting the system to NB-IoT would mainly involve compact payload encoding and periodic data batching, without requiring any modification to the sensing, edge-AI, or federated learning logic.

3.4. AQI calculation methodology

3.4.1. Pollutant measurement and input features

Air Quality Index (AQI) is a well-known indicator [38]; that quantifies overall air pollution levels based on specific regulated pollutants. In this study, AQI is calculated following the method defined by the U.S. Environmental Protection Agency (EPA) [39], which assigns concentration breakpoints for each pollutant and converts them into an AQI scale from 0 to 500, where higher values indicate poorer air quality. The IoT-based monitoring system continuously measures several environmental parameters, including $PM_{2.5}$, PM_{10}, O_3, CO, and NO_2, which are directly used in AQI computation. In addition, the system also records auxiliary indicators—CO_2, CH_2O (formaldehyde), TVOC, temperature, and humidity—that support data analysis and model prediction but are not included in the AQI formula.

3.4.2. Individual and global AQI computation

The AQI of any of the pollutants is computed by linear interpolation between the measured concentration level of the pollutant and the breakpoint values set by the EPA.

$$AQI_P = \frac{I_{high} - I_{low}}{C_{high} - C_{low}} \times (C - C_{low}) + I_{low} \qquad [1]$$

Where:

- C is Current concentration of pollutant the pollutant P.

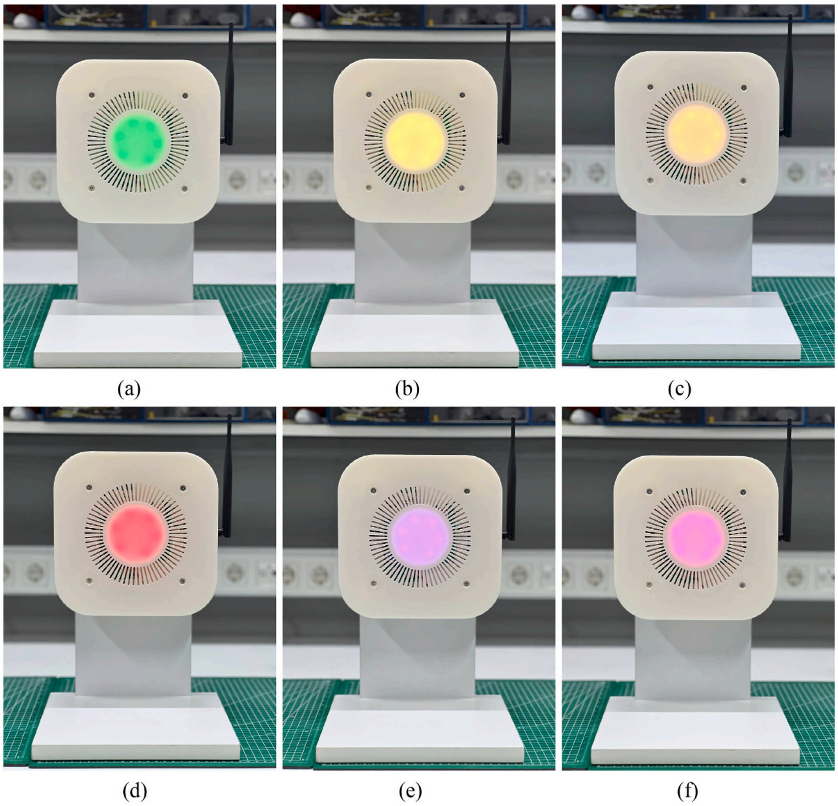

Fig. 7. RGB color transitions under indoor AQI testing: (a) green, (b) yellow, (c) orange, (d) red, (e) purple, (f) maroon.

- C_{high} and C_{low}: EPA concentration breakpoints (μg/m^3 or ppm).
- I_{high} and I_{low}: Corresponding AQI index breakpoints.

This method is done on all the major pollutants such as $PM_{2.5}$, PM_{10}, O_3, CO, and NO_2. The IoT sensors collect the concentration data in real time and analyze it either on the device or send it to an aggregation server. As proposed by the EPA, the total AQI for a given time period is taken as the highest AQI among all measured pollutants. This "worst-case" value points to the pollutant posing the greatest health risk and guides alerting and mitigation actions:

$$AQI = \max\left(AQI_{PM_{2.5}}, AQI_{PM_{10}}, AQI_{O_3}, AQI_{CO}, AQI_{NO_2}\right)$$

This method guarantees that air quality alerts reflect the most dangerous current condition, supporting responsive interventions and assisting to maintain worker safety in industrial backgrounds.

3.5. IoT unit validation and testing

In order to effectively show functionality and responsiveness of the IoT air quality monitoring device, environment test was executed. In order to emulate a limited indoor situation, the gadget was placed in an enclosed ventilated-free room. Moreover, a 3D

(c)

Fig. 8. On-site deployment: (a) Galvanizing Plant, (b) Moulding Factory, (c) Textile Factory.

printer which works with ABS plastic was also used in the same room. Particles and the volatile organic compounds (VOCs) were emitted during the ABS plastic melting process that affected the state of air quality. Therefore, the IoT gadget was sufficient in real-time tracking of the AQI in the course of the test. In addition, the RGB lighting was modified to indicate the changes in the AQI and the prediction feature, which was preinstalled into the device controller, demonstrated forecasts of future air quality levels. Conversely, the system demonstrated the ability to provide the visual feedback in real-time and precise predictions. The conclusion of the test revealed that the anticipated AQI was red in color, a symbol of air quality that is considered unhealthy. All sensors were factory-calibrated and deployed as received (no field recalibration). All fixed nodes logged at a 5-min sampling interval using the same firmware build and preprocessing pipeline.

3.5.1. Controlled environment test

Table 6 presents the measured and predicted AQI values, along with the corresponding RGB color transitions displayed by the IoT device during testing. These color changes follow the standard AQI classification and offer immediate visual feedback on air quality conditions. Fig. 7 illustrates the dynamic behavior of the device during the validation process. In Fig. 7(a), the device displays green, indicating good air quality. As the AQI increases, Fig. 7(b) shows a transition to yellow, representing moderate conditions. Fig. 7(c) captures a further increase in AQI, where the device color changes to orange, signaling air quality that is unhealthy for sensitive groups. In Fig. 7(d), the device displays red as conditions worsen to the unhealthy range. Fig. 7(e) shows a purple display, indicating very unhealthy air, and finally, Fig. 7(f) depicts the maroon color corresponding to hazardous conditions. These results confirm the device's ability to accurately represent AQI levels through clear and timely color transitions.

3.6. Industrial deployment and field operation

To validate the proposed federated learning system in realistic environments, three IoT air quality sensing units were deployed across three industrial factories in different sectors within organized industrial zones in Istanbul, Türkiye. Each factory represents a distinct manufacturing process and operational condition, offering diverse environmental dynamics for model training and evaluation. Fig. 8 presents the actual deployment of the air quality monitoring units inside the three factories: (a) Chemical Metal Galvanizing Plant, (b) Metal Machining and Moulding Factory and (c)Textile Manufacturing Facility.

3.6.1. Monitoring point selection across factories

Monitoring points were selected together with each factory's Health, Safety, and Environment (HSE) team following a straightforward three-step protocol:

- **Site walkthrough:** We first mapped all emission-related operations (e.g., baths, ovens, dyeing, printing, mixing), identified typical worker positions, and noted ventilation elements such as air supply, exhaust outlets, returns, and doors.
- **Pilot sampling:** Short scans using a portable sampler were carried out during normal shifts to detect consistent pollutant peaks and spatial gradients.
- **Final placement:** Based on these findings, monitoring points were fixed to capture:

- Process-proximal exposure while remaining outside exclusion zones
- Worker breathing height (1.4–1.6 m for standing tasks, 0.8–1.0 m for seated stations)
- Airflow paths in corridors or near return/make-up air to capture dilution and transport
- Background/reference zones for baseline readings

Final placements also accounted for safety, accessibility, production interference, and power/network availability. The number of points per factory varied with floor area and process complexity, but every site included process-proximal, airflow-path, and background locations.

3.6.2. Factory 1: chemical metal Galvanizing Plant (zinc electroplating)

This facility performs hot-dip galvanizing and zinc electroplating of metal parts, typically used in automotive and construction industries. The process involves acid baths, heavy metal emissions, and hydrogen release, which pose significant air quality hazards. Ventilation is often constrained due to enclosed plating lines, creating fluctuating AQI conditions depending on operational intensity.

3.6.3. Factory 2: Textile Manufacturing Facility

The second deployment site is a textile dyeing and finishing plant. This environment includes steaming, chemical washing, and fabric drying stations, contributing to volatile organic compound (VOC) emissions, high humidity, and dust from fabric handling. The factory has moderate natural ventilation and operates under multi-shift worker density.

3.6.4. Factory 3: Metal Machining and Moulding Factory

This factory specializes in CNC machining, plastic injection moulding, and metal stamping. These operations produce oil vapors, thermal degradation by-products, and localized particulate matter, especially during tool changes and high-speed operations. The facility uses mechanical exhaust systems with varying ventilation efficiency. Each IoT unit continuously monitored environmental variables including $PM_{2.5}$, PM_{10}, CO_2, CH_2O, CO, O_3, NO_2, temperature, and humidity. Sensor data was stored locally and transmitted over LoRaWAN to a central server for federated aggregation. The sensors operated uninterrupted for more than four months in harsh conditions, capturing over 30,000 timestamped datapoints per site.

3.7. Cascading Fuzzy Federated Aggregation (CFFA)

3.7.1. Rationale and method overview

N-FedAvg [40] is introduced to tackle the main limitation of FedAvg. In FedAvg, client contributions depend only on dataset size, which creates bias when data is unevenly distributed. N-FedAvg assigns equal weights to all clients, reducing this bias. Tests showed that N-FedAvg consistently outperforms FedAvg when data distribution is unbalanced. Building on this, a novel method called Cascading Fuzzy Federated Aggregation (CFFA) is proposed. CFFA goes further by integrating fuzzy logic to adapt aggregation weights based on contextual factors, not just data counts.

The approach uses a two-level fuzzy logic system to optimize the aggregation process in federated learning by considering contextual conditions in the real life. Instead of relying only on the scale of local datasets or presupposing the fair distribution of data, CFFA uses the layered fuzzy inference model to calculate the importance of each client. It accomplishes it by examining such indicators of practical performance as the density of the working population, air quality levels (AQI), and ventilation systems efficiency. CFFA incorporates these dynamic factors of the environment and operations thus introducing more fairness, stability and contextual relevance to the aggregation process- hence the reason why CFFA is best suited when applied in an industrial and high-risk setting.

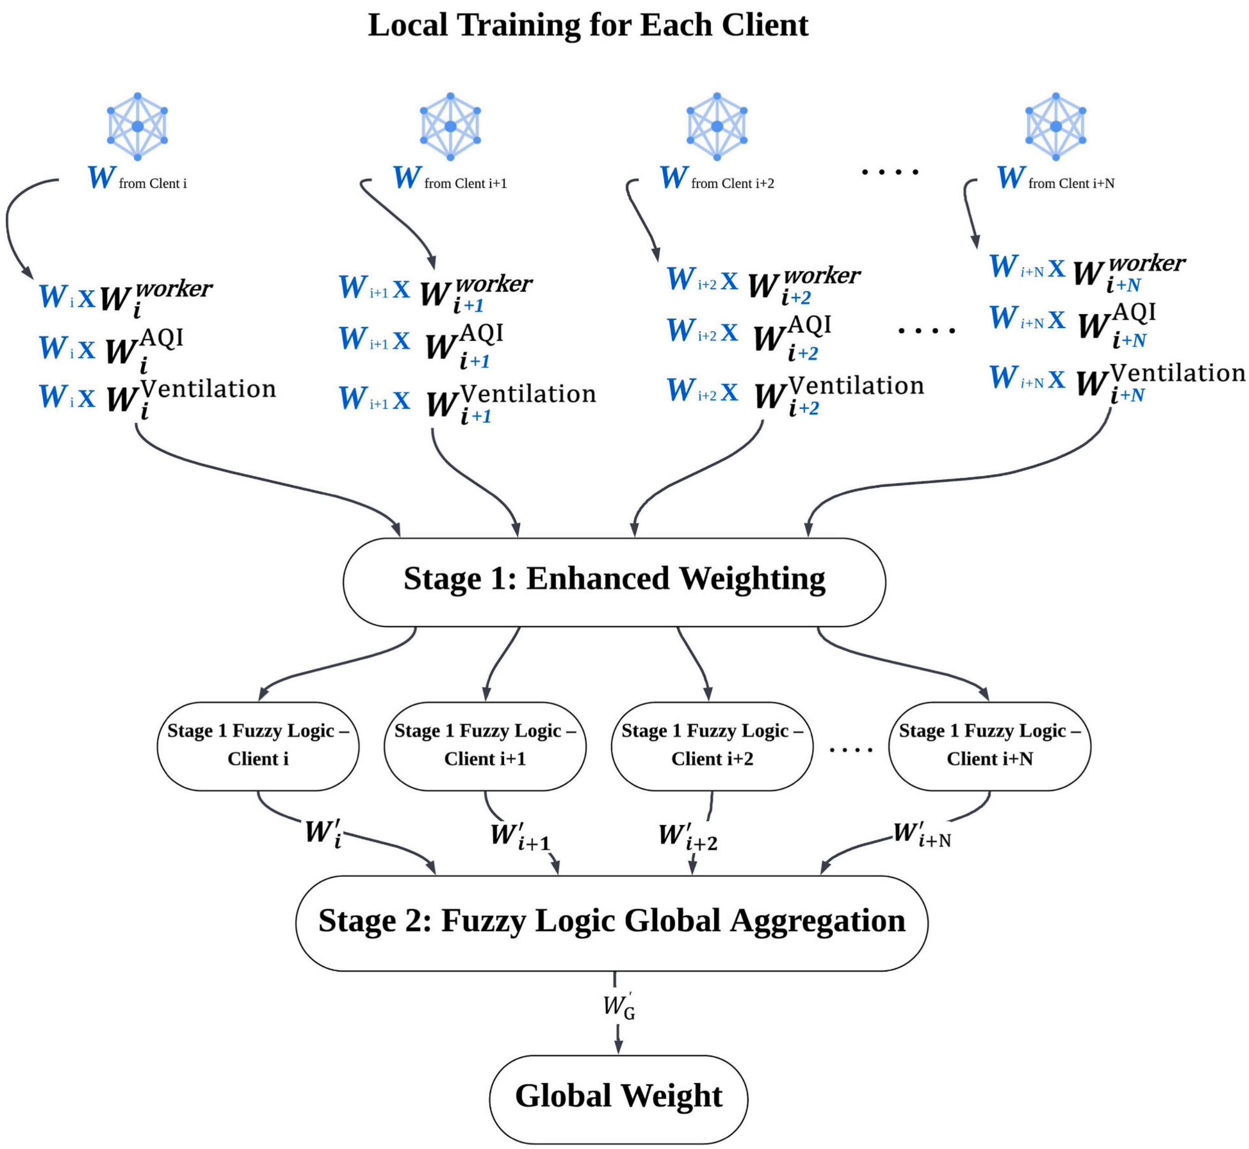

Fig. 9. CFFA block diagram showing the flow from local context to global model aggregation.

3.7.1.1. Two-stage fuzzy aggregation framework. In order to address the limitations of the traditional federated learning methods which are primarily based on the size of the dataset to weight them, we propose a two-stage fuzzy logic framework named Cascading Fuzzy Federated Aggregation (CFFA). This approach introduces the environmental and operational awareness to the aggregation process. During the initial stage, each client will implement its own local version of the fuzzy inference system that will be processing three major inputs: number of workers in the location, the local air quality index (AQI) and the efficiency of the ventilation system, expressed as the rate of air exchange. Such inputs are incorporated into calculating a set of scaled weights which are the operational pressure and sensing conditions of every client. Each of these contextual values is multiplied by the client's raw model weight W_i, resulting in three weighted terms: $\boldsymbol{W_i^{\text{worker}}}$, $\boldsymbol{W_i^{\text{AQI}}}$, $\boldsymbol{W_i^{\text{Ventilation}}}$. These are then fed into the fuzzy system to produce an **Adjusted Client Weight** $\boldsymbol{W_i'}$, which reflects the client's overall contribution trust level.

In Stage 2, the set of adjusted weights $W_i', W_{i+1}', \ldots, W_N'$ from all participating clients is aggregated using a second fuzzy inference system. This system evaluates the collective importance of the inputs and produces a final Global Weight W_G', which is used to scale the aggregated model parameters during the global model update. The entire process is visualized in the block diagram as shown in Fig. 9. It clearly illustrates the flow of raw model weights through contextual enhancement in Stage 1 and global-level fusion in Stage 2, resulting in a trust-aware, adaptive aggregation strategy. This structure improves fairness, robustness, and adaptability, especially in scenarios with heterogeneous sensor reliability or environmental stress.

3.8. Stage 1: Client-level fuzzy evaluation

In the first stage of the CFFA framework, each client uses a local fuzzy inference system to compute an Adjusted Client Weight W_i'. This weight reflects how trustworthy and contextually important the client's model contribution is, based on three inputs:

- Number of workers on-site
- Local air quality index (AQI)
- Ventilation efficiency (air exchange rate)

This phase captures human risk, environmental stress, and infrastructure efficiency through fuzzified input weighting.

3.8.1. Input fuzzification and membership functions

Each input is fuzzified into three linguistic terms using triangular membership functions:

- **Number of Workers (*W*):** e.g., from 0 to 500 workers
- **Air Quality Index (AQI):** e.g., 0 (good) to 500 (hazardous)
- **Ventilation Efficiency (*V*):** e.g., measured as air exchange rate (1–10), higher values mean better ventilation.

Use triangular membership functions (μ) to fuzzify inputs into linguistic variables such as Low, Medium, and High as shown in Figss. 10–12.

Example definitions for each variable:

- **Number of Workers ($\boldsymbol{W_i^{\text{worker}}}$):**
  - ✓ Range: 0 to 500
  - ✓ Linguistic Terms: Low, Medium, High
  - ✓ Membership Functions (Triangular):

$$\mu_{W,\text{low}}(x) = \text{trimf}(x; 0, 0, 250)$$

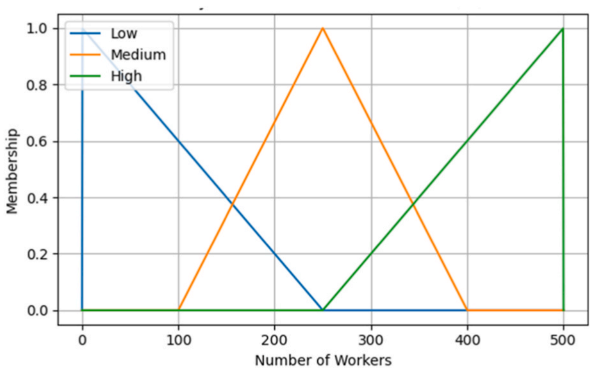

Fig. 10. Fuzzy set for number of workers.

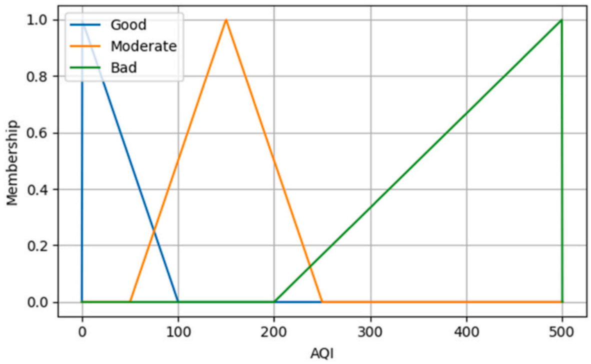

Fig. 11. Fuzzy sets for Air Quality Index.

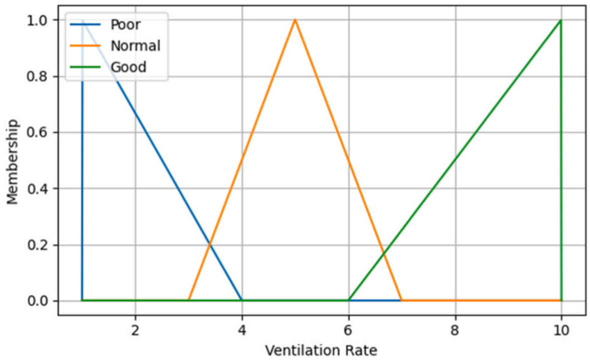

Fig. 12. Fuzzy set of Ventilation Efficiency.

$$\mu_{W,\text{medium}}(x) = \text{trimf}(x; 100, 250, 400)$$

$$\mu_{W,\text{high}}(x) = \text{trimf}(x; 250, 500, 500)$$

- **Air Quality Index (W_i^{AQI}):**
  - ✓ Range: 0 (good) to 500 (hazardous)
  - ✓ Linguistic Terms: Good, Moderate, Bad
  - ✓ Membership Functions:

$$\mu_{AQI,\text{good}}(x) = \text{trimf}(x; 0, 0, 100)$$

$$\mu_{AQI,\text{moderate}}(x) = \text{trimf}(x; 50, 150, 250)$$

$$\mu_{AQI,\text{bad}}(x) = \text{trimf}(x; 200, 500, 500)$$

- **Ventilation ($W_i^{\text{Ventilation}}$):**
  - ✓ Range: 1 to 10 (air exchange rate)
  - ✓ Linguistic Terms: Poor, Normal, Good
  - ✓ Membership Functions:

$$\mu_{V,\text{poor}}(x) = \text{trimf}(x; 1, 1, 4)$$

$$\mu_{V,\text{normal}}(x) = \text{trimf}(x; 3, 5, 7)$$

$$\mu_{V,\text{good}}(x) = \text{trimf}(x; 6, 10, 10)$$

For each input, we first multiply the raw model weight W_i by the corresponding fuzzified contextual value, producing three intermediate weighted inputs:

- $W_i^{worker} = W_i \times Worker_i$
- $W_i^{AQI} = W_i \times AQI_i$
- $W_i^{Vent} = W_i \times Vent_i$

These weighted terms are the actual inputs to the fuzzy rule system.

3.8.2. Output weighting and rule set

Output = A fuzzy Adjusted Factory Weight ∈ [0, 1] for each factory:

$$\boldsymbol{W_i'} \quad (\textbf{AdjustedClientWeight}) = \{\boldsymbol{Low}, \boldsymbol{Medium}, \boldsymbol{High}\}$$

The output for each factory's fuzzy system is Adjusted Factory Weight:

- **Low**: 0.0–0.3
- **Medium**: 0.2–0.7
- **High**: 0.6–1.0

$$\mu_{\text{low}}(x) = \text{trimf}(x; 0, 0, 0.3)$$

$$\mu_{\text{medium}}(x) = \text{trimf}(x; 0.2, 0.45, 0.7)$$

$$\mu_{\text{high}}(x) = \text{trimf}(x; 0.6, 1, 1)$$

The corresponding fuzzy output set for Stage 1 is illustrated in Fig. 13.

3.8.2.1. Define fuzzy rules.

The fuzzy system uses a rule base to combine the three weighted inputs and infer the output $\boldsymbol{W_i'}$. These rules apply universally to all clients. Table 7 presents the rules for computing adjusted client weight.

3.9. Stage 2: Global fuzzy aggregation

In the second stage of the **CFFA framework**, the adjusted client weights from Stage 1 are aggregated using a global fuzzy inference system. The purpose of Stage 2 is to determine a single **Global Aggregation Weight** $\boldsymbol{W_G'}$, reflecting the collective reliability and trustworthiness of all participating clients. This final weight scales the aggregated model parameters, ensuring global model updates are context-aware and balanced.

3.9.1. Aggregation input processing

Stage 2 takes as input the fuzzy outputs generated by Stage 1, specifically the Adjusted Client Weights ($\boldsymbol{W_i'}$) for each factory:

- $\boldsymbol{W_i'}$: AdjustedClientWeight for Factory 1 ∈ [0, 1]
- $\boldsymbol{W_{i+1}'}$: AdjustedClientWeight for Factory 2 ∈ [0, 1]
- $\boldsymbol{W_{i+2}'}$: AdjustedClientWeight for Factory 3 ∈ [0, 1]

Each Adjusted Client Weight has the following properties:

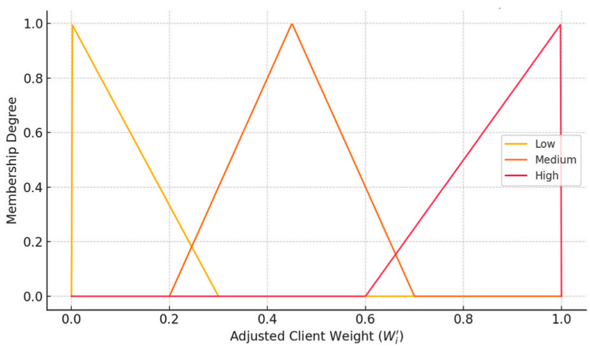

Fig. 13. Output fuzzy set of Stage 1.

Table 7
Fuzzy rules for computing adjusted client weight W_i'

Rule No.	W_i^{worker}	W_i^{AQI}	$W_i^{\text{Ventilation}}$	W_i'
1	L	G	P	L
2	L	G	N	L
3	L	G	G	M
4	L	Mod	P	L
5	L	Mod	N	L
6	L	Mod	G	M
7	L	B	P	M
8	L	B	N	M
9	L	B	G	M
10	M	G	P	L
11	M	G	N	M
12	M	G	G	M
13	M	Mod	P	M
14	M	Mod	N	M
15	M	Mod	G	H
16	M	B	P	M
17	M	B	N	H
18	M	B	G	H
19	H	G	P	M
20	H	G	N	M
21	H	G	G	H
22	H	Mod	P	M
23	H	Mod	N	H
24	H	Mod	G	H
25	H	B	P	H
26	H	B	N	H
27	H	B	G	H

Note: W_i^{worker} is Client's worker-adjusted input weight, W_i^{AQI} is Client's AQI-adjusted input weight, $W_i^{\text{Ventilation}}$ is Client's ventilation-adjusted input weight, and W_i' is Adjusted Client Weight. L = Low, M = Medium, H = High, G = Good, Mod = Moderate, B = Bad, P = Poor, and N = Normal.

✓ Range: 0 to 1
✓ Linguistic Terms: Low, Medium, High
✓ Membership Functions (Triangular)**:**
 o Low: 0.0–0.3
 o Medium: 0.2–0.7
 o High: 0.6–1.0

$$\mu_{\text{low}}(x) = \text{trimf}(x; 0, 0, 0.3)$$

$$\mu_{\text{medium}}(x) = \text{trimf}(x; 0.2, 0.45, 0.7)$$

$$\mu_{\text{high}}(x) = \text{trimf}(x; 0.6, 1, 1)$$

3.9.2. Output interpretation

The output of Stage 2 is the **Final Global Aggregation Weight W_G'**, used for scaling the federated learning global model parameters. This global weight is defined similarly as:

- W_G': Final Aggregated Weight (Global Model Scaling) $\in [0, 1]$
- **Linguistic Terms**: Low, Medium, High
- **Low**: 0.0–0.3
- **Medium**: 0.2–0.7
- **High**: 0.6–1.0
- The membership functions are defined as:

$$\mu_{\text{low}}(x) = \text{trimf}(x; 0, 0, 0.3)$$

$$\mu_{\text{medium}}(x) = \text{trimf}(x; 0.2, 0.45, 0.7)$$

$$\mu_{\text{high}}(x) = \text{trimf}(x; 0.6, 1, 1)$$

Fig. 14 shows the fuzzy output set used in Stage 2. This global fuzzy output ensures the aggregated model reflects a balanced assessment of individual client trustworthiness.

This global fuzzy output ensures the aggregated model reflects a balanced assessment of individual client trustworthiness.

3.9.2.1. Fuzzy rule base for global aggregation. Stage 2 uses a rule base to determine the Final Global Aggregation Weight $\boldsymbol{W}_G'$ by combining the three adjusted client weights from Stage 1Each rule assesses the reliability of a participating client and assigns it a corresponding linguistic term that reflects its contribution to the global aggregation. Table 8 lists the fuzzy rules applied in stage 2 to convert the adjusted client weights into final global weights, while Fig. 15 illustrates how stage 2 merges the three adjusted client weights to compute a unified global aggregation weight.

4. Results and discussion

This part describes an experimental validation of the Cascading Fuzzy Federated Aggregation (CFFA) framework in a scenario of realistic industrial environments. The assessment shows how CFFA is adapting to the volatile operational and environmental conditions across the factory locations and how the adaptation influence the performance of the global model. We look at the quality and quantity of the data gathered, evaluate local and aggregated model accuracies in five distinct scenarios and contrast the performance of CFFA with the performance of the standard aggregation approaches to show its effectiveness in practice.

4.1. Data collection in industrial environments

For a period of four consecutive months, three IoT-based air quality monitoring units were deployed at separate industrial locations to collect continuous time-series environmental data. Each device captured sensor readings at 5-min intervals, monitoring critical air quality variables including $PM_{2.5}$, PM_{10}, CO_2, CO, O_3, NO_2, temperature, and humidity. This deployment was designed to mimic actual industrial conditions, providing realistic data to support and evaluate the proposed federated learning model.

The IoT sensors were installed at the following industrial sites:

- Chemical Metal Galvanizing Facility – This location presented high levels of gas emissions due to the use of acidic substances in the galvanization process.
- Textile Manufacturing Plant – This site was marked by elevated airborne particles resulting from material handling, fabric processing, and dyeing operations.
- Metal Machining and Moulding Factory: Involving metal cutting, shaping, and casting, generating both particulate and gaseous pollutants.

All three units were configured with identical sampling intervals (5 min) and operated continuously, with data stored locally on microSD cards. Table 9 summarizes the total amount of data collected from each site.

This dataset, totaling over 100,000 samples across all factories, served as the basis for local LSTM model training and subsequent federated learning aggregation. The size and consistency of the dataset across different industrial processes enabled robust evaluation of the proposed CFFA method under varied environmental conditions.

4.2. Onboard data preprocessing

To ensure the quality and reliability of the collected sensor data, each IoT unit performs onboard data preprocessing before storage. Every incoming sensor value is first validated against acceptable operating ranges predefined for each gas and environmental parameter. If a value is determined to be within the valid range, it is stored directly on the internal microSD card. However, if the reading falls outside the valid range due to sensor error, interference, or environmental anomalies the unit automatically replaces it

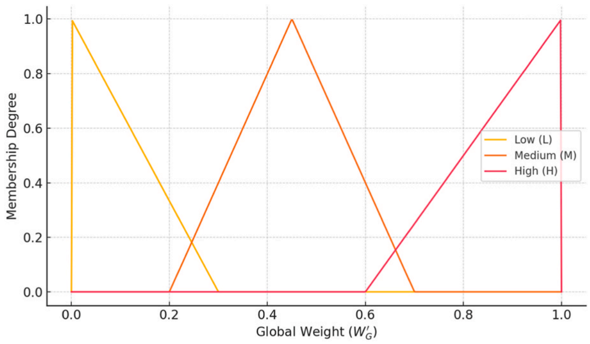

Fig. 14. Fuzzy output set of stage 2.

Table 8
Fuzzy rules used in Stage 2 for converting adjusted client weights into the final global weight W'_G

Rule No.	W'_i	W'_{i+1}	W'_{i+2}	W'_G
1	L	L	L	L
2	L	L	M	L
3	L	L	H	M
4	L	M	L	L
5	L	M	M	M
6	L	M	H	M
7	L	H	L	M
8	L	H	M	M
9	L	H	H	H
10	M	L	L	L
11	M	L	M	M
12	M	L	H	M
13	M	M	L	M
14	M	M	M	M
15	M	M	H	H
16	M	H	L	M
17	M	H	M	H
18	M	H	H	H
19	H	L	L	M
20	H	L	M	M
21	H	L	H	H
22	H	M	L	M
23	H	M	M	H
24	H	M	H	H
25	H	H	L	H
26	H	H	M	H
27	H	H	H	H

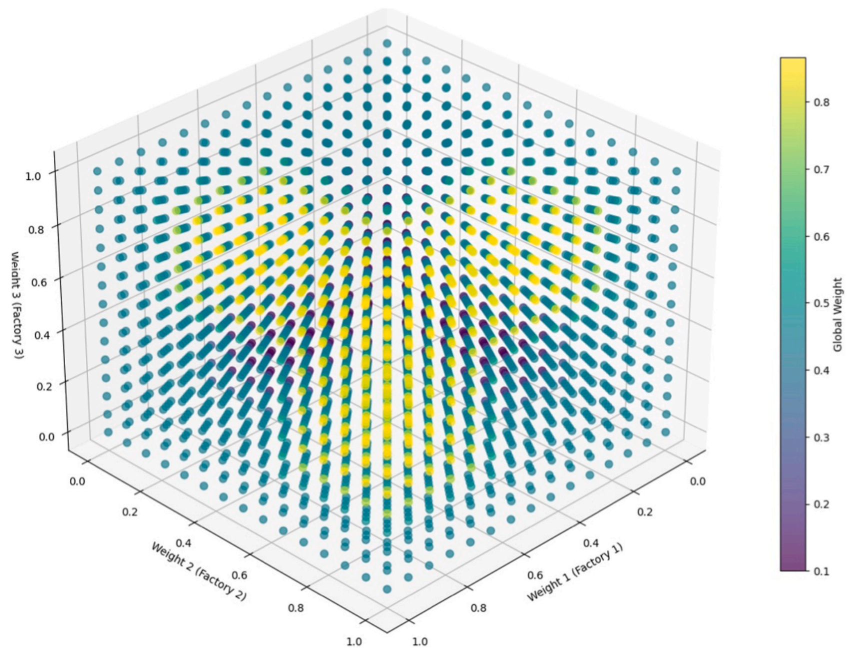

Fig. 15. Illustration of Stage 2 fuzzy aggregation: merging three adjusted client weights to compute a unified global aggregation weight.

with the average of the last three valid readings. This simple method of filtering will address sharp spikes and false readings adequately to cut down noise even before training local models. The whole preprocessing procedure is directly incorporated into the firmware of the microcontroller allowing to clean the data in real-time at the device level without cloud-based processing.

Table 9
Total data samples collected from each factory over 4 Months.

Factory	Sampling Interval (minutes)	Days of Deployment	Samples Collected per Day	Total Samples Collected (4 months)
Chemical Metal Galvanizing Plant	5	120	288	34,560
Textile Manufacturing Facility	5	120	288	34,560
Metal Machining and Moulding Factory	5	120	288	34,560

4.3. Air pollutant trends across factories

The trends in concentration of pollutants obtained at the three industrial sites namely the Galvanizing Plant, Textile Factory and Moulding Facility were quite different depending on the operation of the sites and the various environmental conditions at the sites as shown in Fig. 16. The level of carbon dioxide (CO_2) was not the same in all the places but the Textile Factory had the highest level and, in most cases, it exceeded 1000 ppm. This trend has probably been associated with overcrowding of people and lack of ventilation systems. The Moulding Factory had more favorable and reduced CO_2, whereas the Galvanizing Plant had erratic spikes, which may be related to certain processes that entail heat or chemical treatment during peak shifts. The same occurred in the levels of particulate matter, especially PM2.5 and PM10. The most severe levels were demonstrated in the Galvanizing Plant, as PM2.5 was in some cases above 800ug/m3, and PM10 was up to 1000ug/m 3, probably due to the regular release of the tasks like sanding, dipping, or grinding. Comparatively, the Moulding Factory had moderate levels of particulate matters but the Textile Factory was relatively clean and there was minimal change in the airborne particles.

The same trend was consistent in the formaldehyde (CH_2O) scoring. The Galvanizing Plant once more had the highest readings greater than 1.5 mg/m 3 because of chemical reactions that involve coating and surface treatment. The Moulding and Textile Factories were found to have very low levels of CH_2O of less than 0.3 mg/m 3 which showed that there is no emanation of aldehydes in those sites. Data on ozone (O_3) also showed characteristic environmental patterns of the factories. The Galvanizing Plant also had a lot of very spiky readings, some as much as 0.4 ppm, possibly due to high temperature or electrochemical processes. The Moulding Facility had moderate ozone levels and the Textile Factory had almost zero readings almost all the time which shows that the environment inside the Textile Factory is stable with low reactive emissions.

In contrast, the level of carbon monoxide (CO) was very high at the Moulding Factory and it frequently exceeded 2 ppm. This could have been as a result of combustion sources of equipment located on site or the heating system. The Spikes of the Galvanizing Plant were higher, yet the average level of CO was lower, and the Textile Factory was the most stable with the lowest CO results. Total volatile organic compounds (TVOCs) were also different. The Textile Factory recorded the most TVOC levels in all categories, most probably as a result of frequent applications of chemicals during the dyeing and washing exercises. At the Galvanizing and Moulding Factories, the peaks were more sporadic and sloppier, which might have been related to ventilation periods or occasional releases. The highest and most stable concentration of nitrogen dioxide (NO_2) was recorded in the Textile Factory with an average of about 0.3 ppm- a fact that may be attributed to steam systems or heating with gas. Galvanizing Plant recorded spikes of high NO_2 emissions whereas the Moulding Facility recorded the lowest and the most stable levels during the monitoring period. The overall trend explains the great disparity in the level of environmental exposure and working conditions in the various industrial plants. This is further evidence of the necessity of context-sensitive approaches such as the CFFA framework, where federated learning can be customized to local realities to better perform aggregation of models in a more accurate and fair manner.

4.4. LSTM-based air quality prediction results

The LSTM model established strong short-term prediction capabilities across most pollutants, with outputs closely following the observed trends in several cases. Fig. 17 compares predicted and actual concentrations for CO_2, PM2.5, PM10, CH_2O, O_3, CO, TVOC, and NO_2 over 300-time steps. The model showed high alignment for CO_2, PM2.5, and PM10, with predicted curves closely tracking real values, including peak formations and drop intervals. These pollutants are typically stable and abundant in the dataset, which may have helped the model generalize temporal patterns effectively. For CH_2O, the LSTM model also performed well. It captured both gradual increases and sharp spikes with good temporal alignment, though there was slight underestimation during some peaks. This indicates the model's ability to generalize nonlinear chemical concentration changes.

Prediction accuracy was temperately strong for NO_2, although the model occasionally lagged behind sharp transitions. It followed the general shape of the time series but missed some of the finer oscillations, especially during periods of rapid change. The prediction performance for O_3 was less consistent. The model followed broad trends but failed to capture sharp fluctuations and high-frequency variability. This could be due to the irregular and sparse nature of ozone concentration data or its sensitivity to external factors not captured by input features. CO showed poor predictive performance. The actual values were consistently zero, while the model predicted a constant small non-zero output across all time steps. This suggests a mismatch or data quality issue, likely caused by underrepresentation or faulty measurements in the training data. TVOC predictions followed the general binary grade pattern (0–3), but with noticeable smoothing. The model captured the rising trends and some of the peaks, but failed to replicate sharp vertical transitions seen in the actual data. This is likely due to the categorical/discrete nature of TVOC sensor output, which may not be well-suited to regression-based LSTM prediction.

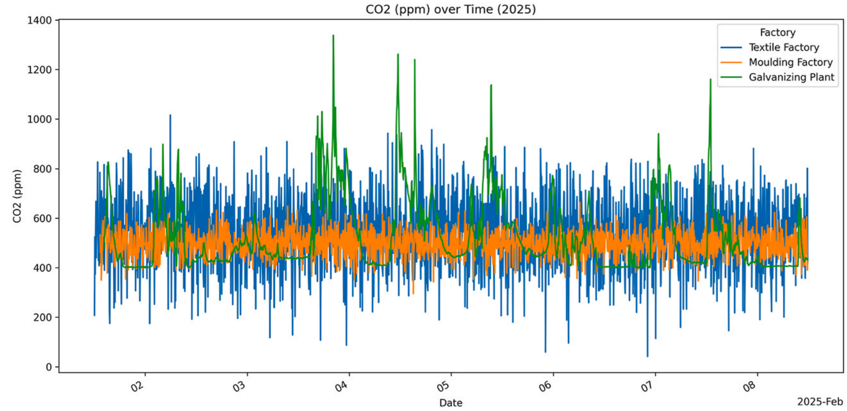

(a)

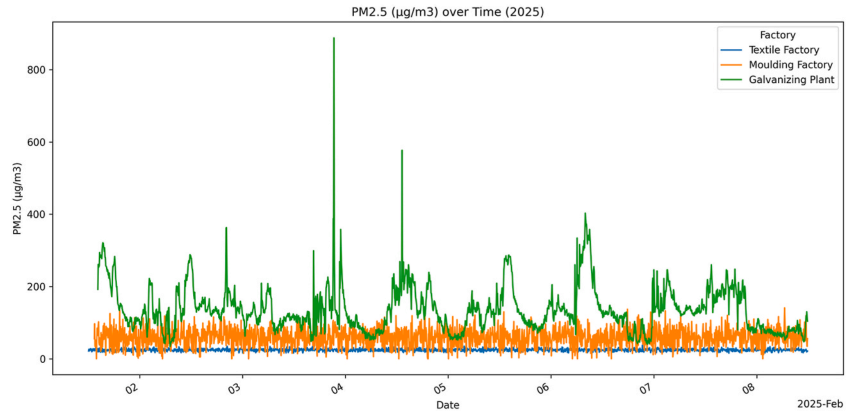

(b)

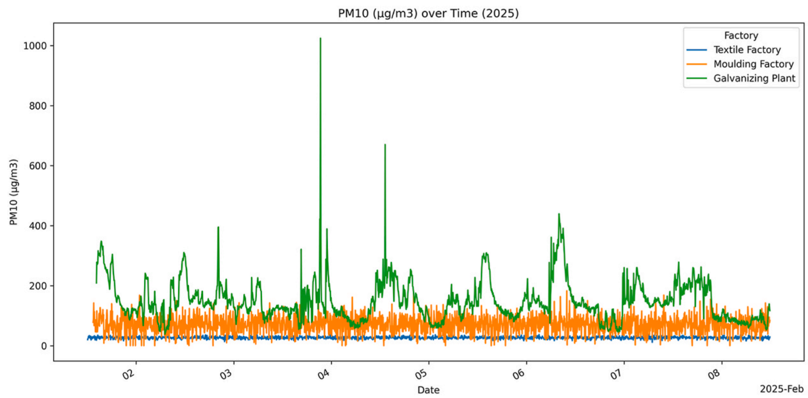

(c)

Fig. 16. Time-series trends of air pollutant concentrations across three factories: (a) CO_2 (ppm), (b) PM2.5 (μg/m^3), (c) PM10 (μg/m^3), (d) CH_2O (mg/m^3), (e) O_3 (ppm), (f) CO (ppm), (g) TVOC (grade), (h) NO_2 (ppm).

In summary, the LSTM model achieved reliable prediction performance for major pollutants (CO_2, PM2.5, PM10, CH_2O), moderate accuracy for NO_2 and O_3, and weak performance on CO and TVOC due to either flat inputs, data imbalance, or sensor quantization. These results highlight both the strengths and the current limitations of applying time-series deep learning to heterogeneous air quality data in industrial settings.

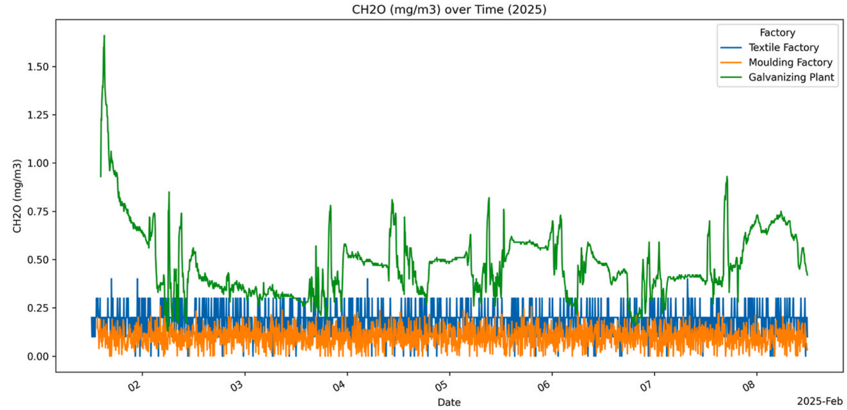

(d)

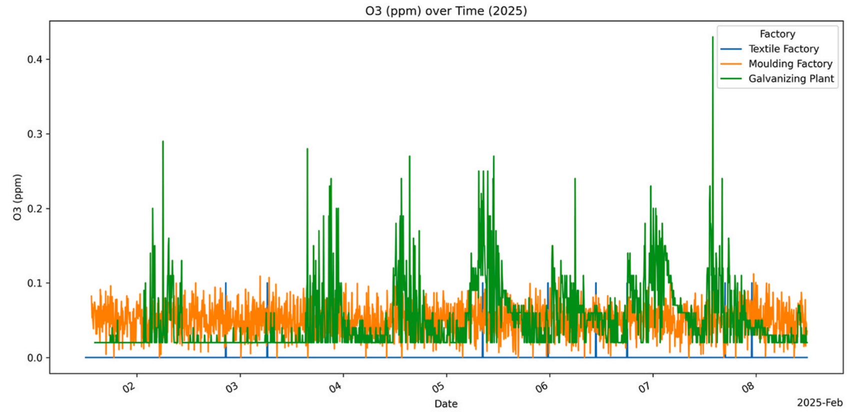

(e)

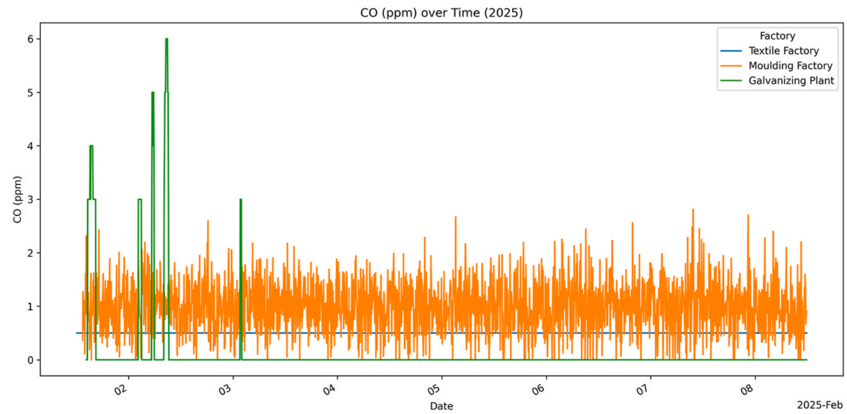

(f)

Fig. 16. (*continued*).

4.5. Federated learning with FedAvg and N-FedAvg

4.5.1. Experimental setup with balanced and unbalanced datasets

FL enables factories to collaboratively train a shared model without sharing sensitive data. However, FL performance is closely tied to the distribution of data across clients. When datasets are balanced, traditional aggregation methods like FedAvg work well. In contrast, when datasets are unbalanced a common scenario in real-world industrial settings FedAvg tends to favor the factory with the largest dataset, leading to biased global models. N-FedAvg addresses this problem by normalizing client contributions, giving equal weight to all participants. To explore this, experiments were conducted across three factory types: Moulding Factory, Textile Factory

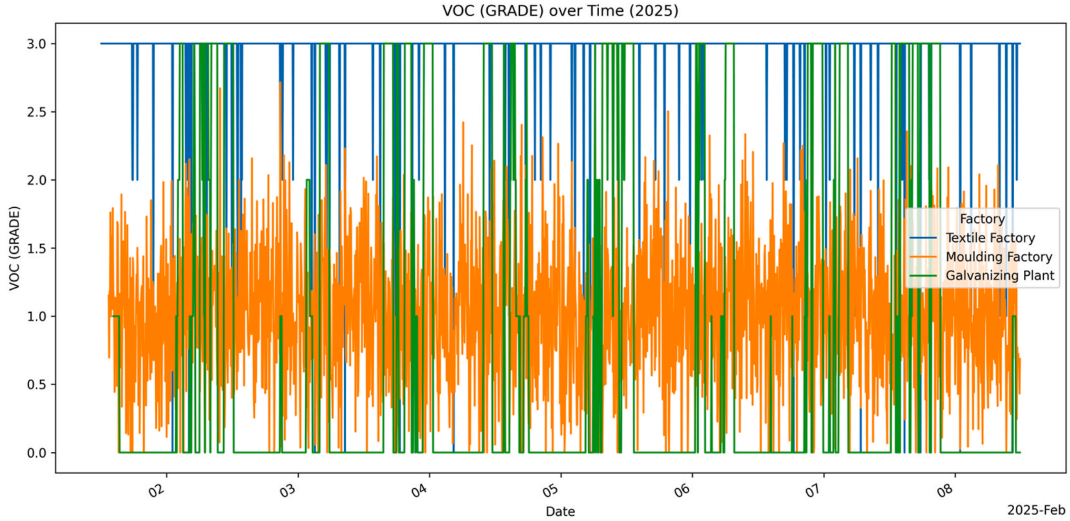

(g)

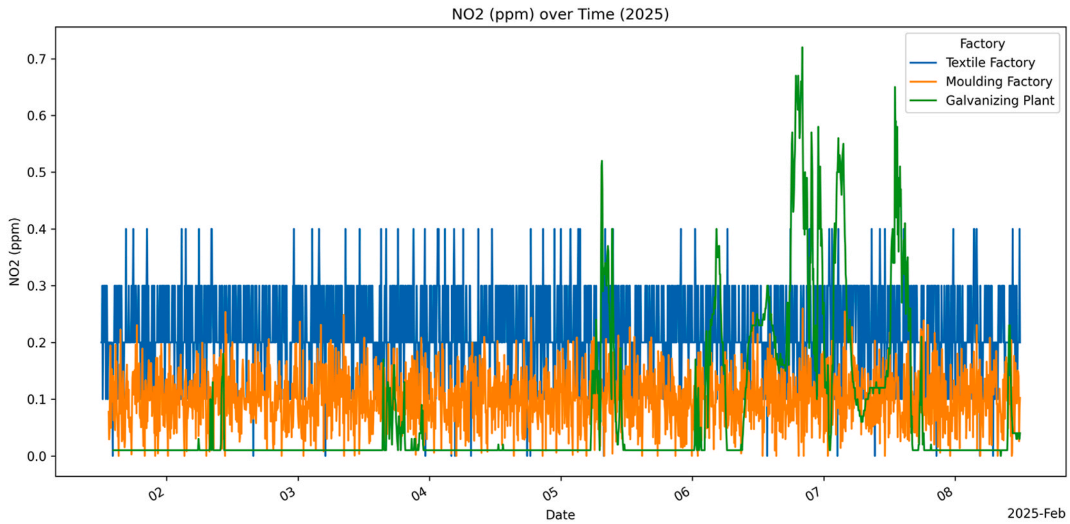

(h)

Fig. 16. (*continued*).

and Galvanizing Factory. We varied dataset sizes while maintaining the total number of samples at 900,000. Two scenarios were tested: balanced datasets, where data size is the same for all factories, and unbalanced datasets, where data sizes differ. The experimental results are summarized in Tables 10 and 11.

4.5.2. Discussion of results

In balanced scenarios (Table 7), both FedAvg and N-FedAvg achieved almost identical global accuracies. For example, in Case 5, both methods achieved 92.4 % global accuracy. This confirms that under ideal, balanced conditions, both aggregation strategies are equally effective. However, the situation changes significantly under unbalanced settings (Table 8). In Case 1 of the unbalanced dataset (300k Moulding Factory, 100k Textile Factory, 500k Galvanizing Factory), FedAvg's global accuracy dropped to 87.7 %, while N-FedAvg achieved 90.5 %. Similarly, across all unbalanced cases, N-FedAvg consistently produced global accuracies closer to the balanced case results, while FedAvg suffered noticeable drops.

The detailed breakdown per factory also shows that FedAvg tends to prioritize whichever factory has the largest dataset, compromising model fairness. In contrast, N-FedAvg maintained much closer, consistent performance across all factories, proving its resilience against imbalance.

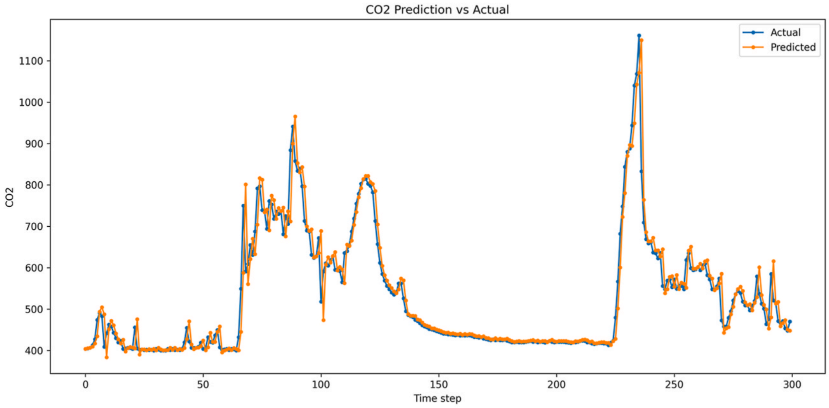

(a)

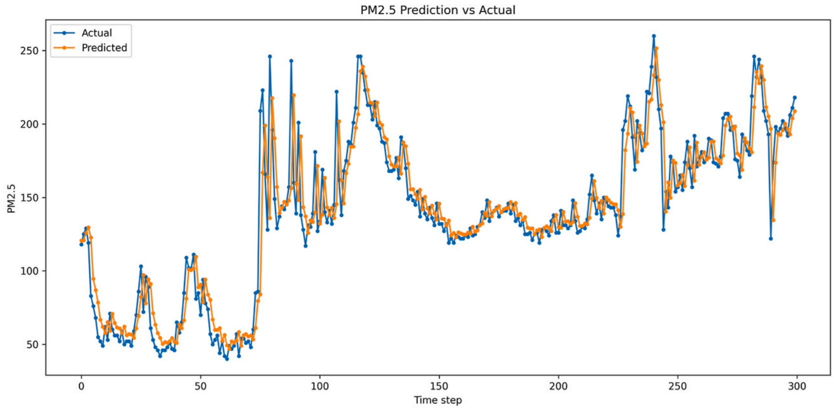

(b)

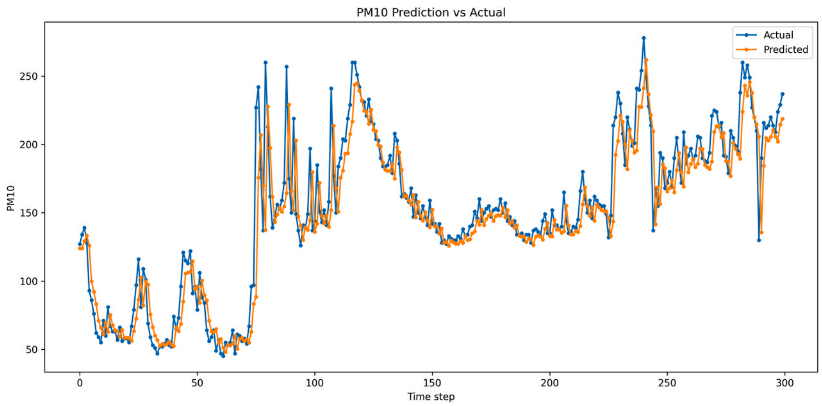

(c)

Fig. 17. Prediction results of the LSTM model versus actual measurements for each air pollutant: (a) CO_2, (b) PM2.5, (c) PM10, (d) CH_2O, (e) O_3, (f) CO, (g) TVOC, and (h) NO_2.

4.5.3. Statistical validation with hypothesis testing

To statistically justify the superiority of N-FedAvg over FedAvg in unbalanced datasets, a paired *t*-test was steered on the universal accuracies of FedAvg and N-FedAvg across the six unbalanced cases.

- Null Hypothesis (H_0): There is no significant difference in performance between FedAvg and N-FedAvg in unbalanced settings.

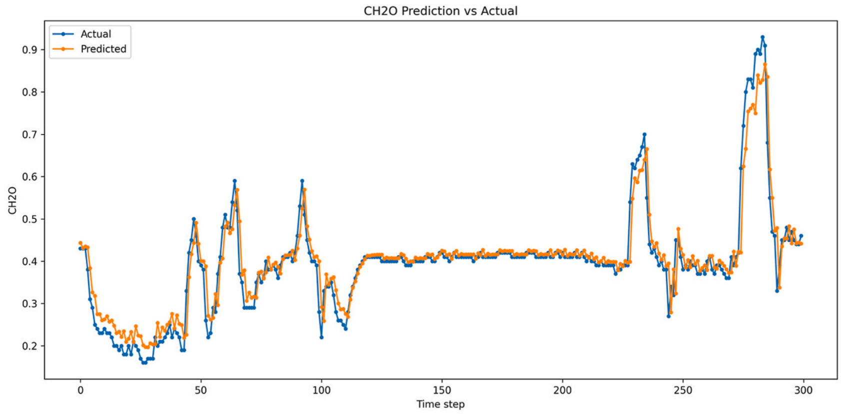

(d)

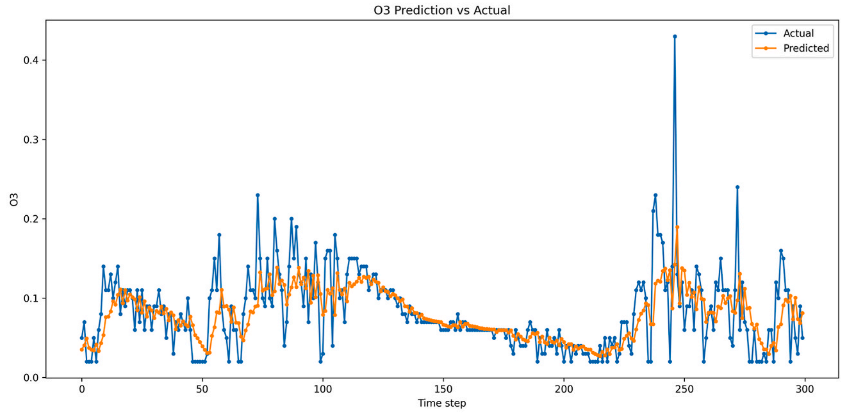

(e)

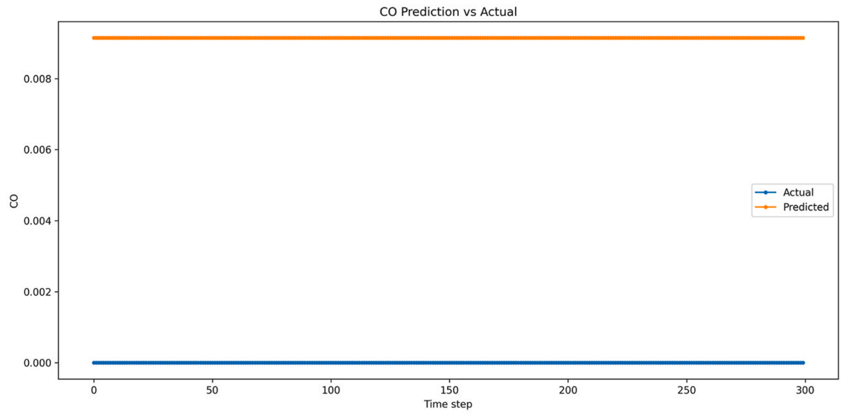

(f)

Fig. 17. (*continued*).

- Alternative Hypothesis (H_1): N-FedAvg notably outperforms FedAvg under unbalanced conditions.

The *t*-test yielded a p-value <0.05, leading to the rejection of the null hypothesis. This indicates that the observed improvement using N-FedAvg is statistically significant and not due to random variation.

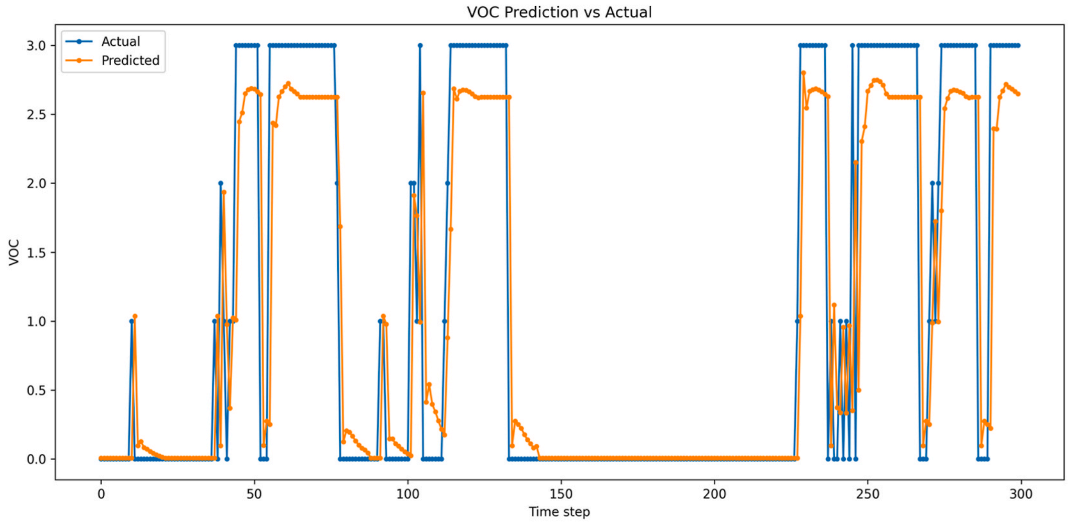

(g)

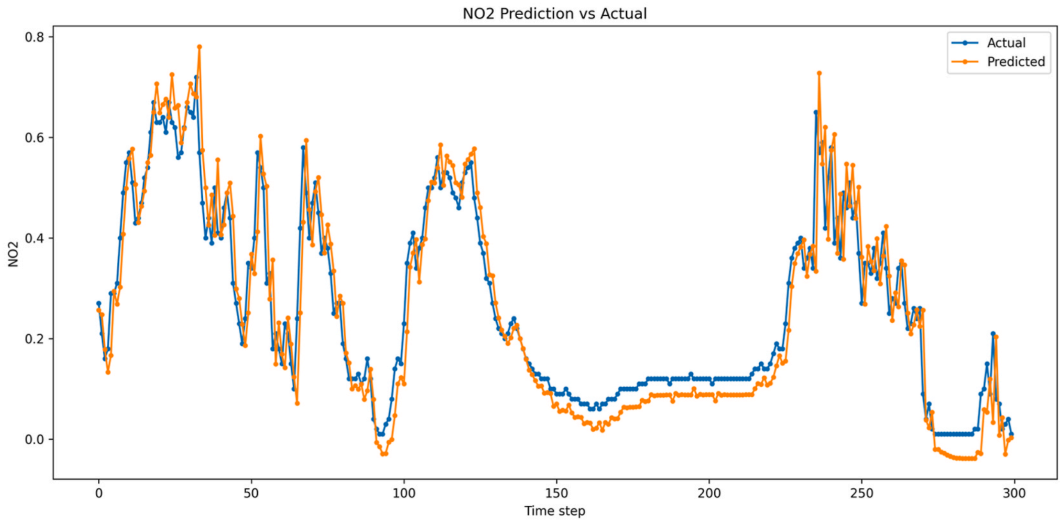

(h)

Fig. 17. (*continued*).

Table 10
Effect of dataset distribution on FL model accuracy (balanced dataset).

Category	Source	Case 1	Case 2	Case 3	Case 4	Case 5
Dataset Size	Moulding Factory	100,000	200,000	300,000	400,000	500,000
	Textile Factory	100,000	200,000	300,000	400,000	500,000
	Galvanizing Factory	100,000	200,000	300,000	400,000	500,000
Accuracy Without FL	Moulding Factory	90 %	91 %	91.7 %	92.3 %	93 %
	Textile Factory	89 %	90 %	90.8 %	91.2 %	91.8 %
	Galvanizing Factory	87 %	88 %	88.8 %	89.2 %	89.9 %
FedAvg Global Accuracy	FedAvg Global	89.2 %	91.2 %	91.5 %	91.9 %	92.4 %
N-FedAvg Global Accuracy	N-FedAvg Global	89.2 %	91.3 %	91.5 %	91.8 %	92.4 %
Accuracy (FedAvg)	Moulding Factory	89 %	90.7 %	91.1 %	91.6 %	92.3 %
	Textile Factory	89 %	91.1 %	91.1 %	91.6 %	92.1 %
	Galvanizing Factory	89.1 %	90.7 %	91.2 %	91.7 %	92.3 %
Accuracy (N-FedAvg)	Moulding Factory	89 %	91.1 %	91.3 %	91.5 %	92.3 %
	Textile Factory	88.9 %	91.1 %	91.2 %	91.6 %	92.2 %
	Galvanizing Factory	88.8 %	91.1 %	91.2 %	91.5 %	92.4 %

Table 11
Effect of dataset distribution on FL model accuracy (unbalanced dataset).

Category	Source	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
Dataset Size	Moulding Factory	300k	300k	100k	500k	100k	500k
	Textile Factory	100k	500k	300k	300k	500k	100k
	Galvanizing Factory	500k	100k	500k	100k	300k	300k
Accuracy Without FL	Moulding Factory	91.7 %	91.7 %	89.5 %	93 %	90 %	93 %
	Textile Factory	89 %	91.8 %	90.8 %	90.8 %	91.8 %	89 %
	Galvanizing Factory	89.9 %	87 %	89.9 %	87 %	88.8 %	88.8 %
FedAvg Global Accuracy	FedAvg Global	87.7 %	91.2 %	89.1 %	89.8 %	90.7 %	90.8 %
N-FedAvg Global Accuracy	N-FedAvg Global	90.5 %	91.6 %	90.7 %	91.5 %	91.3 %	91.2 %
Accuracy (FedAvg)	Moulding Factory	87.5 %	91 %	88.9 %	90 %	90.5 %	91 %
	Textile Factory	87.5 %	91.4 %	88.9 %	89.6 %	90.9 %	90.6 %
	Galvanizing Factory	87.9 %	91 %	89.3 %	89.6 %	90.5 %	90.6 %
Accuracy (N-FedAvg)	Moulding Factory	90.3 %	91.5 %	90.5 %	91.3 %	91.1 %	91.2 %
	Textile Factory	90.3 %	91.5 %	90.5 %	91.3 %	91.1 %	91.1 %
	Galvanizing Factory	90.3 %	91.5 %	90.5 %	91.3 %	91.1 %	91.1 %

4.6. *Proposed Cascading Fuzzy Federated Aggregation (CFFA)*

4.6.1. *Scenario-based evaluation setup*

To assess the effectiveness of the Cascading Fuzzy Federated Aggregation (CFFA) approach, we developed five test scenarios that reflect realistic changes in factory environments. Each scenario involves three distinct clients the Moulding Factory, Textile Factory, and Galvanizing Factory each operating under unique combinations of worker population, air quality index (AQI), and ventilation efficiency. The purpose of these tests is to examine how CFFA dynamically adjusts each client's contribution to the global model based on its specific local conditions, and how this adaptation impacts the overall model accuracy.

4.6.2. *Aggregated accuracy using CFFA*

Table 12 delineates the operational variables and corresponding local model accuracies for all five scenarios. The accuracy of each local model remained constant throughout the evaluation: 93.0 % for the Moulding Factory, 91.8 % for the Textile Factory, and 89.9 % for the Galvanizing Factory. This setup isolates the impact of the aggregation method itself by removing the influence of varying local accuracy. Using these operational inputs, the CFFA method was applied to compute a global model accuracy per scenario. Stage 1 generated trust-adjusted weights for each client using fuzzy rules based on workload and environmental stress. These weights were then aggregated in Stage 2 using a second fuzzy logic layer to determine the global weight applied during model update.

The results are shown in Table 13. CFFA achieved a global accuracy of 90.40 % in Case 1, where the Galvanizing plant client operated under extreme stress (500 workers, AQI 500, ventilation rate 1), lowering its influence in the aggregation. As the scenarios shifted toward more balanced or favorable input conditions, the accuracy improved. In Case 5, where conditions were optimal for Moulding factory and Textile factory and less severe for Galvanizing factory, the CFFA method achieved its highest global accuracy of 92.57 %.

4.6.3. *Comparative evaluation against FedAvg and N-FedAvg*

To evaluate the benefit of CFFA over standard aggregation strategies, we selected Case 5 because it reflects a full dataset training scenario for all participating factories. In this case, each client contributed its entire dataset under diverse but realistic operational conditions. As shown in Table 14, both FedAvg and N-FedAvg achieved a global accuracy of 92.4 %, whereas CFFA slightly

Table 12
Operational inputs used for fuzzy aggregation in five evaluation scenarios.

Case	Factory	Workers	AQI	Ventilation Rate	Local Accuracy (%)	Local MAE	Local RMSE
1	Moulding Factory	100	120	9	93	0.070	0.265
	Textile Factory	50	150	9	91.8	0.082	0.286
	Galvanizing Factory	500	500	1	89.9	0.101	0.318
2	Moulding Factory	180	140	7	93	0.070	0.265
	Textile Factory	230	180	4	91.8	0.082	0.286
	Galvanizing Factory	380	250	2	89.9	0.101	0.318
3	Moulding Factory	200	150	6	93	0.070	0.265
	Textile Factory	260	200	4	91.8	0.082	0.286
	Galvanizing Factory	400	280	2	89.9	0.101	0.318
4	Moulding Factory	220	170	6	93	0.070	0.265
	Textile Factory	280	210	3	91.8	0.082	0.286
	Galvanizing Factory	420	300	1	89.9	0.101	0.318
5	Moulding Factory	500	500	1	93	0.070	0.265
	Textile Factory	30	50	10	91.8	0.082	0.286
	Galvanizing Factory	50	120	10	89.9	0.101	0.318

Table 13
Aggregated accuracy results using CFFA across five scenarios.

Scenario	CFFA Accuracy (%)	MAE	RMSE
Case 1	90.40	0.096	0.31
Case 2	91.07	0.089	0.299
Case 3	91.14	0.089	0.298
Case 4	91.27	0.087	0.295
Case 5	92.57	0.074	0.273

Table 14
Comparison of global model accuracy across aggregation methods.

Scenario	Method	Accuracy %	MAE	RMSE
Case 5 - Full dataset training	FedAvg	92.4	0.076	0.276
	N-FedAvg	92.4	0.076	0.276
	CFFA	92.57	0.074	0.273

outperformed them with an accuracy of 92.57 %.

While the numerical gain is modest, the value of CFFA lies in its fairness and realism. FedAvg assigns weight purely based on dataset size, and N-FedAvg uses uniform weighting. In contrast to the traditional practices, CFFA considers the most important operational aspects, including the workforce load, ambient air conditions and ventilation efficiency in calculating the client contributions. This enables the system to minimize weight of the clients who are currently working under stressful or unreliable environments despite having large datasets. Consequently, CFFA is particularly useful in industrial, high-risk environments, whose variable environmental conditions may heavily influence the quality of obtained data and the reliability of the global model.

5. Conclusion

This paper introduced Cascading Fuzzy Federated Aggregation (CFFA), a context-aware federated learning framework for accurate air quality prediction in industrial environments. The approach integrates multi-sensor IoT units, edge-deployed LSTM models, and a two-stage fuzzy logic system that adjusts client contributions based on local operational context, including AQI levels, worker count, and ventilation conditions. Real-world deployments in three Istanbul factories—galvanizing, textile, and metal moulding—generated over 100,000 time-series samples. The LSTM models delivered reliable predictions for major pollutants such as CO_2, PM2.5, PM10, and CH_2O, with moderate performance on NO_2 and lower accuracy for CO and TVOC due to sensor limitations and data imbalance. Comparative analysis showed that N-FedAvg improved model fairness over FedAvg under unbalanced data distributions. However, CFFA outperformed both by producing more stable and realistic global model updates. In a full dataset training scenario, CFFA achieved 92.57 % accuracy, exceeding FedAvg and N-FedAvg, both of which reached 92.4 %. A paired *t*-test confirmed that CFFA's improvements in unbalanced settings were statistically significant. Future work will focus on integrating reinforcement learning to dynamically control ventilation based on AQI forecasts and scaling the system to multi-city deployments with automated edge model updates. Overall, CFFA offers a robust and adaptable foundation for context-sensitive, privacy-preserving federated learning in industrial IoT applications.

CRediT authorship contribution statement

Montaser N.A. Ramadan: Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Mohammed A.H. Ali:** Writing – review & editing, Supervision. **Mohammad Alkhedher:** Writing – original draft.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowlegement

Authors would like to thank Universiti Malaya and Ministry of High Education-Malaysia for supporting this work via Fundemental Research Grant Scheme (FRGS/1/2023/TK10/UM/02/3).

Appendix A. — Preliminary Fuzzy-Threshold Sensitivity Analysis

This appendix presents a compact sensitivity analysis evaluating how variations in the triangular membership thresholds of the

Stage-1 fuzzy system affect the final Adjusted Client Weights W_i' and the Final Global Aggregation Weight W_G'. All notation and membership definitions strictly follow Section 4.

A.1 Sensitivity Setup

The Stage-1 fuzzy system uses three fuzzified inputs:

- **Number of Workers: W_i^{worker}**

  Memberships: $\mu_{W,\text{low}}, \mu_{W,\text{medium}}, \mu_{W,\text{high}}$

- **Air Quality Index: W_i^{AQI}**

  Memberships: $\mu_{AQI,\text{good}}, \mu_{AQI,\text{moderate}}, \mu_{AQI,\text{bad}}$

- **Ventilation: $W_i^{\text{Ventilation}}$**

  Memberships: $\mu_{V,\text{poor}}, \mu_{V,\text{normal}}, \mu_{V,\text{good}}$
  All membership functions use a triangular form:

$$\mu(x) = \text{trimf}(x;\ a, b, c)$$

where (a,b,c) are the thresholds defining the triangle.
To evaluate robustness, all triangular breakpoints in Section 4 were perturbed by:

$$p \in \{-30\%, -20\%, -10\%, +10\%, +20\%, +30\%\}$$

For each membership function:

$$(a', b', c') = (a(1+p), b(1+p), c(1+p))$$

This generates perturbed versions of all nine Stage-1 membership functions:

- $\mu_{W,low}, \mu_{W,medium}, \mu_{W,high}$
- $\mu_{AQI,\text{good}}, \mu_{AQI,\text{moderate}}, \mu_{AQI,\text{bad}}$
- $\mu_{V,\text{poor}}, \mu_{V,\text{normal}}, \mu_{V,\text{good}}$

Each configuration produces new values of:

$$W_i^{\text{worker}}, W_i^{\text{AQI}}, W_i^{\text{Ventilation}}$$

These feed into the same Stage-1 rule base (Table 6), yielding updated adjusted client weights:

$$W_1', W_2', W_3'$$

Stage-2 then aggregates these using its own rule base (Table 7) to compute the final global weight:

$$W_G'$$

All perturbed models were evaluated on the same validation datasets used in Table 12.

A.2 Quantitative Sensitivity Results

Average results across the three factories are shown below:

Threshold Shift	Δ MAE (%)	Δ RMSE (%)	Δ Accuracy (points)
±10 %	1.1–2.0	1.8–2.9	−0.7 to −1.1
±20 %	1.9–3.4	2.5–3.9	−1.1 to −1.5
±30 %	2.8–3.9	3.4–4.6	−1.4 to −1.9

Interpretation of results:

- RMSE changes remain below 5 % even at ±30 % threshold variation, indicating stability.

- MAE changes remain below 4 %.
- Global accuracy drops by only 0.7–1.9 points, showing mild degradation.
- The ordering of the five operational cases (Table 12) remains unchanged.
- No membership function perturbation caused rule inversions or unstable output

W_i' or W_G'.

A.3 Why CFFA Remains Stable

CFFA exhibits stability for three key reasons:

1. Overlapping triangular membership functions prevent abrupt output changes.
2. Rule bases (Tables 6 and 7) combine contextual signals redundantly, reducing sensitivity to any single threshold.
3. The Stage-2 output W_G' smooths local deviations, acting as a global stabilizer.

Thus, CFFA does not depend on finely tuned membership thresholds—important in industrial environments where worker counts, AQI, and ventilation measurements naturally fluctuate.

Data availability

The authors do not have permission to share data.

References

[1] R.C. Karpagalakshmi, J. Lenin, P. Rajaram, V. Balaji, R. Rangasamy, A. Sungheetha, Adaptive anomaly detection for secure federated learning in IIoT environments, Procedia Comput. Sci. 258 (2025) 536–551.
[2] B.S. Diba, et al., Open problems and challenges in federated learning for IoT: a comprehensive review and strategic guide, Comput. Electr. Eng. 126 (2025) 110515.
[3] Z.-S. Tan, E.W.K. See-To, K.-Y. Lee, H.-N. Dai, M.-L. Wong, Privacy-preserving federated learning for proactive maintenance of IoT-empowered multi-location smart city facilities, J. Netw. Comput. Appl. 231 (2024) 103996.
[4] M.N.A. Ramadan, M.A.H. Ali, H. Jaber, M. Alkhedher, Blockchain-secured IoT-federated learning for industrial air pollution monitoring: a mechanistic approach to exposure prediction and environmental safety, Ecotoxicol. Environ. Saf. 300 (2025) 118442.
[5] J. Hao, R. Naja, D. Zeghlache, Adaptive federated reinforcement learning for critical realtime communications in UAV assisted vehicular networks, Comput. Netw. 247 (2024) 110456.
[6] X. Li, K. Huang, W. Yang, S. Wang, Z. Zhang, On the convergence of FEDAVG on non-IID data. International, Conference on Learning Representations (ICLR (2020) , ICLR, 2020, pp. 1–26.
[7] S. Yu, K. Zhu, F. Liang, J. Wang, K. Kant, L. Yin, Robust multimodal federated learning for non-IID multimodal data with incompleteness, Future Gener. Comput. Syst. (2025) 107948.
[8] X. Han, B. Liu, C. Hu, D. Cheng, OWL: Worker-assisted server bandwidth optimization for efficient communication federated learning, J. Parallel Distr. Comput. 198 (2025) 105027.
[9] S. Bharati, M. Mondal, P. Podder, V.B. Prasath, Federated learning: applications, challenges and future directions, Int. J. Hybrid Intell. Syst. 18 (1–2) (2022) 19–35.
[10] T. An, L. Ma, W. Wang, Y. Yang, J. Wang, Y. Chen, Consideration of fedprox in privacy protection, Electronics 12 (20) (2023) 4364.
[11] J. Wang, Q. Liu, H. Liang, G. Joshi, H.V. Poor, Tackling the objective inconsistency problem in heterogeneous federated optimization, Adv. Neural Inf. Process. Syst. 33 (2020) 7611–7623.
[12] D. Kim, K. Oh, Y. Lee, H. Woo, Overview of fair federated learning for fairness and privacy preservation, Expert Syst. Appl. (2025) 128568.
[13] T. Li, M. Sanjabi, A. Beirami, V. Smith, Fair resource allocation in federated learning, International, Conference on Learning Representations (ICLR, ICLR, 2020, pp. 1–27 (2020).
[14] Z. Qu, R. Duan, L. Chen, J. Xu, Z. Lu, Y. Liu, Context-aware online client selection for hierarchical federated learning, IEEE Trans. Parallel Distr. Syst. 33 (12) (2022) 4353–4367.
[15] R.A. Alharbey, F. Jamil, Federated learning framework for real-time activity and context monitoring using edge devices, Sensors 25 (4) (2025) 1266.
[16] J. Xu, C. Zhang, L. Jin, C. Su, A trust-aware incentive mechanism for federated learning with heterogeneous clients in edge computing, Journal of Cybersecurity and Privacy 5 (3) (2025) 37.
[17] Z. Song, H. Sun, H.H. Yang, X. Wang, Y. Zhang, T.Q.S. Quek, Reputation-based federated learning for secure wireless networks, IEEE Internet Things J. 9 (2) (2021) 1212–1226.
[18] T. Skripcak, et al., Creating a data exchange strategy for radiotherapy research: towards federated databases and anonymised public datasets, Radiother. Oncol. 113 (3) (2014) 303–309.
[19] Y.I. Daradkeh, I. Tvoroshenko, Technologies for making reliable decisions on a variety of effective factors using fuzzy logic, Int. J. Adv. Comput. Sci. Appl. 11 (5) (2020).
[20] A. Celikyilmaz, I.B. Türksen, Modeling Uncertainty with Fuzzy Logic, Springer, 2009.
[21] M. Wang, J.N.K. Liu, Fuzzy logic-based real-time robot navigation in unknown environment with dead ends, Robot. Autonom. Syst. 56 (7) (2008) 625–643.
[22] A. Wilbik, P. Grefen, Towards a federated fuzzy learning system, in: 2021 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE, IEEE, 2021, pp. 1–6.
[23] N. Cha, et al., Fuzzy logic based client selection for federated learning in vehicular networks, IEEE Open Journal of the Computer Society 3 (2022) 39–50.
[24] R. Llasag Rosero, C. Silva, B. Ribeiro, B.F. Santos, Label synchronization for hybrid federated learning in manufacturing and predictive maintenance, J. Intell. Manuf. 35 (8) (2024) 4015–4034.
[25] X. Cheng, F. Shi, Y. Liu, X. Liu, L. Huang, Wind turbine blade icing detection: a federated learning approach, Energy 254 (2022) 124441.
[26] Z. Zhou, L. Chu, C. Liu, L. Wang, J. Pei, Y. Zhang, Towards fair federated learning, in: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2021, pp. 4100–4101.

[27] Y. Guo, K. Huang, J. Chen, Wcl: client selection in federated learning with a combination of model weight divergence and client training loss for internet traffic classification, Wireless Commun. Mobile Comput. 2021 (1) (2021) 3381998.
[28] A. Alwabli, Federated learning for privacy-preserving air quality forecasting using iot sensors, Eng. Technol. Appl. Sci. Res. 14 (4) (2024) 16069–16076.
[29] Z. Wu, M.M. Amiri, R. Raskar, B.K.H. Low, Incentive-aware federated learning with training-time model rewards, in: The Twelfth International Conference on Learning Representations, 2024.
[30] S. Yarham, M. Behjati, H.A.H. Alobaidy, A.P.P.A. Majeed, Y. Zheng, Enhancing air quality monitoring: a brief review of federated learning advances, in: International Conference on Intelligent Manufacturing and Robotics, Springer, 2024, pp. 489–501.
[31] T. Van Quang, D.T. Doan, J. Ngarambe, A. Ghaffarianhoseini, A. Ghaffarianhoseini, T. Zhang, AI management platform for privacy-preserving indoor air quality control: review and future directions, J. Build. Eng. 100 (2025) 111712.
[32] S. Abimannan, A. John, S. Shukla, D. Satheesh, Federated learning for improved air pollution prediction: a combined LSTM-SVR approach, in: 2023 IEEE 4th Annual Flagship India Council International Subsections Conference (INDISCON), IEEE, 2023, pp. 1–7.
[33] Y. Feng, Application of a federated learning model with incentive mechanism in air prediction, in: Proceedings of the 2024 4th International Conference on Big Data, Artificial Intelligence and Risk Management, 2024, pp. 746–750.
[34] S. Beriwal, J. Ayeelyan, Decoding pollution: a federated learning-based pollution prediction study with health ramifications using causal inferences, Electronics 14 (2) (2025) 350.
[35] J. Haxhibeqiri, E. De Poorter, I. Moerman, J. Hoebeke, A survey of LoRaWAN for IoT: from technology to application, Sensors 18 (11) (2018) 3995.
[36] A. Maier, A. Sharp, Y. Vagapov, Comparative analysis and practical implementation of the ESP32 microcontroller module for the internet of things, in: 2017 Internet Technologies and Applications (ITA), IEEE, 2017, pp. 143–148.
[37] A. Aimi, S. Rovedakis, F. Guillemin, S. Secci, ELoRa: End-to-end emulation of massive IoT LoRaWAN infrastructures, in: NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium, IEEE, 2023, pp. 1–3.
[38] S.A. Horn, P.K. Dasgupta, The Air Quality Index (AQI) in historical and analytical perspective a tutorial review, Talanta 267 (2024) 125260.
[39] U.S. Environmental Protection Agency (EPA), Technical Assistance Document for the Reporting of Daily Air Quality – the Air Quality Index (AQI), EPA-454/B-18-007, Office of Air Quality Planning and Standards, Research Triangle Park, NC, September 2018 [Online]. Available: https://www.airnow.gov/sites/default/files/2020-05/aqi-technical-assistance-document-sept2018.pdf.
[40] M.N.A. Ramadan, M.A.H. Ali, S.Y. Khoo, M. Alkhedher, SecureIoT-FL: a federated learning framework for privacy-preserving real-time environmental monitoring in industrial IoT applications, Alex. Eng. J. 114 (2025) 681–701.