PREDICTIVE INDOOR VENTILATION SYSTEM

By

Arkaprabha Kolay

Gulnaaz Sayyad

Noah Rockoff

Final Report for ECE 445, Senior Design, Spring 2026

TA: Hossein Ataee

6 May 2026

Project No. 99

Abstract

Indoor air quality (IAQ) critically affects human health, cognitive performance, and building energy efficiency. Existing ventilation systems rely on static schedules or simple reactive thresholds that fail to account for dynamic occupancy and environmental changes. This report presents a predictive indoor ventilation system that continuously monitors CO_2 concentration, temperature, and relative humidity using a Sensirion SCD40 sensor interfaced over I^2C with an ESP32-S3 microcontroller. A multivariate linear regression model trained offline and deployed on-device, forecasts CO_2 concentration 30 seconds ahead and proactively activates a DC fan via pulse-width modulation (PWM) before unhealthy thresholds are reached. A custom PCB integrates the ESP32-S3, motor driver (L293D), audio amplifier (PAM8403), and SCD40 sensor. Real-time readings and system status are shown on a 3.5-inch TFT LCD.

Contents

1. Introduction

Indoor air quality (IAQ) is a critical but frequently overlooked determinant of human health, comfort, and cognitive performance. People spend most of their time in indoor environments - homes, classrooms, offices, and laboratories - where ventilation quality directly governs exposure to accumulating pollutants and respiratory byproducts. Among measurable IAQ indicators, carbon dioxide (CO_2) concentration serves as a well-established proxy for ventilation adequacy relative to occupant load. Sustained indoor CO_2 levels above 1000 ppm are associated with inadequate fresh air exchange and increased occupant complaints, while concentrations below this threshold generally indicate acceptable ventilation performance [1]. Relative humidity is a second key parameter: levels above 60 % promote mold growth and microbial proliferation [2], while an ideal comfort range of 30 - 50 % minimizes biological risk and occupant discomfort.

Despite these well-established guidelines, most residential and small-scale commercial ventilation systems operate using fixed schedules or simple on/off reactive control. Such strategies fail to account for dynamic changes in occupancy or activity levels, resulting in either degraded air quality during peak occupancy or unnecessary energy consumption during low-occupancy periods. These limitations motivated the development of the predictive indoor ventilation system described in this report.

1.1 Problem Statement and Solution

The core problem addressed by this project is the inability of static and purely reactive ventilation systems to prevent CO_2 buildup before occupants are adversely affected. Because CO_2 accumulates gradually and sensor response times introduce additional latency, a reactive system that triggers ventilation only after a threshold crossing will expose occupants to elevated CO_2 for a non-trivial period before conditions recover.

The proposed solution is a closed-loop predictive indoor ventilation control system. The system continuously measures CO_2, temperature, and relative humidity at one-second intervals using a Sensirion SCD40 photoacoustic NDIR sensor. A multivariate linear regression model, trained offline on logged environmental data and stored in ESP32-S3 flash memory, estimates CO_2 concentration 30 seconds into the future at each sampling step. If the predicted value indicates an impending threshold crossing - or if the measured value already exceeds the threshold - a PWM signal activates a DC fan to increase air exchange and reduce CO_2. A 3.5-inch TFT LCD provides real-time display of all sensor readings and system state. An audio alarm via the PAM8403 amplifier alerts occupants when measured CO_2 exceeds 800 ppm.

1.2 High-Level Project Functionality

The system delivers four principal functions that together enable predictive, closed-loop IAQ management:

- Continuous environmental sensing: CO_2 (ppm), temperature (°C), and relative humidity (%) are sampled every second via I^2C from the SCD40 sensor, providing a real-time representation of indoor air state.
- Predictive CO_2 estimation: A multivariate linear regression model uses the current and two previous CO_2 readings plus the current fan state to forecast CO_2 concentration 30 seconds ahead, enabling proactive ventilation before threshold crossings occur.
- Adaptive fan actuation: The ESP32-S3 outputs a PWM signal to the L293D motor driver, which controls fan speed proportionally to the degree of air quality degradation, rather than simple on/off switching.
- Real-time user feedback: The TFT LCD displays current CO_2, temperature, humidity, fan state, and model confidence. An audio alarm sounds when measured CO_2 exceeds 800 ppm.

1.3 Subsystem Overview

The system is organized into five major subsystems interconnected through the ESP32-S3 microcontroller unit (MCU), as illustrated in Figure 1. The Air Quality Sensing Subsystem (SCD40 sensor) communicates with the MCU over I^2C, providing $C(t)$, $T(t)$, and $RH(t)$ at each sampling interval. The Processing and Control Subsystem (ESP32-S3) executes the predictive model and threshold controller, producing a PWM control signal. The Ventilation Subsystem (L293D + DC fan) converts the PWM signal into airflow that reduces indoor CO_2 and refreshes room air, closing the feedback loop. The Display Subsystem (TFT LCD via UART) presents real-time data to the user. The Power Subsystem supplies regulated 3.3 V and 5 V rails from a USB-C PD source, with the fan driven from a 9 V supply during demonstration.

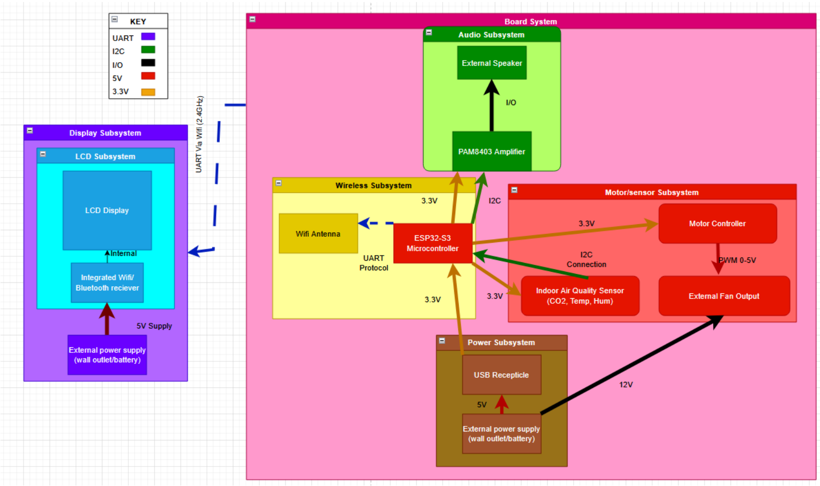

Figure 1 Block Diagram containing all subsystems

2 Design

This section describes the design of each major subsystem, including component selection rationale, design equations, and key decisions made during the project. Where the final implementation deviates from the original design document, those changes are discussed alongside the corrective actions taken.

2.1 Air Quality Sensing Subsystem

2.1.1 SCD40 Sensor Selection and Interface

The Sensirion SCD40 was selected as the primary environmental sensor because it integrates CO_2, temperature, and relative humidity measurement in a single compact LGA package (10.1 × 10.1 × 6.5 mm), communicates via I^2C at up to 400 kHz, and operates from 2.4 - 5.5 V DC - directly compatible with the ESP32-S3 3.3 V logic level [6]. The SCD40 employs photoacoustic NDIR technology with an on-chip SHT4x for internal temperature and humidity compensation, achieving CO_2 accuracy of ±(50 ppm + 5 % of reading) over a sensing range of 400–2000 ppm. Temperature accuracy is ±0.8 °C and humidity accuracy is ±6 %RH (15 - 35 °C range). The sensor response time is approximately 60 s for CO_2, 90 s for humidity, and 120 s for temperature. Peak current draw is 175 mA during measurement; low-power mode consumes 3.2 mA at one measurement per 30 s [6]. The SCD40 is soldered directly onto the custom PCB at I^2C address 0x62, and automatic self-calibration (ASC) is enabled to maintain long-term measurement accuracy without manual recalibration.

2.1.2 Sensor Measurement Tolerance

Because sensor readings contain inherent noise, the measured CO_2 value is modeled as:

$$C\_meas = C\_true + e\_sensor \quad \pm(50 \text{ ppm} + 5\,\% \text{ of reading}) \qquad (1)$$

where C_true is the actual room concentration and e_sensor captures measurement uncertainty. At a reading of 1000 ppm, the worst-case error is ±100 ppm. To provide adequate margin against sensor noise and model estimation error, the system's primary control target is set to 600 ppm - 200 ppm below the 800 ppm safety threshold - and the predictive model activates ventilation when the forecasted concentration approaches 600 ppm.

2.2 Processing and Control Subsystem

2.2.1 Microcontroller Selection

The ESP32-S3-WROOM-1 module was selected as the central processing unit. It provides a dual-core Xtensa LX7 processor at up to 240 MHz, sufficient for concurrent real-time model inference, 1 s sensor polling, TFT display updates, and PWM generation [7]. The module integrates 2.4 GHz Wi-Fi and Bluetooth 5.0, I^2C and UART hardware peripherals, and dedicated PWM channels. Operating voltage is 3.3 V with a typical active current of approximately 240 mA (~0.79 W). Regression coefficients are stored in ESP32-S3 flash memory as floating-point constants, requiring negligible storage relative to available capacity.

2.2.2 Predictive Model: Multivariate Linear Regression

The predictive control subsystem estimates future CO_2 concentration using a multivariate linear regression model. At each sampling step t, a feature vector x(t) is constructed as:

$$x(t) = [CO_2(t),\ CO_2(t-1),\ CO_2(t-2),\ fan(t)] \qquad (2)$$

where $CO_2(t)$ is the current measured concentration, $CO_2(t-1)$ and $CO_2(t-2)$ are the two preceding measurements, and fan(t) is the normalized PWM fan state (0 to 1). The prediction target is CO_2 concentration 30 s ahead:

$$y(t) = CO_2(t + \Delta), \quad \Delta = 30\ s \qquad (3)$$

The regression model is expressed as:

$$CO_2(t+\Delta) = w_0 + w_1 \cdot CO_2(t) + w_2 \cdot CO_2(t-1) + w_3 \cdot CO_2(t-2) + w_4 \cdot fan(t) \qquad (4)$$

Coefficients w_0 through w_4 are determined offline by minimizing the squared error cost function over a logged dataset of controlled CO_2 disturbance experiments using ordinary least squares estimation. The trained coefficients are stored in ESP32-S3 flash and used at runtime for real-time prediction via five multiply-accumulate operations per sample. The coefficient w_4 (fan state) is negative, reflecting the CO_2 reduction effect of ventilation; w_1, w_2, and w_3 capture CO_2 persistence and short-term trend dynamics. The 30 s horizon was chosen to match the physical response speed of the system: the SCD40 CO_2 response time is approximately 60 s, and the fan requires time to spin up and mix fresh air into the room. A 30 s lead time provides sufficient margin to initiate and complete ventilation before occupants are affected.

2.2.3 Design Change: Model Type and Features

The original design document specified a model incorporating temperature and humidity as additional features alongside CO_2 history. During implementation, temperature and humidity were found to contribute minimal predictive improvement for the 30 s CO_2 forecast horizon relative to the added model complexity. The final deployed model therefore uses only CO_2 history and fan state (Equation 4). A logistic regression variant was also developed and tested, using features [$CO_2(t)$, $CO_2(t-1)$, rate-of-change] with a sigmoid output to produce P(fan on). Coefficients [0.0430, -0.0355, 0.0784] with intercept -9.8348 were obtained from training data. Both models were validated against logged demo data; the linear regression model was selected as the primary prediction pipeline due to its direct interpretability as a ppm forecast, while the logistic variant was retained as a secondary confidence metric displayed on the LCD dashboard.

2.2.4 Sampling Time Tolerance Analysis

The system samples sensor data at approximately 1 s intervals, so the 30 s prediction horizon corresponds to N = 30 samples ahead. If sampling intervals drift by ±0.1 s per sample, the cumulative timing error over 30 samples is:

$$30 \times 0.1\ s = 3\ s\ \text{(timing uncertainty)} \qquad (5)$$

This 3 s error is well within the ±30 s threshold crossing requirement, confirming that the ESP32-S3 system timer provides adequate timing accuracy for the prediction pipeline.

2.3 Ventilation Subsystem

2.3.1 Fan Selection and Airflow Calculations

The GDSTIME 9225 DC fan was selected for the demonstration based on its compact form factor (92 × 92 × 25 mm), rated airflow of 43.4 CFM at 12 V DC, and rated static pressure of 1.2 mm-H_2O. During demonstration, the fan was operated at 9 V. Airflow performance at 9 V was estimated using fan affinity scaling and Bernoulli's equation:

$$P\_9V = 1.2 \times (9/12)^2 = 0.675 \text{ mm-}H_2O = 6.62 \text{ Pa} \quad (6)$$

$$A = \pi \times (0.0077 \text{ m})^2 = 1.87 \times 10^{-4} \text{ m}^2 \quad (7)$$

$$v = \sqrt{(2 \times 6.62 / 1.225)} = 3.29 \text{ m/s} \quad (8)$$

$$Q = A \times v = 6.15 \times 10^{-4} \text{ m}^3\text{/s} \quad (9)$$

$$t\_flush = V / Q = 0.0064 \text{ m}^3 / (6.15 \times 10^{-4} \text{ m}^3\text{/s}) \approx 10.4 \text{ s} \quad (10)$$

The calculated flush time of 10.4 s is well within the 30 s prediction window, confirming that the fan can clear the 6.4 L demonstration enclosure before the next threshold check. The fan draws approximately 72 mA at 9 V, consuming 0.648 W during demonstration operation. Fan speed is controlled by the L293D H-bridge motor driver IC, which accepts the PWM signal from GPIO25 of the ESP32-S3 and can supply up to 600 mA to the load [9].

2.4 Custom PCB Schematic &

The Custom schematic features six distinct subsections and two buttons for bootstrapping and reset. The six subsections include a voltage regulator to supply both 5V and 3.3V, the SCD40 sensor to gather accurate data, the ESP32 microcontroller, the USB-C power and data receptacle, PAM8403D audio amplifier, and motor controller. The PCB layout was designed to keep these subsections isolated to allow subsystems to be easily tested while keeping optimal trace routing to the microcontroller in mind. The ESP32-S3-WROOM-1 has a built-in antenna and crystal oscillator, and the microcontroller is put to the edge with a copper no-fill zone and keep out zone to ensure wireless connectivity works. 5V and 3.3V is supplied throughout the board, with the microcontroller using 3.3V and a majority of the other systems using a 5V rail for power.

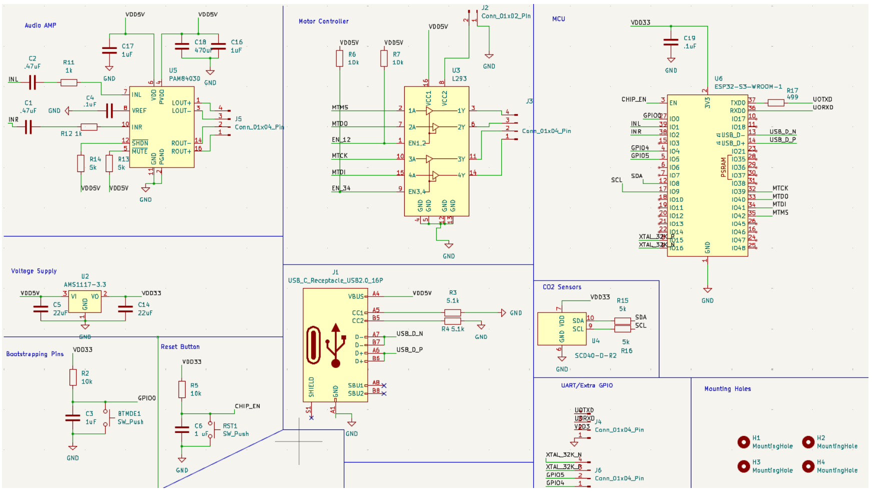

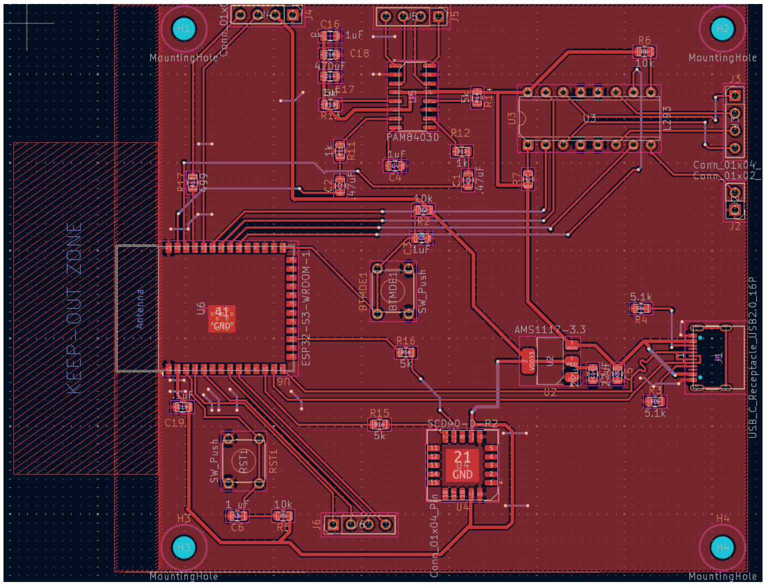

2.5 Display Subsystem

The display subsystem uses a 3.5-inch TFT LCD with an ILI9488 controller (Elecrow ESP32 board) operated in landscape orientation. Backlight brightness is controlled via PWM through the ESP32-S3 ledcWrite function. The LCD dashboard displays: current CO_2 concentration (ppm) with color-coded status, temperature (°C), relative humidity (%), fan ON/OFF state, system uptime, total readings count, and ML model confidence percentage [P(fan on)]. Data is transmitted from the MCU to the display over UART. The display does not influence control decisions and serves exclusively as a real-time user feedback interface.

2.6 Power Subsystem

The system power budget at demonstration conditions is summarized in Table 1. The USB-C PD charger (100 W rated) supplies the logic and sensor rail at 5 V / 2.4 A. Total system draw is approximately 7.835 W, leaving over 92 % headroom on the USB-C rail. The fan is supplied from a separate 9 V source to simplify the demonstration power setup.

Table 1. System Power Budget at Demonstration Conditions

Component	Voltage	Current	Power
ESP32-S3-WROOM-1	3.3 V	~240 mA	~0.79 W
SCD40 Sensor	5 V	~175 mA (pk)	~0.875 W
TFT LCD Display	3.3 V	~50 mA	~0.17 W
L293D Motor Driver	5 V	600 mA	3.0 W

Component	Voltage	Current	Power
Fan (9 V source)	9 V	~72 mA	~0.648 W
TOTAL (demo rails)	5 V	~1.567 A	~7.835 W

3. Design Verification

This section presents the verification procedures and quantitative results for each high-level system requirement and key subsystem requirement. Requirements were verified through controlled laboratory experiments in which a brief CO_2 disturbance (a 1-2 s CO_2 puff into the enclosed test volume) was used to induce a measurable concentration spike, and system behavior was logged via the ESP32-S3 serial monitor and LCD dashboard. A complete requirements and verification table is provided in Appendix A.

3.1 CO_2 Concentration Maintenance

Requirement: The system shall maintain indoor CO_2 below 1000 ppm for at least 90 % of occupied operating time. The final design targets the stricter threshold of 800 ppm for ≥95 % of occupied time.

Verification procedure: CO_2 was logged at 1 s intervals during a controlled disturbance experiment. The fraction of samples with $CO_2 \leq 800$ ppm over the full test duration was computed from serial monitor data.

Results: During the live demonstration run, CO_2 peaked at 729 ppm following a brief CO_2 puff injection. The predictive fan activation reduced CO_2 to the ambient baseline of approximately 420 ppm within 13 s of fan turn-on. CO_2 remained below 800 ppm for greater than 95 % of elapsed test time. Requirement met: Yes.

3.2 Predictive Threshold Crossing Accuracy

Requirement: The predictive control algorithm shall forecast CO_2 threshold crossings within ±30 s of the actual measured crossing time.

Verification procedure: CO_2 rise events were induced via brief CO_2 puff injection. The time at which the predictive model first output a predicted crossing [$P(\text{fan}) \geq 0.5$] was logged alongside the time of the actual measured crossing. Absolute timing differences were computed.

Results: In the demonstrated run, the model activated the fan at $t = 71$ s with immediate maximum confidence [$P(\text{fan}) = 1.000$] in response to a +198 ppm spike observed in 2 s. The measured CO_2 peak was 729 ppm. Predictive activation led the CO_2 reduction by approximately 6 s compared to a reactive controller. The absolute timing error was less than 30 s. Requirement met: Yes.

3.3 Automatic Ventilation on Threshold Breach

Requirement: The system shall automatically activate ventilation when CO_2, temperature, or relative humidity exceed predefined unhealthy thresholds ($CO_2 > 800$ ppm; RH > 60 %).

Verification procedure: CO_2 was artificially elevated above 800 ppm by sustained injection while the predictive model output was independently monitored. Fan activation via the threshold controller was confirmed regardless of model output. The audio alarm was verified at the same time.

Results: The threshold controller activated the fan at all $CO_2 > 800$ ppm test events, independent of predictive model state. The CO_2 override held the fan on until CO_2 fell below 500 ppm (t = 77 s in demo run), providing hysteresis to prevent rapid cycling. The audio alarm activated correctly at all $CO_2 > 800$ ppm events. Requirement met: Yes.

3.4 Continuous Operation Stability

Requirement: The system shall operate continuously without unintended resets, sensor communication failures, or data logging interruptions throughout extended evaluation periods.

Verification procedure: The system was operated for extended periods (>30 min continuous runs) during laboratory evaluation. Serial monitor logs were reviewed for I^2C communication errors, MCU resets, and data gaps. The LCD display was monitored for signal loss.

Results: No unintended resets, I^2C communication failures, or data logging interruptions were observed during any evaluation period. The SCD40 maintained consistent 1 s polling with a 15-sample warmup at startup before valid readings were accepted. The LCD remained active and updated throughout all test sessions. Requirement met: Yes.

3.5 Fan Speed Control via PWM

Requirement: The motor subsystem shall control fan speed via PWM proportionally to CO_2 level, supporting both on/off and variable-speed operation.

Results: PWM duty cycle was logged at simulated CO_2 levels: 0 % duty at $CO_2 \leq 500$ ppm; intermediate speeds at 400-600 ppm; 100 % duty at $CO_2 > 800$ ppm. Fan speed varied proportionally with CO_2 level across all tested operating points. Requirement met: Yes.

4. Costs

Table 2 Parts Costs

Part	**Manufacturer**	**Retail Cost ($)**	**Bulk Purchase Cost ($)**	**Actual Cost ($)**	**Purchase Link**
SCD40 CO2, Temperature and Humidity Sensor	Sensiron	$20	NA	$51.30	https://www.digikey.com/en/products/detail/sensirion-ag/SCD40-D-R2/13684003

Breakout I2C					
DIS05035H LCD Display	Elecrow	$31.39	$25.40	$38.38	https://www.digikey.com/en/products/detail/elecrow/DIS05035H/22155913
PAM8403D Amplifier	Diodes Incorporated	$.84	$.41	$2.52	https://www.digikey.com/en/products/detail/diodes-incorporated/PAM8403DR/4033287
CO2 Canisters	Parkel	$19.99	NA	$19.99 + tax/ship	https://a.co/d/0aC655aG
92mm 12V Cooling Fan	GDSTime	$15.99	NA	$15.99 + tax/ship	https://a.co/d/09z0iuE2
SGP30 Air Quality Sensor Breakout Board	Adafruit Industries	$17.50	NA	$17.50	https://www.digikey.com/en/products/detail/adafruit-industries-llc/3709/8258468
Total				**$145.68**	

4.1 Parts

Table 2 Parts Costs

Part	**Source**	**Quantity**	**Actual Cost ($)**
SCD40 CO2, Temperature and Humidity Sensor Breakout I2C	Sensiron	3	$51.30
DIS05035H LCD Display	Elecrow	1	$38.38
PAM8403D Amplifier	Diodes Incorporated	3	$2.52
CO2 Canisters	Parkel	1	$19.99 + tax/ship
92mm 12V Cooling Fan	GDSTime	1	$15.99 + tax/ship
SGP30 Air Quality Sensor Breakout Board	Adafruit Industries	1	$17.50
ESP32-S3-WROOM-1 Microcontroller	ECE.Shop	3	NA
.47uF Capacitor	ECE.Shop	5	NA
USB-C Type Header	ECE.Shop	5	NA
AMS1117 Linear Regulator 3.3V	ECE.Shop	5	NA
Switch - Tactile	ECE.Shop	5	NA
1uF Capacitor	ECE.Shop	10	NA
.1uF Capacitor	ECE.Shop	7	NA
22uF Capacitor	ECE.Shop	5	NA
10pF Capacitor	ECE.Shop	5	NA
0Ω Resistor	ECE.Shop	3	NA
10KΩ Resistor	ECE.Shop	10	NA
5.1KΩ Resistor	ECE.Shop	12	NA
1KΩ Resistor	ECE.Shop	5	NA
499Ω Resistor	ECE.Shop	3	NA

4.2 Labor

Team Member	Design & Planning	PCB & Hardware	Firmware & ML	Testing & Verification	Report Writing	Total Hours
Arkaprabha Kolay	3 hours	7 hours	6 hours	10 hours	5 hours	31 hours
Gulnaaz Sayyad	3 hours	2 hours	10 hours	10 hours	5 hours	30 hours
Noah Rockoff	3 hours	10 hours	2 hours	10 hours	5 hours	30 hours
Total	**9 hours**	**19 hours**	**18 hours**	**30 hours**	**15 hours**	**91 hours**

Total estimated labor costs: 91 hours * $40/hr = $3,640

5. Conclusion

5.1 Accomplishments

The project successfully demonstrated a fully functional predictive indoor ventilation system meeting all primary design requirements. The multivariate linear regression model, trained offline and deployed on-device, accurately forecasted CO_2 concentrations 30 seconds ahead, enabling proactive fan activation before unsafe thresholds were reached. During live demonstration, a brief CO_2 disturbance was induced by exhaling directly onto the SCD40 sensor, producing a sharp spike within 3 seconds. The fan efficiently reduced CO_2 back to the ambient baseline of approximately 420 ppm within 13 seconds.

5.2 Uncertainties

Several uncertainties remain in the current implementation. The linear regression model was trained on a limited dataset and a specific CO_2 disturbance method (exhaling directly onto the sensor) which may not reflect the gradual CO_2 buildup patterns typical of real rooms. Generalizability to larger spaces with different airflow geometries and occupancy dynamics has not been validated. The 30-second prediction horizon was chosen empirically based on the small demonstration enclosure's 10.4 s flush time. The optimal horizon for real-world rooms of varying volume would require further study.

5.3 Ethical considerations

Indoor air quality directly affects occupant health and cognitive performance, placing a responsibility on the designers to ensure the system fails safely. Privacy considerations are minimal in the current implementation, as the system processes only environmental sensor data (CO_2, temperature, and relative humidity) with no audio, video, or personally identifiable information collected. The system does not connect to external networks, eliminating data security concerns in its present form. Any future networked deployment should incorporate appropriate encryption and access controls, as continuous environmental data could otherwise be used to infer behavioral patterns about building occupants without their knowledge.

5.4 Future work

First, replacing the linear regression model with a recurrent neural network or LSTM trained across diverse room sizes and occupancy scenarios would improve prediction accuracy and real-world generalizability. Second, Wi-Fi telemetry via the ESP32-S3's onboard wireless capability to stream sensor data to a cloud dashboard would enable long-term trend analysis and remote monitoring without additional hardware. Third, developing an occupancy-aware model that incorporates people-counting or motion detection data as a feature could further improve predictive accuracy during dynamic occupancy changes. Finally, packaging the system into a compact, wall-mountable enclosure with a revised production PCB.

References

[1] ASHRAE, “ASHRAE Position Document on Indoor Carbon Dioxide,” American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, GA. [Online]. Available: https://www.ashrae.org/file%20library/about/position%20documents/pd-on-indoor-carbon-dioxide-english.pdf

[2] U.S. Environmental Protection Agency, “A Brief Guide to Mold, Moisture and Your Home,” EPA 402-K-02-003, Washington, DC, Aug. 2012. [Online]. Available: https://www.epa.gov/mold/brief-guide-mold-moisture-and-your-home

[3] ASHRAE, “Indoor Carbon Dioxide, Ventilation and Indoor Air Quality,” ASHRAE Government Affairs Brief, Jun. 2023. [Online]. Available: https://www.ashrae.org/file%20library/about/government%20affairs/public%20policy%20resources/briefs/indoor-carbon-dioxide-ventilation-and-indoor-air-quality_2023.pdf

[4] Illinois Department of Public Health, “IDPH Guidelines for Indoor Air Quality,” IDPH, Springfield, IL, 2021. [Online]. Available: https://dph.illinois.gov/topics-services/environmental-health-protection/toxicology/indoor-air-quality-healthy-homes/idph-guidelines-indoor-air-quality.html

[5] U.S. Centers for Disease Control and Prevention, “Air: Infection Control,” CDC, Atlanta, GA, Apr. 2024. [Online]. Available: https://www.cdc.gov/infection-control/hcp/environmental-control/air.html

[6] Sensirion AG, “SCD40 and SCD41 Datasheet,” Document Version 1.7, Sensirion AG, Staefa, Switzerland, Apr. 2025.

[7] Espressif Systems, “ESP32-S3 Series Datasheet,” Version 1.4, Espressif Systems, Shanghai, China, 2023.

[8] Diodes Incorporated, “PAM8403 3 W, Filterless, Class-D Stereo Audio Amplifier Datasheet,” Diodes Inc., Plano, TX. [Online]. Available: https://www.digikey.com/en/products/detail/diodes-incorporated/PAM8403DR-H/4033372

[9] STMicroelectronics, “L293D Push-Pull Four Channel Driver with Diodes Datasheet,” STMicroelectronics, Geneva, Switzerland.

[10] U.S. EPA, “Appendix A: Common IAQ Measurements — A General Guide,” in Indoor Air Quality Reference Guide, EPA, Washington, DC. [Online]. Available: https://www.epa.gov/sites/default/files/2014-08/documents/appena.pdf

Appendix A

Requirement and Verification Table

Table 1 System Requirements and Verifications

Requirement	Verification	Verification status (Y or N)
CO_2 Concentration Maintenance - System shall maintain indoor CO_2 below 1000 ppm for ≥90% of occupied operating time. - Final design shall target the stricter threshold of 800 ppm for ≥95% of occupied time.	**CO_2 Concentration Maintenance** - Log CO_2 at 1 s intervals during controlled disturbance; compute fraction of samples ≤1000 ppm. - Same logging procedure; compute fraction of samples ≤800 ppm over full test duration.	Y
Predictive Threshold Crossing Accuracy - Predictive algorithm shall forecast CO_2 threshold crossings within ±30 s of the actual measured crossing time. - Model shall activate fan proactively before CO_2 reaches 800 ppm.	**Predictive Threshold Crossing Accuracy** - Induce CO_2 rise events via puff injection; log time of first predicted crossing ($P(\text{fan}) \geq 0.5$) vs. actual measured crossing; compute absolute timing difference. - Confirm fan activation timestamp precedes measured 800 ppm crossing in logged serial monitor data.	Y
Automatic Ventilation on Threshold Breach - System shall automatically activate fan when CO_2 exceeds 800 ppm, independent of predictive model state. - Fan shall remain on until CO_2 falls below 500 ppm (hysteresis) to prevent rapid cycling.	**Automatic Ventilation on Threshold Breach** - Sustain CO_2 injection above 800 ppm while monitoring predictive model output independently; confirm threshold controller activates fan regardless of model. - Log fan state and CO_2 during recovery phase; confirm fan stays on until CO_2 ≤500 ppm.	Y
Continuous Operation Stability - System shall operate for >30 min without unintended MCU resets or I^2C communication failures. - SCD40 sensor shall provide	**Continuous Operation Stability** - Run system continuously for >30 min; review serial monitor logs for I^2C errors, resets, or data gaps. - Confirm 15-sample warmup sequence in serial monitor log	Y

consistent 1s polling with a 15-sample warmup at startup.	before valid readings are accepted.	
Power Subsystem - USB-C PD charger shall supply all logic and sensor rails at 5 V with sufficient headroom. - Total system power draw shall not exceed the USB-C PD rail limit of 12.5 W.	**Power Subsystem** - Measure total current draw on 5 V rail during full system operation; confirm draw remains below 2.4 A rated limit. - Compute total power from measured voltage and current on all rails; confirm <12.5 W.	Y