Scientific Reports

https://doi.org/10.1038/s41598-026-40990-w

Article in Press

Presence-gated VOC sensing for urban search and rescue applications

Received: 12 November 2025

Accepted: 17 February 2026

Published online: 24 February 2026

Cite this article as: Tanggono E.N., Mokhtarzadeh A.A., Balasubramaniam K. *et al.* Presence-gated VOC sensing for urban search and rescue applications. *Sci Rep* (2026). https://doi.org/10.1038/s41598-026-40990-w

Ervan N. Tanggono, Amir A. Mokhtarzadeh, Kaveendran Balasubramaniam, Jing Pin Zou, Hera Naeem, Parvez Mosharof & Milon Selvam Dennison

Presence-Gated VOC Sensing for Urban Search and Rescue Applications

Ervan N. Tanggono[1], [1]Faculty of Mechanical and Materials Engineering, Huaiyin Institute of Technology, No.1 Mei Cheng Road, Huai'an, Jiangsu Province, China, Email: ervannovt@gmail.com

Amir A. Mokhtarzadeh[2], [2]Faculty of Computer Science and Software Engineering, Huaiyin Institute of Technology, No.1 Mei Cheng Road, Huai'an, Jiangsu Province, China, Email: amir@rooz.co.uk

Kaveendran Balasubramaniam*[1], [1]Faculty of Mechanical and Materials Engineering, Huaiyin Institute of Technology, No.1 Mei Cheng Road, Huai'an, Jiangsu Province, China, Email: kaveendran.b@hyit.edu.cn

Jing Pin Zou[2], [2]Faculty of Computer Science and Software Engineering, Huaiyin Institute of Technology, No.1 Mei Cheng Road, Huai'an, Jiangsu Province, China, Email: 1647297275@qq.com

Hera Naeem[2], [2]Faculty of Computer Science and Software Engineering, Huaiyin Institute of Technology, No.1 Mei Cheng Road, Huai'an, Jiangsu Province, China, Email: heranaeem@hotmail.com

Parvez Mosharof[2], [2]Faculty of Computer Science and Software Engineering, Huaiyin Institute of Technology, No.1 Mei Cheng Road, Huai'an, Jiangsu Province, China, Email: parvesk918@gmail.com

Milon Selvam Dennison*[3], [3]Department of Mechanical Engineering, School of Engineering and Applied Sciences, Kampala International University, Western Campus, Ishaka-Bushenyi, Uganda, Email: milon.selvam@kiu.ac.ug

*Corresponding Author: kaveendran.b@hyit.edu.cn; milon.selvam@kiu.ac.ug

Abstract

Urban search and rescue (USAR) operations require sensing modalities that remain effective under conditions where visibility-dependent techniques are degraded by dust, smoke, debris, or occlusion. This work presents a presence-gated volatile organic compound (VOC) sensing approach intended to support triage decisions in USAR scenarios. The system combines an actively sampled metal-oxide gas sensor array with environmental compensation and embedded inference, while a 24 GHz FMCW radar is used to verify human presence through

micro-motion cues and suppress chemistry-only false alarms. The sensing pipeline is designed for microcontroller-class execution, employing deterministic timing, duty-cycled operation, and compact feature extraction suitable for resource-constrained platforms. Offline evaluation was conducted using tethered data collection to assess instrument-level separability across controlled exposure conditions, including fresh air, surrogate injury-related VOC mixtures, and reference gases. Results are reported in terms of confusion matrices, ROC and precision–recall curves, and feature-importance analysis, emphasizing system behavior and feasibility rather than diagnostic performance. The findings demonstrate that presence-gated VOC sensing can provide a complementary information channel for USAR applications, while large-scale field validation and biological specificity remain topics for future investigation.

Keywords: electronic nose, metal-oxide gas sensors, ESP32, TinyML, FMCW radar presence gating, urban search-and-rescue, Random Forest, volatile organic compounds.

Introduction

In contemporary urban search-and-rescue operations, victim localization typically relies on a combination of optical cameras, thermal imaging, acoustic sensing, and radar-based techniques. While these modalities are highly effective under favorable conditions, their performance can degrade in post-disaster environments characterized by dust, smoke, debris occlusion, confined voids, and thermal equilibration. In such scenarios, visibility-dependent sensors may fail to provide reliable cues, particularly during early search phases. Volatile organic compound (VOC) based sensing offers a complementary information channel by exploiting chemical signatures associated with human presence that can diffuse through cracks and porous debris where line-of-sight sensing is ineffective. Rather than competing with established modalities, VOC sensing is best viewed as an auxiliary triage cue that can guide responders or robotic platforms toward regions of interest, where confirmatory sensing and rescue actions can then be prioritised.

Timely identification of live victims in urban search-and-rescue (USAR) requires sensing that remains effective under obstructed sight lines and constrained access. While line-of-sight modalities (RGB cameras, thermal imagers) are foundational, their performance degrades with occlusion and thermal equilibration; consequently, visibility-independent cues have been investigated. Unlike optical or thermal sensing, which can degrade under occlusion, smoke, or thermal equilibration, VOC-based sensing remains viable in visually inaccessible voids, making it a complementary modality rather than a replacement for existing USAR tools. Among these, airborne chemical signatures from human bioeffluents are promising: prior studies demonstrate that compact multi-sensor arrays can detect VOC mixtures indicative of live human presence and entrapment under chamber and field-like conditions [1,2].

VOC associated with human presence and entrapment have been systematically investigated as potential markers for urban search and rescue. A comprehensive analytical review by Güntner *et al.* highlighted that human breath, skin emissions, and wound-related volatilomes contain reproducible patterns of compounds such as ammonia, ketones, aldehydes, and reduced sulfur species that can persist and propagate through debris and confined spaces [1]. Importantly, this work emphasized that USAR-relevant detection relies on multivariate VOC fingerprints rather than single-compound specificity, motivating the use of sensor arrays coupled with pattern-recognition techniques rather than selective gas detectors.

Biological olfaction motivates this approach. Humans express on the order of ~350–400 functional olfactory receptor (OR) genes, enabling high-dimensional, combinatorial encoding of odors; receptor usage, expression, and perceptual impact vary across individuals [3,4]. Beyond receptor diversity, active sniffing and nasal conditioning (humidification/temperature) turn smell into a controlled sampling process rather than passive exposure [5]. Engineered electronic noses operationalize a pattern-recognition paradigm using arrays of semi-selective sensors coupled to multivariate statistical or machine-learning models, rather than pursuing

single-analyte selectivity; this approach traces to early artificial-olfaction work [6]. Evidence from other domains reinforces this generality: a biomimetic e-nose with nano-sensor arrays and machine-learning algorithms has been used for coffee detection [7].

Translating biology into robust devices demands careful engineering because metal-oxide semiconductor (MOX) sensors are cross-sensitive and drift with temperature, humidity, and flow. Recent reviews emphasize chamber geometry and active sampling (fan-assisted intake) to stabilize exposures, plus environmental compensation at the firmware layer to reduce confounding practices that materially improve repeatability and signal-to-noise in portable systems [8–10]. Biomimetic, directed-flow chambers homogenize sensor exposure and reduce response times in portable e-noses [11].

Within this context, the present work introduces a presence-gated VOC sensing system for urban search and rescue applications built on an electronic-nose core: a multi-sensor MOX array in an actively sampled micro-chamber, firmware-level temperature/humidity compensation, and embedded inference sized for microcontrollers (TinyML-style), keeping decisions local to minimize latency and network dependence [12,13]. The biochemical targets align with trauma and early-infection literature ammonia/biogenic amines, reduced sulfur species (e.g., H_2S), ketones, and mid-chain aldehydes which repeatedly appear across wound and whole-body volatilome studies [14,15]. In parallel, proof-of-concept studies demonstrate that low-cost MOX arrays can discriminate wound-related microbial VOC fingerprints, supporting a fingerprint-over-single-marker strategy [16].

To confirm presence and guide responders, the device couples to a compact FMCW radar front end, which consists range/angle bins and micro-motion (respiration/heartbeat) features pass through debris and thin occlusions, enabling presence-gated chemical alarms and coarse localization when sweeping voids or tunnels [17–19]. This integrated perspective, active-sampling MOX e-nose with embedded analytics, fused with radar presence cues, focuses on deployable architecture aimed at triage support rather than laboratory discovery.

Prior USAR-oriented demonstrations further motivate feasibility, as the selective, nanostructured MOX array has tracked breath/skin tracers like acetone, ammonia, and isoprene at parts-per-billion concentrations in simulated entrapment, indicating sensitivity compatible with fieldable tools [1]. More recently, a purpose-built platform quantified responses to decomposition/starvation surrogates and discussed practical aspects of portable deployment in post-disaster settings [2]. Building on this evidence base, the present work contributes a live human body detector that integrates an actively sampled MOX array, firmware-level temperature/humidity compensation, on-device TinyML inference on an ESP32, and 24 GHz FMCW presence gating, emphasizing device architecture and synchronized processing rather than new algorithms. This framing provides a reproducible foundation upon which subsequent trials can focus on operational validation in standardized scenarios and dataset expansion.

Related works

E-nose systems designed for Urban Search and Rescue and for detecting VOC fingerprints associated with injuries or infections.

Early USAR-oriented studies established that compact chemoresistive arrays can detect human-presence mixtures (breath/skin VOCs; e.g., acetone, ammonia, and isoprene) under chamber and field-like conditions [1]. More recent disaster-context experiments showed discrimination of starvation vs. decomposition signatures relevant to mass-casualty scenes, reinforcing feasibility for collapsed-structure search [2]. Device-forward prototypes aimed explicitly at confined-space victim search linked sensor selection, airflow, and I/O choices to USAR workflows, indicating practical design lanes for portable platforms [20]. From a biochemical standpoint, clinical and volatilomics literature consistently highlights amines

(including NH_3), reduced sulfur species (e.g., H_2S), ketones, and C_6–C_{10} aldehydes as elevated in major trauma, infected/non-healing wounds, and whole-body emissions, motivating fingerprint-based (pattern-recognition) detection rather than single-marker selectivity [14,15]. Proof-of-concept sensing with low-cost MOX arrays shows these devices can separate wound-related microbial volatilomes, supporting the feasibility of broad-spectrum arrays for injury/infection indication [16].

Beyond laboratory-scale experimentation, electronic-nose technologies have been increasingly deployed in real-world and end-user applications across disaster response, environmental monitoring, and industrial quality control. In the context of urban search and rescue, Anyfantis *et al.* reported a mobile e-nose system integrated into a USAR platform for real-time victim localization and hazard detection, explicitly evaluating usability and operational constraints under realistic field conditions [20]. Mobile and aerial deployments have also been demonstrated in environmental monitoring, where the RHINOS system was used for real-time odor quantification over active wastewater treatment facilities, highlighting the feasibility of lightweight, portable e-nose instruments mounted on robotic platforms [21]. In industrial settings, portable e-nose devices have been adopted for on-site quality assurance, such as the discrimination of cork taint compounds in the wine industry, enabling rapid screening outside laboratory environments [22]. More recently, electronic-nose systems have been employed for in-situ environmental methane monitoring, demonstrating robustness under outdoor operating conditions and supporting continuous field deployment [23]. Portable breath-analysis e-nose prototypes have also been evaluated with large user cohorts in online or point-of-care settings, further illustrating the transition of e-nose technology toward practical, user-facing applications [24]. Collectively, these studies indicate that electronic-nose systems have matured beyond proof-of-concept demonstrations and are increasingly being engineered to meet the constraints and requirements of real-world deployment.

Device-engineering patterns for deployable e-noses

The process involves active sampling and chamber design. Controlled airflow is repeatedly linked to faster response, greater signal amplitude, and improved repeatability; reviews and experiments recommend micro-fans/pumps and guided chamber geometries to stabilize exposure and mitigate boundary-layer effects [8,9]. These findings motivate our use of a micro-fan and a short, low-dead-volume path to the array.

The system controls environmental sensitivity and drift. MOX responses vary with temperature/humidity (T/H) and exhibit drift; firmware-level compensation and calibration materially stabilize outputs while remaining MCU-friendly. Studies demonstrate regression-based T/H correction and random-forest–based multi-variable compensation that reduce baseline wander and improve classification robustness, and complementary modeling describes humidity response power laws useful for lightweight compensation [25,26]. These techniques align with our pipeline.

On-device machine learning operates within the constraints of the MCU. To avoid cloud latency and fragile links during early search passes, recent portable e-nose designs execute classification on the microcontroller, validating real-time decision-making in handheld, battery-powered form factors [13]. This practice fits within the broader TinyML-style trend edge inference under tight RAM/flash and milliwatt budgets [27]. Our device follows this pattern by keeping inference local and streaming only compact risk scores/flags.

FMCW radar for presence confirmation and coarse localization

Compact FMCW radars can simultaneously recover range/angle and micro-motion (respiration/heartbeat), enabling a single sensor to both confirm human presence and provide coarse victim location in low-visibility, partially occluded scenes [19]. At 24 GHz, device-free schemes have demonstrated indoor localization and posture estimation from the radar cube using deep models, illustrating that practical angular/range bins are sufficient for human-level

geolocation in clutter [28]. At mmWave (120 GHz), narrow-beam FMCW systems have achieved non-contact respiratory and heartbeat monitoring, underscoring the headroom for micro-motion sensitivity relevant to triage [29].

For a portable e-nose device, these capabilities align naturally with our design goals: (i) use radar-derived range/angle bins and vital-sign surrogates to gate chemical alarms to verified human presence (reducing false positives from ambient plumes), and (ii) surface coarse bearing/range to guide responders or a robot toward likely victims both functions supported by FMCW literature that couples localization with vital-sign extraction [19,28].

Gas-distribution mapping

Gas dispersion in built environments is turbulent and intermittent, so robotic olfaction typically reconstructs spatial concentration fields from sparse samples (Gas-Distribution Mapping, GDM) and searches for likely sources via probabilistic/infotaxis-style policies; this provides a theoretical path to integrate e-nose payloads with exploration when mapping is needed [30,31]. Modern surveys in field robotics summarize odor-source localization (OSL) strategies gradient, bio-inspired/infotaxis, probabilistic, and learning-based demonstrating that portable gas sensors can be mounted on UGV/UAV platforms for hazardous reconnaissance [32]. In our paper, GDM/OSL remains theoretical context: the present device is portable and radar is used to assist users (presence confirmation and coarse bearing) rather than to perform radar-based mapping; however, the design choices (active sampling, synchronized timestamps, deterministic I/O) align with established GDM/OSL frameworks should future robotic integration be pursued [31,32].

Building on these findings, the present work introduces the proposed sensing platform in an active-sampling microchamber with on-device inference on an ESP32 and a 24 GHz FMCW radar for presence gating. Air is drawn by a micro-fan through a short, low-dead-volume path; signals are baseline-corrected and temperature/humidity-compensated, features are computed on a controlled hold window, and a J48 model issues a per-cycle risk score locally consistent with TinyML-style deployment [9,13]. The radar supplies micro-motion and range cues to suppress chemistry-only false positives and to provide coarse operator guidance [19]. Section 3 details the system architecture.

System architecture

Overview and operating modes

The SmellTec unit is a proposed system, multimodal detector designed for rapid casualty indication in visually degraded scenes. It combines different types of gas sensors that detect specific chemicals related to injuries or infections, measures environmental factors like temperature and humidity to adjust its readings, and uses 24 GHz FMCW radar to confirm if a person is present and give rough distance and direction information to help responders The enclosure supports handheld surveys and robot-mounted use (Wi-Fi telemetry to a base station or robot controller), with core computation on an ESP32-class MCU and tethered to computer for heavier inference and data logging. Keeping inference on the MCU follows established edge/TinyML-style practice for latency and resilience, while a companion SBC extends capability without blocking real-time operation [13,27].

Mechanically, the device is split into a front sensing module (intake, chamber, array) and a rear service bay (MCU, optional SBC, battery, fan driver, OLED/LED UI, Wi-Fi antenna). The device operates using a three-phase pneumatic cycle consisting of intake, hold, and purge phases, with timing controlled by the MCU to ensure repeatable exposure during analysis.; this is a widely used pattern in portable e-noses because it improves repeatability versus passive sniffing [13]. The radar sits behind the chamber wall and streams presence and micro-Doppler features over UART/GPIO for gating and user assistance (presence confirmation and coarse bearing), rather than for full mapping/SLAM [19,28].

Recent USAR-focused reviews have surveyed the practical application of gas sensor technologies for locating victims in mass-disaster scenarios. The review by Sunnucks *et al.* synthesizes laboratory studies, field trials, and prototype deployments, concluding that portable electronic-nose systems can provide early victim indication when conventional modalities fail [2]. However, the review also highlights recurring challenges, including sensor drift, environmental confounders, false positives from ambient plumes, and the lack of auxiliary modalities to confirm human presence [2]. Building on these findings, the present work focuses on device architecture rather than new chemical markers. Specifically, we address two gaps identified in prior USAR-oriented studies: (i) the need for deterministic, low-power embedded execution suitable for handheld or robot-mounted deployment, and (ii) the suppression of chemistry-only false positives through an independent physical presence cue. The proposed system integrates an actively sampled MOX array with on-device inference and 24 GHz FMCW radar presence gating, aiming to improve operational reliability while remaining compatible with microcontroller-class platforms.

Sensor suite: targets, roles, and rationale

Target chemical families. The array is designed to capture amines (incl. NH_3), reduced sulfur species (e.g., H_2S), ketones (e.g., acetone), and C_6–C_{10} aldehydes families repeatedly associated with trauma, infected/non-healing wounds, and whole-body emissions [14,15]. Instead of single-analyte selectivity, we use fingerprint-based detection (pattern recognition across channels), which is the recommended strategy for low-cost MOX e-noses in complex air [13,33]. The sensor suite and its nominal target gases are summarized in Table 1.

Table 1. Sensor array and key characteristics used for VOC data acquisition.

Sensors	Description	Advantages	Detected Gases
MQ-135	The gas-sensitive material used in the gas sensors is tin dioxide, which has low conductivity in clean air. When a polluting gas is present in the environment where the sensor is located, the sensor's conductivity increases with the concentration of polluting gases in the air.	monitoring of smoke and other harmful gases.	NH_3, Sulphide, Benzene vapor
MQ-136		High sensitivity to sulphur-containing organic vapours.	H_2S
MQ-137		High sensitivity to other organic ammonia.	NH_3
SGP30	The sensor reports TVOC and eCO estimates from a MOX cell using an on-chip model; it does not directly measure CO.	The device offers easy integration into a wide range of applications, low power consumption, and long life.	Volatile organic compounds and carbon dioxide
BME680	The BME680 contains a Metal Oxide Semiconductor sensor that detects VOCs in the air. The sensor enables you to qualitatively understand the total concentration of volatile organic compounds/pollutants in the surrounding air.	High precision, versatility, low power consumption and a wide range of applications.	Temperature, humidity, atmospheric pressure, gases, volatile organic compounds
MQ2	Its primary applications are fire detection, gas leak monitoring and general air quality assessment, the	The system’s sensitivity is enhanced to detect early-stage tissue decay	Methane, Propane, Butane,

	sensor can contribute meaningfully to a system designed to detect injured human bodies or related biological indicators.	or microbial activity.	Hydrogen, Smoke and Alcohol

Note: NH_3 = ammonia; H_2S = hydrogen sulphide.

Sensors are arranged radially in a cylindrical chamber with front intake and rear extraction so each element experiences a similar boundary layer and residence time, improving repeatability under flow. The arrangements of the device can be seen on Fig. 1. This geometry is consistent with reports that flow field and chamber shape materially affect response amplitude/latency, and that guided/bionic chambers improve uniform exposure [8,9].

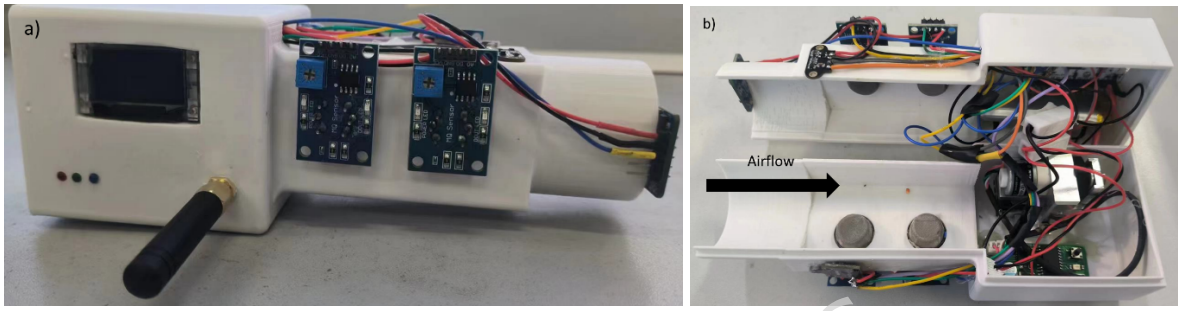

Figure 1. SmellTec prototype. (a) External view. (b) Internal layout: active-sampling chamber, MOX array and BME680, micro-fan/hatch, ESP32, 24-GHz radar, battery, and power/I-O. Airflow direction is indicated on the picture.

Reason why we pick MOX with environmental sensors because MOX arrays are sensitive to temperature/humidity and drift, so companion T/H/P channels enable firmware-level compensation that stabilizes baselines and reduces false alarms in variable environments, a proven approach for embedded devices [26]. While combining broad MOX with environmental context is standard practice in deployable e-noses [13].

The MCU samples gas and environmental channels during the hold phase of each cycle and timestamps the radar presence/micro-motion features for sensor-level fusion. Synchronized timing is recommended in portable e-noses to avoid phase errors between concentration transients and compensation variables [13].

Sampling chamber and airflow

The sensing core is a front-open cylindrical chamber with the gas and environmental sensors placed radially around the inner circumference; a rear zone houses the actuator hardware (fan with hatch). This layout exposes all MOX channels to a similar boundary layer and residence time, improving sampling consistency across the array. Empirically, flow field and chamber shape are known to affect response amplitude and latency in e-noses; guided geometries that reduce stagnation and deliver uniform exposure improve repeatability [8,9].

During operation, the MCU drives a three-phase pneumatic cycle intake to hold and then purge using a rear-mounted microfan for controlled inflow and a servo-actuated hatch to isolate the sample during analysis. Active sampling is standard in portable e-noses and helps stabilize exposure versus passive sniffing; paired with short, guided chambers it reduces dispersion artifacts and increases signal-to-response at the sensor surface [8,9,13]. Three-phase cycle idea will be detailed as the following:

- Intake. The fan ramps to a fixed set-point (nominal 0.5 $L \cdot min^{-1}$) while the rear hatch is open to draw ambient air through the front aperture; the goal is to reach a repeatable chamber renewal time (τ) before analysis. Actively controlled inflow reduces the influence

of external drafts and head-on approach angles—factors flagged as physical confounders in e-nose hardware [8].

- Hold. After a fixed intake duration (e.g., 1–2 s), the servo closes the hatch, freezing the sample for 0.5–2 s while the array is read at high rate; this minimizes dilution and pressure perturbations that distort rising edges.
- Purge. The hatch re-opens and the fan drive a washout (for about 1-3 seconds) to shorten tail memory; purge duration is auto-extended if baseline hasn't re-stabilized. Short, repeatable purge cycles align with chamber-optimization results showing faster recovery when flow is controlled [9].

During the intake phase, airflow is generated by a micro-fan mounted downstream of the sensing chamber. The nominal airflow rate used in this study was approximately 0.5 $L \cdot min^{-1}$, selected to ensure repeatable chamber renewal while limiting turbulence and power consumption. The airflow rate is adjustable in the range of approximately 0.3–0.8 $L \cdot min^{-1}$ through pulse-width modulation of the fan drive, enabling adaptation to different chamber volumes or operating conditions.

The chamber co-locates a BME680 for T/H/P and uses its readings both to guard fan set-points (avoid condensation regimes) and to feed firmware-level T/H compensation alongside baseline correction, best practice for MOX arrays where temperature/humidity and flow act as dominant confounders [8]. The design keeps short, low-dead-volume paths and avoids absorbent foams/filters in the analysis path, consistent with chamber studies that caution against VOC memory effects in porous media [9].

Mechanical layout and field serviceability are deliberate to protect sensing quality and simplify upkeep. The rear bay consolidates the ESP32, the optional Raspberry Pi, battery, fan driver, Wi-Fi antenna, and OLED/LED UI, with the hatch servo mounted to an internal bulkhead; keeping heat-generating electronics and moving actuators behind the chamber wall helps stabilize the sensing volume and makes board or battery swaps straightforward. The separation between the sensing module and service bay also helps keep temperature and airflow changes inside the chamber to a minimum, leading to more consistent readings during active sampling. Calibration and cleaning are supported by a fixture-friendly front end so bench work mirrors field operation. The intake rim accepts a push-fit collar for tubing during exposures, enabling guided inflow and repeatable fixture-based tests which is good for portable, fan-controlled e-noses designed for real-time inference [13]. Timing, telemetry, and synchronization are also very important. The firmware timestamps each intake to hold then purge cycle, acquires gas and T/H primarily during hold, and ingests radar presence/micro-motion features as a continuous stream over UART/GPIO for fusion and gating; clocked acquisition and coordinated fan control are emphasized in portable e-nose implementations precisely to avoid phase errors between transients and compensation variables [13].

These engineering choices matter because disaster air is turbulent, nonstationary, and unforgiving. Uncontrolled drafts, temperature swings, and intermittent plumes can inflate false alarms or bury weak injury-linked VOC fingerprints; combining active sampling, guided chamber geometry, and T/H compensation directly targets those error sources, as shown in engineering reviews and bionic-chamber experiments , while retaining a compact, battery-fit pneumatic stack suitable for handheld passes or mounting on a mobile base [8,9,13].

Radar-assisted presence & coarse localization

The radar in SmellTec is included to assist the user: it verifies that a human is actually nearby (via micro-motion) and supplies a coarse bearing/range cue that helps an operator or a host robot orient toward a suspected victim, rather than attempting full mapping or SLAM. This design choice is grounded in FMCW studies showing that a single unit can recover range/angle while simultaneously extracting vital-sign surrogates (respiration/heartbeat) in cluttered indoor environments and that 24 GHz setups can perform device-free human localization from radar-cube features at practical standoff distances [19,28]. Functionally, the

24 GHz FMCW module emits linear chirps and processes reflections into range–Doppler (and angle, with multiple receive channels) bins; slow chest motion yields low-Doppler micro-motion energy at respiratory rates (~0.1–0.5 Hz), with occasional heartbeat fundamentals/harmonics under favorable SNR, a well-documented phenomena in the FMCW literature [19]. Work at mmWave (≈120 GHz) further confirms non-contact vital-sign tracking with narrower beams and higher resolution, underscoring the headroom for micro-motion sensitivity when line-of-sight is partially occluded by dust or fabric [29]. Radar presence-detection performance under representative conditions is summarized in Table 2.

Table 2. Radar performance

Scenario	Presence detected
Standing subject	✓
Seated subject	✓
No person	✗ (No live being in the range)
Person behind obstacle	✓ (reduced confidence)

SmellTec uses the FMCW radar channel primarily to contextualize and constrain chemical decisions rather than to perform mapping. Radar micro-motion cues gate chemical alarms, so alerts are only raised when evidence of nearby human presence is detected, which reduces susceptibility to unrelated background plumes. The radar also provides simple directional guidance by reporting the most active range bin and, when available, coarse bearing information inferred from angle cues or device yaw during a brief sweep. Similar 24 GHz FMCW approaches have been demonstrated for device-free localization and posture estimation using learned feature representations [19,28]. Overall, radar improves decision robustness and guides operator search, while the e-nose contributes the VOC-based signature used for injury or infection indication [19,28].

From an implementation standpoint, the radar streams presence flags and compact feature summaries (e.g., respiration-stability over short windows, range-bin indices) at tens of hertz over a simple UART/GPIO link. The MCU aligns the radar data with the e-nose's intake–hold–purge cycle to ensure that gas-phase transients and motion evidence are fused without phase error, a synchronization practice emphasized in embedded FMCW vital-sign/localization work. This interface keeps the radar lightweight and deterministic in the power budget while delivering the user-assistance functions the device requires, not a full radar-mapping stack [19,28].

Embedded computing and power hierarchy

The compute architecture prioritizes deterministic, on-device decisions on an ESP32; training and heavier analysis run on a PC. In operation, the ESP32 performs multi-channel ADC sampling, fan/hatch control, baseline and temperature/humidity (T/H) compensation, feature extraction, and local classification at ~1–2 Hz. This mirrors best practice in portable e-nose systems that keep inference local to avoid cloud or link latency and follows TinyML patterns for running models within tight RAM/flash and milliwatt budgets [13,27].

Model training is performed off-device. Logged sessions are used to learn a J48 decision tree and a Random Forest (RF), choose and calibrate thresholds, and export a compact tree for deployment to the ESP32. Training large models on a desktop and deploying a compressed or interpretable version to the MCU balances accuracy and latency [34,35]. When a tethered PC

is present (USB/serial or Wi-Fi), the ESP32 can mirror features for host-side RF inference and logging while still issuing a local, real-time decision; the on-device classifier remains the single source of truth for the UI/buzzer to avoid I/O back-pressure [13,27].

Power budgeting treats the fan and 24 GHz radar as dominant loads. Both are duty-cycled around the intake–hold–purge schedule so airflow and radar micro-motion snapshots are available when needed, without violating timing guarantees for synchronized fusion [13]. This keeps the device battery-fit while preserving the determinism expected of embedded instruments at the edge [27]. A lightweight host utility mirrors hold-window features for training and optional tethered inference; a baseline session is shown in Fig. 2 and tethered run during propane exposure is shown in Fig. 3.

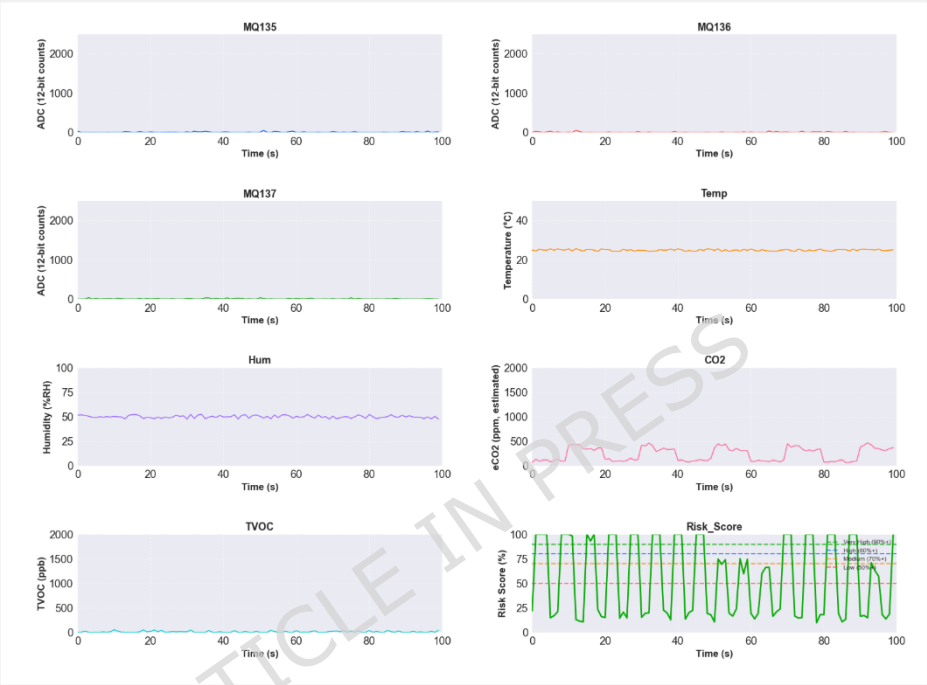

Figure 2. Representative baseline session in tethered development mode. Real-time streams from MQ135/136/137 (12-bit ADC), temperature, humidity, eCO_2 (estimated), and TVOC are shown during repeated intake–hold–purge cycles under clean-air conditions. The risk score is computed once per cycle from the hold-window feature vector and is shown to illustrate the end-to-end pipeline and timing synchronization.

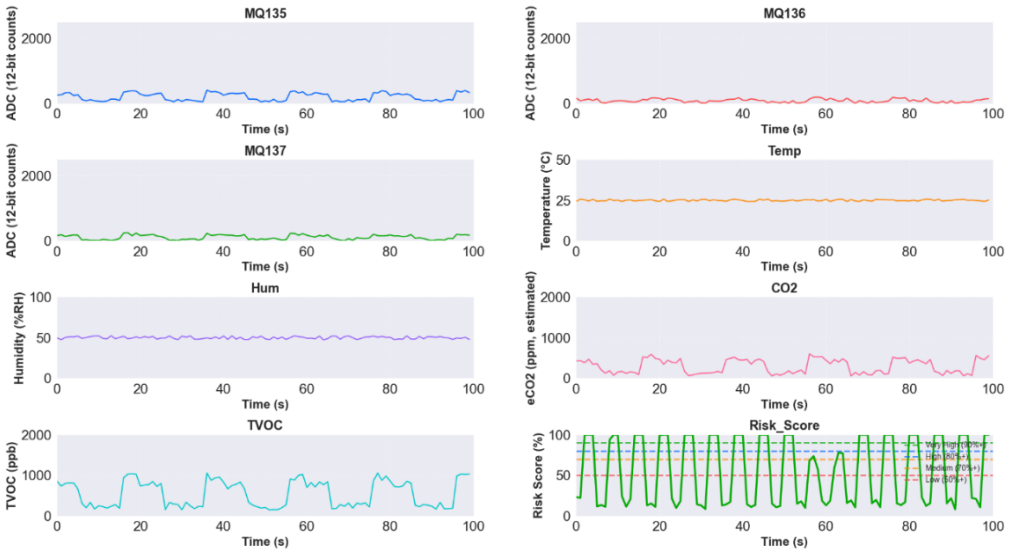

Figure 3. Tethered host snapshot showing real-time streams from MQ135/136/137 (ADC, 12-bit), temperature (°C), humidity (%RH), eCO_2 (ppm), and TVOC (ppb) during exposure to propane gas at 1 Hz. The right panel controls capture and training; the Risk Score trace reports the per-cycle J48 probability.

The real-time pipeline is synchronized to the pneumatic cycle so that each cycle yields one fused decision. During intake, the fan reaches a fixed flow. During hold, the ESP32 applies baseline correction and T/H compensation using the environmental channels, steps repeatedly shown to stabilize MOX arrays under real-world variability [10]. Still within hold, the firmware extracts a windowed feature set (level/median, σ or MAD, slope or rise time, max or integrated area, and inter-channel ratios), which is a feature style validated in portable, embedded e-noses [13]. The end-to-end presence-gated decision pipeline is summarized in Fig. 4.

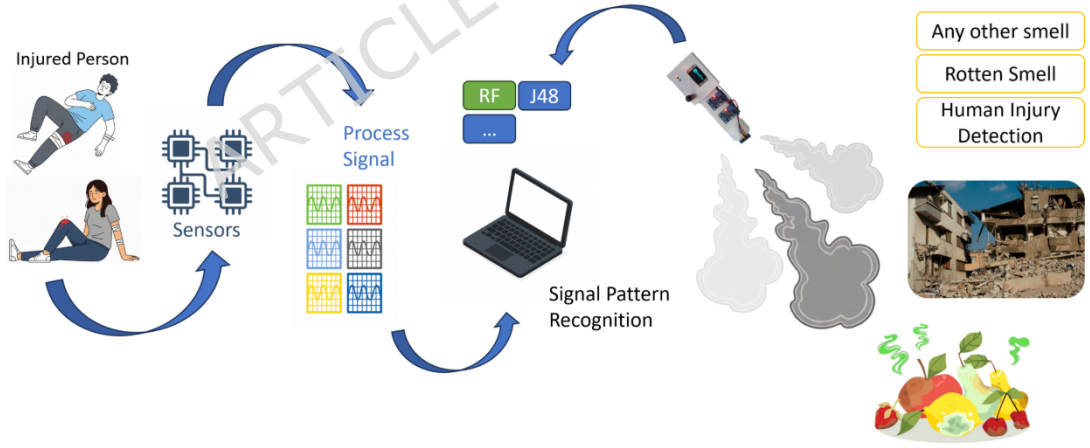

Figure 4. Illustrates the end-to-end workflow for presence-gated chemical detection. 24 GHz radar supplies presence/micro-motion evidence; the MOX array with active sampling produces a windowed feature vector; an on-device J48 model issues a per-cycle risk score, optionally fused with a host-side Random Forest. The final alert is gated by radar presence.

In parallel, the 24 GHz FMCW radar provides a presence flag, a short-window respiration-stability score, and a range-bin index. Chemistry and motion are time-aligned; a no-presence state gates chemical alarms to zero, consistent with FMCW studies that pair localization with micro-motion for human verification. The classifier computes a J48 probability on the ESP32

for low latency. When tethered, the ESP32 can mirror features to the PC for RF inference and logging. If a PC is absent, the on-device J48 alone drives the UI/buzzer. The alert threshold is chosen at training using Youden's to balance sensitivity and specificity [36]. Each cycle outputs a risk score, a binary alert, and, from radar, coarse bearing/range cues surfaced to the operator. The main hardware components and key specifications are summarized in Table 3.

Table 3. Hardware Components and specification

Component	Model / Series	Key Specifications	Role
MCU	ESP32-WROOM-32	240 MHz, 520 KB SRAM, Wi-Fi	Main controller
Radar	LD2410C	24 GHz FMCW microwave Presence radar, Rx channels, UART	Presence gating
SBC (optional)	Raspberry Pi 4 Model B	4 GB RAM, Linux	Tethered inference
Fan	40x20mm PWM Fan	0.2–1.0 L·min^{-1}, X mA	Active sampling
Battery	Li-ion 18650	2600 mAh, 5V	Power
Display	OLED SSD1306	0.96", 128 x 64 Pixels, I^{2}C	UI
Antenna	2.4 GHz PCB antenna	2 dBi	Wireless

Recent work on low-power prescriptive and resource-aware edge systems emphasizes deterministic execution, explicit control of computational budgets, and local decision-making to reduce dependence on cloud connectivity and unstable links. Frameworks such as LP-OPTIMA demonstrate that enforcing resource constraints at design time improves system reliability and energy efficiency in embedded IoT deployments. The SmellTec architecture follows a similar philosophy by prioritizing MCU-level inference, fixed per-cycle timing, and duty-cycled sensing components, ensuring predictable behavior under constrained power and compute budgets.

On-device signal processing & features

We treat every decision as a single "hold-window" snapshot which means that all computations happen on the short, stable analysis interval that follows intake and precedes purge, so the math is tied to a reproducible piece of data rather than a drifting stream. This separation between a controlled window and the signal math here mirrors portable, battery-powered e-nose practice for reliable embedded inference and pairs naturally with guided chambers that reduce sampling artifacts.

Baseline removal and slow drift control use a robust reference from the pre-intake ambient window W_0 and a slow EWMA to track hours-scale drift. For sensor i with raw signal $X_i(t)$,

$$b_i^{(k)} = \lambda b_i^{(k-1)} + (1 - \lambda)\text{median}_{t \in w_0} x_i^{(k)}(t) \quad (1)$$

$$d_i(t) = X_i(t) - b_i^{(k)} \quad (2)$$

We then carry only the detrended signal $d_i(t)$ forward. Explicit drift handling is recommended because MOX baselines change over time and can degrade diagnostic reliability if left unchecked [37,38].

Lightweight temperature/humidity (T/H) compensation is applied either at the sample level or at the feature level, using absolute humidity H_a for better physical invariance than RH. We compute H_a with a Tetens-style saturation vapor pressure (SVP) model:

$$\text{SVP}(T) = 6.112\exp\left(\frac{17.67T}{T + 243.5}\right)\text{hPa} \quad (3)$$

$$H_a = \frac{216.7(RH/100)\text{SVP}(T)}{T + 273.15}\,\text{gm}^{-3} \quad (4)$$

and correct each channel with either a linear or humidity power-law form:

$$\grave{d}_i(t) = \beta_{0,i} + \beta_{1,i} d_i(t) + \beta_{2,i} T(t) + \beta_{3,i} H_a(t) \quad (5)$$

Humidity power-law compensation for MOS gas sensors has been validated on VOC tasks, while multivariable learning-based compensation is also reported and can remain MCU-feasible with constrained model size [10,25]. Outlier-robust window statistics trim the top/bottom 2–5% of samples in the hold window before feature computation, limiting the impact of brief pneumatic or ADC spikes without adding latency, an embedded trick consistent with portable e-nose implementations [13].

The feature vector emphasizes MCU-friendly descriptors that capture level, spread, and kinetics of $\grave{d}_i(t)$ at the sampling rate f_s over ahold window of length L:

- Level/dispersion: $\mu_i = \frac{1}{L}\sum \grave{d}_i; \sigma_i(\text{std.}); \wedge \text{MAD}$
- Extrema/integral: $m_i^{max} = \max \grave{d}_i; A_i = \sum_t \grave{d}_i \Delta t$
- Kinetics: least-squares slope S_i vs. time; rise time $t_{90,i}$. if the trace is not monotonically rising, we set $t_{90,i} = \text{NaN}$ and pass a validity bit to keep inference stable.
- Cross-channel structure: stabilized log-ratios $r_{ij} = \log\frac{\mu_i + \varepsilon}{\mu_j + \varepsilon}$ that separate broadband from targeted responses (e.g. MQ-135 vs MQ-137), with ε set to 1% of the median absolute training magnitude to ensure scale robustness.

Such time-domain features remain strong baselines for embedded e-noses and are well documented in surveys and moving-window variants [39,40].

Tight synchronization with the 24 GHz FMCW radar ensures chemistry and motion describe the same moment. From the radar micro-motion stream $r(t)$ we form a compact tuple aligned to the hold window: a presence flag (energy threshold in a human range bin), a respiration-stability score ρ uses a Welch PSD over the hold window:

$$\rho = \frac{\max_{f\in[0.1,0.5]} P_r(f)}{\sum_{f\in[0,1.5]} P_r(f) + \in} \quad (6)$$

optionally multiplied by the peak's prominence above a local median. FMCW vital-sign studies establish the micro-Doppler content in the 0.1–0.5 Hz band and show that 24 GHz systems can yield device-free human localization cues [19,28]. Quantization and scaling are treated as part of the algorithm so inference stays deterministic on the ESP32: we learn per-feature μ_f,σ_f offline, apply on-device $z_f = (f - \mu_f)/\sigma_f$ and then quantize to Q15 fixed-point before classification, an established TinyML deployment pattern that fits RAM/flash and power budgets [27]. Resource-aware numeric budgets demonstrate MCU feasibility and guide reproducibility which the hold window is 0.5–2.0 s at 25–50 Hz (12–100 samples/channel), the feature vector stays ≤ 80 features including radar tuple, and the on-device J48 has depth ≤ 6 with all arithmetic in Q15. These choices mirror published portable e-nose runtimes and edge-ML constraints [27]. Putting everything together, the final result will be

$$\left[\mu,\sigma,\mathrm{MAD},m^{\max},A,s,t_{90},r_{ij} \vee \mathrm{presence},p,k_r\right] \quad (7)$$

The ESP32 delivers the current hold window to the on-device J48 decision tree for a real-time risk score; when tethered, it may also mirror the same features to a PC for heavier analysis without delaying the local decision.

Hybrid j48–random forest model and gated decision logic

The decision layer is constructed to preserve deterministic, low-latency inference on an ESP32 while enabling heavier ensemble analysis on a tethered PC without perturbing real-time behavior. The on-device classifier is a compact J48 decision tree, a choice that is widely used for embedded, interpretable inference on tabular sensor features [34]. Training larger models off-device and deploying a compact version to the MCU follows established TinyML-style edge ML practice under tight RAM/flash and power constraints [27]. For robustness, a Random Forest (RF) is trained on the same features and queried only when a PC is connected, reflecting the strong empirical performance of tree ensembles on noisy, tabular sensing data [35].

Model development proceeds off-device on logged hold-window instances with labels corresponding to the target alarm state. To avoid over-optimism from slow environmental drift or session effects, splits are grouped by session/day so that no group contributes to both train and test a standard leakage control strategy for time-correlated sensing data [41]. Hyperparameters are selected by inner validation (tree depth, minimum leaf size; number of trees, feature subsampling for RF), and the final model is retrained on the union of training folds before export.

Both J48 and RF produce probability-like scores; where necessary, probabilities are calibrated by isotonic or Platt-style post-processing on a held-out set, which is known to improve probability quality for tree-based models [42]. The operating threshold τ^* is chosen by maximizing Youden's index $J = \mathrm{TPR} + \mathrm{TNR} - 1$ on the validation set, a criterion that balances sensitivity and specificity in binary alarms, with ROC curves reported for completeness [36,41]. The selected τ^* is stored alongside the exported model.

At runtime, after section 4 features are computed for the current hold window, the ESP32 evaluates the J48 tree to obtain a probability p_{j48} [34]. In parallel, the radar front-end supplies a presence flag and respiration-stability metric aligned to the same window; micro-Doppler signatures in the 0.1–0.5 Hz band and device-free localization cues at 24 GHz support the reliability of these surrogates for human presence [19,28]. The system then enforces presence gating so that a chemical alarm cannot fire without verified micro-motion in the same window:

$$\text{Alert} = 1 \Longleftrightarrow (\text{presence} = 1) \wedge \left(p_{j48} \geq \tau^*\right) \quad (8)$$

When the radar is unavailable or degraded, the pathway degrades gracefully to the chemical classifier alone while surfacing a reduced-confidence indicator; this maintains deterministic operation on the MCU while acknowledging the missing physics channel [19,27]. If a PC connection is present, the ESP32 mirrors the same feature vector to the host for RF scoring. When an RF probability p_{RF} is returned within the cycle budget, the system computes a probability-level fusion

$$p_{hyb} = \alpha p_{j48} + (1 - \alpha) p_{RF}, \alpha \in [0,1] \quad (9)$$

then applies the same presence gate and threshold τ^*. Linear score fusion across complementary classifiers is a well-studied, stable method for combining different models [43]. If the host response is late or unavailable, the decision for that cycle remains based on p_{j48} to preserve MCU-level latency guarantees. To reduce unnecessary host queries and stabilize UX, an early-exit band δ around τ^* is defined from validation statistics. If $p_{j48} \leq \tau^* - \delta$ or $p_{j48} \geq \tau^* + \delta$, the device finalizes locally; only ambiguous windows $\left|p_{j48} - \tau^*\right| < \delta$ trigger a host RF request. This "edge-first, host-assist" policy is consistent with latency-aware TinyML deployment patterns. The OLED presents a three-level indicator (green/amber/red) with fixed thresholds $\tau_{warn} < \tau^* \wedge \tau_{high} > \tau^*$ derived from the same ROC analysis, so that "warn" favors sensitivity for triage while "high" favors precision for resource commitment. Each cycle's features, probabilities, gate state, and final alert are timestamped and logged, enabling future offline retraining to refine $\alpha, \tau^*, \wedge \delta$ without altering the embedded pipeline—a device-development pattern recommended for portable e-nose systems and edge ML [27].

Under typical duty-cycled operation, including periodic fan activation and radar presence checks synchronized with the sampling cycle, the system draws an average current on the order of 80–120 mA. When powered by a 2600 mAh lithium-ion cell, this corresponds to an estimated operating time of approximately 18–30 hours, depending on sampling frequency and radar duty cycle. These values provide an order-of-magnitude estimate of energy consumption and are intended to demonstrate embedded feasibility rather than to represent optimized field endurance.

Device verification & offline model evaluation

While the ESP32 executes the full sensing and inference pipeline in real time, tethered operation was used in this study to log features and evaluate model performance offline. During development the ESP32 executes the full timing loop: intake–hold–purge control, multi-channel ADC sampling, baseline and temperature/humidity compensation, and feature extraction. Each hold-window feature vector is mirrored to a host PC for scoring and logging. This device-first verification approach is standard in portable e-nose builds and aligns with TinyML guidance to prove latency, RAM/flash, and timing on the target MCU even when training and early evaluation run on a host [13,27].

A fully assembled unit integrates a multi-sensor MOX array, a BME680 for environmental channels, a micro-fan with a short, low-dead-volume path, an ESP32 with an OLED UI, and a 24 GHz FMCW presence radar. In a representative baseline session, the host utility shows stabilized 1 Hz streams and returns a per-cycle risk score based on the mirrored features while the MCU timestamps intake, hold, and purge for repeatability. This setup matches handheld architectures that combine compact pneumatics with embedded inference to avoid link latency [13]. The operational definitions of the exposure classes are summarized in Table 4. A representative tethered log illustrating synchronized sensor streams and per-cycle risk scores is shown in Fig. 5.

Table 4. Operational definition of dataset classes used for model training and evaluation

Class	Description	Source
FreshAir	Ambient laboratory air	Baseline
Injured	VOC surrogate mixture (NH_3-dominant)	Controlled
NH_3	Diluted ammonia	Gas bottle
Propane	Controlled propane exposure	Gas bottle

The 'Injured' label refers to a controlled surrogate VOC condition intended to emulate elevated human-related chemical signatures rather than verified biological injury.

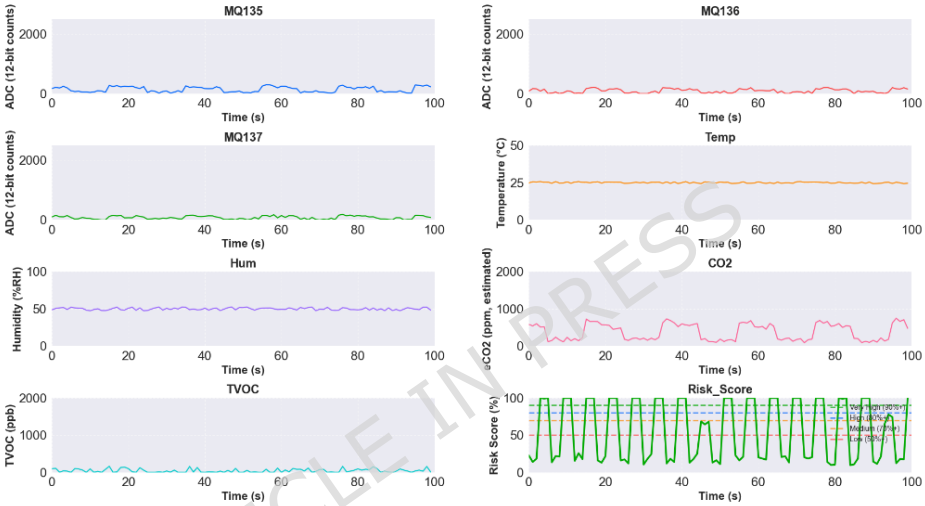

Figure 5. This tethered-mode snapshot shows synchronized streams from MQ135, MQ136, MQ137 (ADC, 12-bit), temperature (°C), humidity (%RH), eCO_2 (ppm, estimated), TVOC (ppb), and the per-hold risk □ produced by the on-device pipeline. Subtle cycle markers indicate the repeating intake–hold–purge schedule. All signals were logged while the ESP32 controlled airflow while exposed to NH_3 gas sample.

The decision policy is presence-gated. For each hold window the classifier outputs a risk □ that is compared with an operating threshold τ^* chosen on validation using Youden's □. A chemical alert is surfaced only when the radar reports human presence and □ ≥ τ^*. Fig. 6 illustrates this timeline from tethered-mode logs: binary presence is aligned with the same hold windows, alerts appear only when presence is one and the risk clears τ^*, and cycle bands make the intake–hold–purge schedule explicit. The use of FMCW micro-motion and device-free localization to confirm human presence is supported by prior 24 GHz studies in cluttered indoor scenes [19].

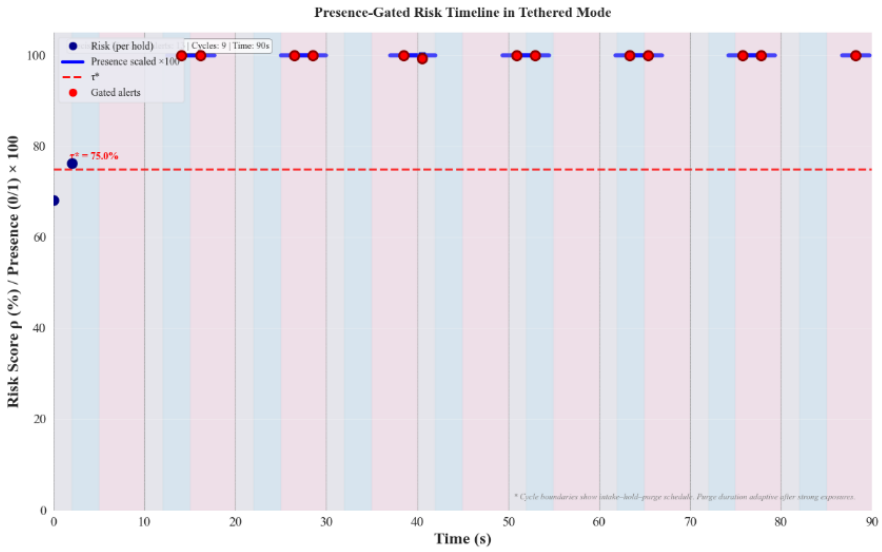

Figure 6. Presence-gated decision timeline from tethered logs. Each hold yields one risk ρ; the dashed line is τ^*. Alerts (red) appear only where presence=1 and ρ ≥ τ^*; cycle bands show intake–hold–purge. Signal on the graph were logged during exposure to NH3.

Presence gating is verified with a stationary volunteer or a low-amplitude motion surrogate at ~0.2–0.3 Hz; the 24 GHz FMCW module supplies a presence flag and a short-window respiration-stability score aligned to the same hold window as chemistry, and chemical alerts are suppressed when micro-motion is absent, consistent with reports of micro-Doppler vital-sign extraction and device-free localization at 24 GHz in cluttered environments [18,19,28]. When radar is unavailable the pathway reverts to chemistry-only classification and surfaces reduced confidence, preserving deterministic MCU operation.

Offline evaluation on the development dataset complements the device checks. The host PC reports holdout test-set metrics for J48 and Random Forest classifiers trained on 1,400 samples and evaluated on 600 held-out samples. Performance is characterized through ROC and precision-recall curves with AUC/AP scores, confusion matrices, and per-class precision/recall/F1 metrics. The 70/30 train-test split employs session-level grouping to prevent temporal leakage, ensuring that samples from the same acquisition session remain together during splitting [44,45]. These offline results quantify separability on logged data and are not presented as clinical sensitivity or specificity; feature-importance charts are included only to aid engineering interpretation, not to claim single-analyte specificity.

The offline evaluation presented in Figures 6 and 7 is based on a dataset comprising 2,000 hold-window samples collected under controlled laboratory conditions using a tethered eNose system. The dataset consists of four balanced classes, with 500 samples per class: FreshAir (ambient laboratory air baseline), Ammonia (controlled NH_3 exposure at 200-250 ppm), Propane (controlled propane exposure at 300-450 ppm), and Injured (replicated VOC profile based on published trauma biomarker studies, with concentrations adjusted for MQ sensor detection range). Each class was collected across 5 independent acquisition sessions conducted on different days. Within each session, 100 samples were extracted from synchronized intake–hold–purge cycles (2s intake, 3s hold, 5s purge). Each sample represents sensor readings during the 3-second hold window when gas concentration stabilizes. To avoid data leakage from temporal correlation, training and validation splits were grouped by session: all samples from the same session were kept together during train/test partitioning. The dataset was split 70% training (1,400 samples) and 30% testing (600 samples), with stratification ensuring balanced class representation.

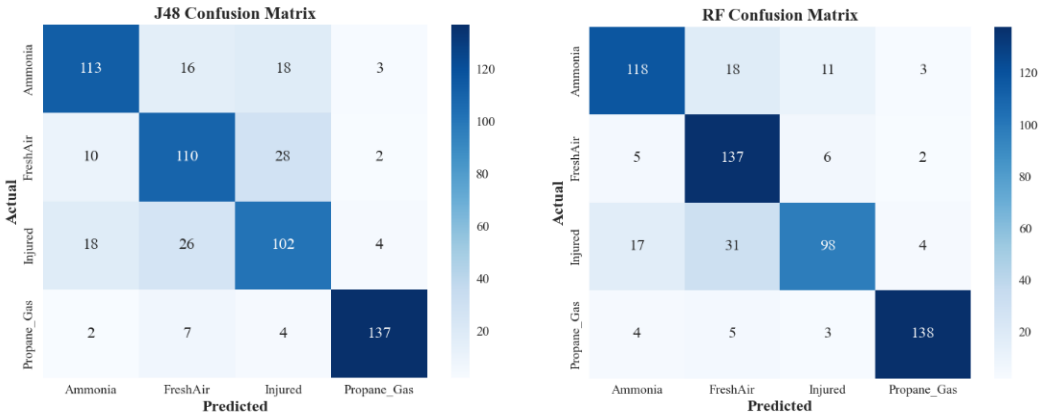

Figure 7. Offline evaluation showing confusion matrices for J48 and Random Forest classifiers. Results computed over 2,000 hold-window samples (500 per class) across four balanced exposure classes collected over 5 independent sessions per class. Dataset split 70% training / 30% testing with session-level grouping to prevent temporal leakage. Matrices illustrate instrument-level separability under controlled laboratory conditions and are not intended to represent clinical or diagnostic performance.

Fig. 7 summarizes the classification performance of the evaluated models using confusion matrices, ROC curves, and precision–recall curves. These results are intended to assess instrument-level separability between controlled exposure conditions rather than to report diagnostic sensitivity or specificity. Clear separation is observed for chemically distinct classes such as Propane, while confusion between FreshAir and Injured reflects overlapping low-concentration VOC backgrounds under controlled conditions.

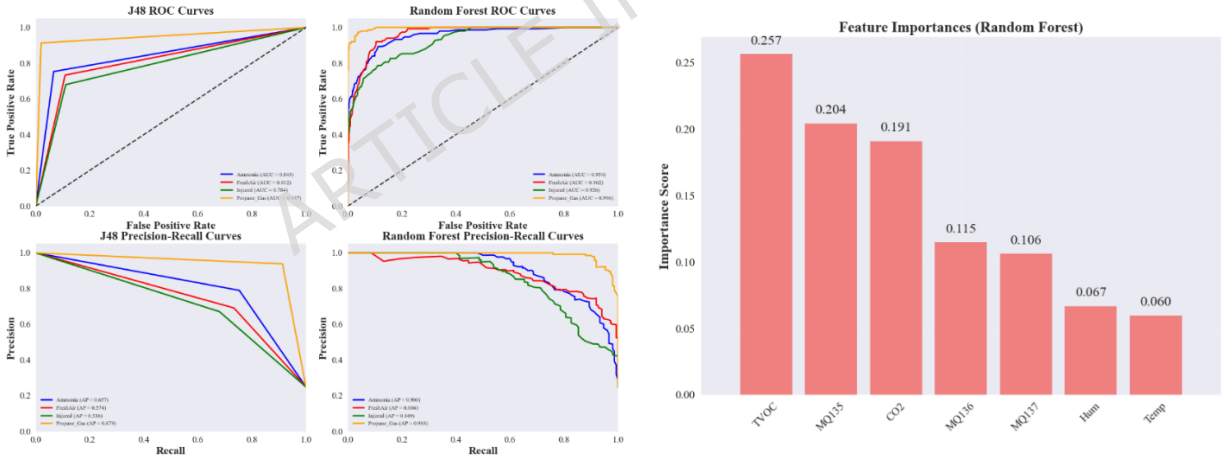

Figure 8. Feature-importance analysis is used to interpret sensor and environmental contributions at the instrument level.

Fig. 8 presents feature-importance and model-behavior analyses, illustrating that environmental variables and sensor kinetics contribute substantially to class discrimination. These plots provide insight into the sensing and feature-extraction pipeline and support the feasibility of compact, embedded inference, rather than serving as evidence of biological specificity.

Conclusion

This work demonstrates device-level feasibility rather than clinical detection performance. The platform unifies a metal-oxide sensor array targeted at injury/infection-linked VOC families, an active-sampling chamber that stabilizes exposure during an intake–hold–purge cycle, on-device inference on an ESP32 using a compact decision tree (with PC-side training), and 24 GHz FMCW micro-motion sensing to verify human presence and supply coarse bearing/range cues. The engineering choices-controlled airflow and short, guided chambers for repeatable responses, embedded inference under TinyML-style constraints and FMCW micro-Doppler for presence confirmation collectively produce a deterministic, battery-fit instrument suitable for handheld or robot-mounted reconnaissance. Instrument-level verification demonstrates that synchronized sampling, compensation, feature extraction, and gated decisions can be performed at the edge with bounded latency and resource footprint, establishing technical viability of the device concept ahead of larger field validations.

Baseline trials under clean air provide the reference needed to interpret gas-specific responses. Under identical cycle settings, ammonia exposure produced 2.7–4.0× larger peak-to-peak excursions than baseline across MOX channels (MQ135, MQ136, MQ137), with an average signal-to-noise ratio of 6.5 dB. Propane exposure produced 2.6–6.5× larger excursions across MOX channels, with an average signal-to-noise ratio of 7.8 dB. These ratios confirm that observed responses significantly exceed measurement noise. The dominant error mode for both J48 and Random Forest classifiers is misclassification between FreshAir and Injured classes, reflecting overlapping low-level VOC signatures rather than sensor failure. The Injured class represents a laboratory-replicated VOC profile with concentrations elevated for sensor detectability, whereas real breath biomarkers occur at ppb levels. This overlap is expected given the similar magnitude of baseline environmental VOCs and the simulated injury signature.

In operational deployment, this limitation is mitigated through sensor fusion with radar-based presence detection: chemical alarms are suppressed unless concurrent micro-motion evidence confirms human presence. Consequently, FreshAir ↔ Injured misclassification does not trigger false victim localization in unoccupied spaces, preserving operational safety and reducing false alarm rates. This multi-modal approach demonstrates how complementary sensor modalities can compensate for individual limitations while maintaining high sensitivity for actual victim detection.

Limitations and future work

This paper evaluates a prototype, device-focused platform for presence-gated detection of human-related volatile signatures. The results characterize instrument behavior, including timing, on-device processing, and offline separability, rather than diagnostic sensitivity or specificity. To move beyond feasibility, we will run controlled exposures with known concentrations, standardized scenarios, and field trials with verified ground truth. For deployment, operating points will be selected from ROC curves using Youden’s index so the sensitivity–specificity trade-off is clear.

The Low-cost MOX sensors on the sensing front end are broad-band and affected by temperature and humidity; compensation reduces drift and cross-sensitivities but does not remove them. Absolute sensitivity is lower than professional instruments for the same chemical families. For example, photoacoustic spectroscopy reaches sub-ppb ammonia in laboratory settings and compact PAS devices report low-ppm to sub-ppm limits in fieldable formats, while real-time mass-spectrometric tools such as SIFT-MS and PTR-MS provide ppb-level VOC quantitation with compound identification. We accept this trade to gain portability, faster response, and low power for triage use.

The verification and data collection were performed in tethered mode, where the ESP32 handles sampling, compensation, and feature extraction, while a host computer performs classification and logs results. We will next freeze the model on the MCU, repeat the same validation untethered, and report end-to-end latency, duty-cycle power, and battery life. On the tethered-mode dataset (N = 2,000 hold-window samples; four balanced classes with 500 samples each: FreshAir, Ammonia, Propane, and Injured), data were collected across five independent sessions per class (100 samples per session) to ensure temporal diversity. Training employed a 70/30 train-test split with session-level grouping to prevent temporal leakage. The confusion matrices in Fig. 7 show J48 at 77% overall accuracy and Random Forest at 82%. The main error mode for both models is FreshAir ↔ Injured (J48 off-diagonals 26 and 28; RF 31 and 6), whereas Propane_Gas is consistently isolated (J48 137/150, RF 138/150 correct). Macro-F1 tracks these results (≈ 0.77 for J48; ≈ 0.82 for RF) and aligns with the ROC/PR curves. Values shown are counts with rows = actual and columns = predicted.

Radar-based presence gating has edge cases. Apnea, posture, occlusion, or low signal-to-noise can mask micro-motion; when radar is unavailable or degraded, the system falls back to chemistry-only decisions and shows reduced confidence. We plan to add adaptive presence thresholds, brief multi-angle sweeps, and to test higher-frequency mmWave modules for improved micro-motion sensing. For the timing and recovery part, it needs more polishing. Active flow speeds up rise and increase signal magnitude, but full washout after strong exposures can take more than one purge, especially at high humidity. Routine reporting will include rise time, peak amplitude, and washout time as instrument metrics, and adaptive purge will be added so recovery is extended only when needed. Several auxiliary channels are indicative rather than absolute. TVOC and eCO_2 from the IAQ device are algorithmic estimates intended for indoor-air-quality trending, not calibrated single-gas concentrations, so we use them only as additional features.

Recent studies on resilient and self-healing edge systems highlight the importance of adaptive resource management and fault tolerance in long-term deployments. While the current SmellTec prototype employs deterministic execution and graceful degradation (e.g., fallback to chemistry-only operation when radar is unavailable), future work will investigate prescriptive and self-healing strategies, such as adaptive duty cycling, confidence-aware sensor suppression, and automatic recalibration, in line with emerging resilient edge-system frameworks.

The development dataset is small and was collected on a limited number of devices and sites. We will expand to multi-site, multi-device logs, use leakage-safe splits at the session and device level, and evaluate probability calibration so risk scores remain stable across conditions. Further gains in selectivity are planned. Small hybrid arrays that add one or two electrochemical channels for high-leverage families such as amines or reduced sulfur will be evaluated, along with simple cascades that allow early negative exit on clean backgrounds while reserving richer models for harder cases. Recalibration intervals and drift management will be measured over longer periods.

In summary, the current system demonstrates that a portable, presence-gated chemical sensing device can run deterministically on an MCU and provide useful triage cues. Next milestones include untethered verification, broader and more diverse datasets, targeted sensor upgrades such as electrochemical add-ons and calibrated CO_2 via NDIR, and stronger machine-learning capability through better calibration, domain adaptation, confidence estimation, and on-device model compression.

Abbreviations

ADC: Analog to Digital Converter

AP: Average Precision

AUC: Area Under the Curve

CO: Carbon Monoxide

CO_2: Carbon Dioxide

CV: Cross-Validation

eCO_2: Equivalent Carbon Dioxide Estimate

ESP32: Espressif ESP32 Microcontroller

EWMA: Exponentially Weighted Moving Average

F1: F1 Score

FMCW: Frequency Modulated Continuous Wave (radar)

GDM: Gas Distribution Mapping

GPIO: General Purpose Input/Output

H_2S: Hydrogen Sulfide

I/O: Input/output

IAQ: Indoor Air Quality

J48: C4.5 Decision Tree (J48)

LED: Light Emitting Diode

MCU: Microcontroller Unit

mmWave: Millimeter Wave

MOX: Metal Oxide (Semiconductor)

NDIR: Non-dispersive Infrared

NH_3: Ammonia

OLED: Organic Light Emitting Diode

OR: Olfactory Receptor

OSL: Odor Source Localization

PAS: Photoacoustic Spectroscopy

PC: Personal Computer

PR: Precision Recall

PSD: Power Spectral Density

PTR-MS: Proton Transfer Reaction Mass Spectrometry

Q15: Q1.15 fixed-point format (16-bit signed; 1 integer bit, 15 fractional bits)

RF: Random Forest (Algorithm)

RH: Relative Humidity

RGB: Red Green Blue (camera)

ROC: Receiver Operating Characteristic

SBC: Single-Board Computer

SIFT-MS: Selected Ion Flow Tube Mass Spectrometry

SLAM: Simultaneous Localization and Mapping

SNR: Signal-To-Noise Ratio

SVP: Saturation Vapor Pressure

TVOC: Total Volatile Organic Compounds

UART: Universal Asynchronous Receiver–Transmitter

UAV: Unmanned Aerial Vehicle

UGV: Unmanned Ground Vehicle

UI: User Interface

USAR: Urban Search and Rescue

USB: Universal Serial Bus

UX: User Experience

VOC: Volatile Organic Compound

Wi-Fi: Wireless Fidelity

ppb: Parts Per Billion

ppm: Parts Per Million

Acknowledgements

We extend our sincere gratitude to the Faculty of Mechanical and Materials Engineering and Faculty of Computer Science and Software Engineering at Huaiyin Institute of Technology for providing an enriching academic environment and invaluable resources.

Author contributions

Ervan N. Tanggono - Methodology, Software, Validation, Formal Analysis, Resources, Data Curation, Writing—Original Draft Preparation, Project Administration.

Amir A. Mokhtarzadeh - Conceptualization, Software, Validation, Resources, Writing—Review & Editing, Supervision, Project Administration.

Kaveendran Balasubramaniam - Validation, Resources, Supervision.

Jing Pin Zou - Investigation, Data Curation, Visualization, Project Administration.

Hera Naeem - Formal Analysis, Investigation, Visualization.

Parvez Mosharof - Writing—Original Draft Preparation, Writing—Review & Editing, Visualization.

Milon Selvam Dennison - Writing—Review & Editing.

All authors have read and agreed to the published version of the manuscript.

Data availability

The datasets generated and/or analyzed during the current study are publicly available on Mendeley Data and can be accessed at https://data.mendeley.com/datasets/jfsgrp5kzc/2 (DOI: 10.17632/jfsgrp5kzc.2). Requests for additional information regarding the dataset should be directed to the corresponding author, Milon Selvam Dennison at milon.selvam@kiu.ac.ug.

Declarations

Funding

This research received no external funding from any public or private agencies.

Institutional review board statement

Not applicable.

Consent to publish declaration

Not applicable.

Clinical Trial

Not applicable.

Competing interests

The authors declare no competing interests.

Consent to publish declaration

Not applicable.

Ethics and consent to participate declaration

Not applicable

Informed consent statement

Not applicable

References

1. Güntner, A. T. *et al.* Sniffing Entrapped Humans with Sensor Arrays. *Anal. Chem.* **90**, 4940–4945 (2018).

2. Sunnucks, E. J. *et al.* Performance of a Novel Electronic Nose for the Detection of Volatile Organic Compounds Relating to Starvation or Human Decomposition Post-Mass Disaster. *Sensors* **24**, 5918 (2024).
3. Trimmer, C. *et al.* Genetic variation across the human olfactory receptor repertoire alters odor perception. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 9475–9480 (2019).
4. Malnic, B., Godfrey, P. A. & Buck, L. B. The human olfactory receptor gene family. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 2584–2589 (2004).
5. Amann, A. *et al.* The human volatilome: volatile organic compounds (VOCs) in exhaled breath, skin emanations, urine, feces and saliva. *J. Breath Res.* **8**, 034001 (2014).
6. Persaud, K. & Dodd, G. Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature* **299**, 352–355 (1982).
7. Geng, J., Mokhtarzadeh, A. A. & Ding, N. Biomimetic Electronic Nose for Coffee DetectionBased on Nanosensor Arrays and MachineLearning Algorithms. in *2023 International Conference on the Cognitive Computing and Complex Data (ICCD)* 181–184 (IEEE, Huaian, China, 2023). doi:10.1109/ICCD59681.2023.10420808.
8. Robbiani, S., Lotesoriere, B. J., Dellacà, R. L. & Capelli, L. Physical Confounding Factors Affecting Gas Sensors Response: A Review on Effects and Compensation Strategies for Electronic Nose Applications. *Chemosensors* **11**, 514 (2023).
9. Shi, Y., Shi, Y., Niu, H., Liu, J. & Sun, P. Structure Optimization and Data Processing Method of Electronic Nose Bionic Chamber for Detecting Ammonia Emissions from Livestock Excrement Fermentation. *Sensors* **24**, 1628 (2024).
10. Yan, M. *et al.* Humidity compensation based on power-law response for MOS sensors to VOCs. *Sensors and Actuators B: Chemical* **334**, 129601 (2021).
11. Chen, P. *et al.* Structure Design and Performance Study of Bionic Electronic Nasal Cavity. *Biomimetics* **10**, 555 (2025).
12. Prakash, S. *et al.* Is TinyML Sustainable? *Commun. ACM* **66**, 68–77 (2023).

13. Huang, Y., Doh, I.-J. & Bae, E. Design and Validation of a Portable Machine Learning-Based Electronic Nose. *Sensors* **21**, 3923 (2021).
14. Rankin-Turner, S. & McMeniman, C. J. A headspace collection chamber for whole body volatilomics. *Analyst* **147**, 5210–5222 (2022).
15. Reeves, N. *et al.* Volatile Organic Compounds in the Early Diagnosis of Non-healing Surgical Wounds: A Systematic Review. *World j. surg.* **46**, 1669–1677 (2022).
16. Salinas Alvarez, C., Sierra-Sosa, D., Garcia-Zapirain, B., Yoder-Himes, D. & Elmaghraby, A. Detection of Volatile Compounds Emitted by Bacteria in Wounds Using Gas Sensors. *Sensors* **19**, 1523 (2019).
17. Pei, T., Liao, T., Wan, X., Wang, B. & Hao, D. Spatial Blind Source Estimation of Respiratory Rate and Heart Rate Detection Based on Frequency-Modulated Continuous Wave Radar. *Sensors* **25**, 1198 (2025).
18. Lee, H., Kim, B.-H., Park, J.-K. & Yook, J.-G. A Novel Vital-Sign Sensing Algorithm for Multiple Subjects Based on 24-GHz FMCW Doppler Radar. *Remote Sensing* **11**, 1237 (2019).
19. Sacco, G., Piuzzi, E., Pittella, E. & Pisa, S. An FMCW Radar for Localization and Vital Signs Measurement for Different Chest Orientations. *Sensors* **20**, 3489 (2020).
20. Anyfantis, A. & Blionas, S. Design and Development of a Mobile e-nose platform for Real Time Victim Localization in Confined Spaces During USaR Operations. in *2020 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)* 1–6 (IEEE, Dubrovnik, Croatia, 2020). doi:10.1109/I2MTC43012.2020.9129247.
21. Burgués, J., Esclapez, M. D., Doñate, S. & Marco, S. RHINOS: A lightweight portable electronic nose for real-time odor quantification in wastewater treatment plants. *iScience* **24**, 103371 (2021).
22. Meléndez, F. *et al.* Portable Electronic Nose Based on Digital and Analog Chemical Sensors for 2,4,6-Trichloroanisole Discrimination. *Sensors* **22**, 3453 (2022).

23. Domènech-Gil, G. *et al.* Electronic Nose for Improved Environmental Methane Monitoring. *Environ. Sci. Technol.* **58**, 352–361 (2024).
24. Chen, Q. *et al.* A Mobile E-Nose Prototype for Online Breath Analysis. *Advanced Sensor Research* **3**, 2300018 (2024).
25. Xu, P. *et al.* Temperature and Humidity Compensation for MOS Gas Sensor Based on Random Forests. in *Intelligent Computing, Networked Control, and Their Engineering Applications* (eds Yue, D. et al.) vol. 762 135–145 (Springer Singapore, Singapore, 2017).
26. Abdullah, A. N. *et al.* Correction Model for Metal Oxide Sensor Drift Caused by Ambient Temperature and Humidity. *Sensors* **22**, 3301 (2022).
27. Rajapakse, V., Karunanayake, I. & Ahmed, N. Intelligence at the Extreme Edge: A Survey on Reformable TinyML. *ACM Comput. Surv.* **55**, 1–30 (2023).
28. Lee, J., Park, K. & Kim, Y. Deep Learning-Based Device-Free Localization Scheme for Simultaneous Estimation of Indoor Location and Posture Using FMCW Radars. *Sensors* **22**, 4447 (2022).
29. Lv, W., He, W., Lin, X. & Miao, J. Non-Contact Monitoring of Human Vital Signs Using FMCW Millimeter Wave Radar in the 120 GHz Band. *Sensors* **21**, 2732 (2021).
30. Vergassola, M., Villermaux, E. & Shraiman, B. I. 'Infotaxis' as a strategy for searching without gradients. *Nature* **445**, 406–409 (2007).
31. Gongora, A. *et al.* Information-Driven Gas Distribution Mapping for Autonomous Mobile Robots. *Sensors* **23**, 5387 (2023).
32. Francis, A., Li, S., Griffiths, C. & Sienz, J. Gas source localization and mapping with mobile robots: A review. *Journal of Field Robotics* **39**, 1341–1373 (2022).
33. Li, B., Zhou, Q., Peng, S. & Liao, Y. Recent Advances of SnO2-Based Sensors for Detecting Volatile Organic Compounds. *Front. Chem.* **8**, 321 (2020).
34. Quinlan, J. R. Induction of decision trees. *Mach Learn* **1**, 81–106 (1986).
35. Breiman, L. Random Forests. *Machine Learning* **45**, 5–32 (2001).
36. Youden, W. J. Index for rating diagnostic tests. *Cancer* **3**, 32–35 (1950).

37. Bosch, S. *et al.* Electronic Nose Sensor Drift Affects Diagnostic Reliability and Accuracy of Disease-Specific Algorithms. *Sensors* **22**, 9246 (2022).
38. Betty, C. A., Choudhury, S. & Girija, K. G. Discerning specific gas sensing at room temperature by ultrathin SnO2 films using impedance approach. *Sensors and Actuators B: Chemical* **173**, 781–788 (2012).
39. Yan, J. *et al.* Electronic Nose Feature Extraction Methods: A Review. *Sensors* **15**, 27804–27831 (2015).
40. Guo, X. *et al.* A Novel Feature Extraction Approach Using Window Function Capturing and QPSO-SVM for Enhancing Electronic Nose Performance. *Sensors* **15**, 15198–15217 (2015).
41. Fawcett, T. An introduction to ROC analysis. *Pattern Recognition Letters* **27**, 861–874 (2006).
42. Niculescu-Mizil, A. & Caruana, R. Predicting good probabilities with supervised learning. in *Proceedings of the 22nd international conference on Machine learning - ICML '05* 625–632 (ACM Press, Bonn, Germany, 2005). doi:10.1145/1102351.1102430.
43. Vergara, L. & Salazar, A. On the Optimum Linear Soft Fusion of Classifiers. *Applied Sciences* **15**, 5038 (2025).
44. Raschka, S. Model Evaluation, Model Selection, and Algorithm Selection in Machine Learning. Preprint at https://doi.org/10.48550/ARXIV.1811.12808 (2018).
45. Rehman, S. U., Ali, A., Khan, A. M. & Okpala, C. Human Activity Recognition: A Comparative Study of Validation Methods and Impact of Feature Extraction in Wearable Sensors. *Algorithms* **17**, 556 (2024).