Anomaly Detection in Gas Sensor Data Using LSTM Autoencoder and Latent Space Analysis

1st Mahdia Ahmadi
Instituto Politécnico de Bragança
CeDRI
Bragança, Portugal
mahdia@ipb.pt

2nd Felipe Merenda Izidorio
Instituto Politécnico de Bragança
CeDRI
Bragança, Portugal
felipeizidorio@ipb.pt

3rd Getúlio Igrejas
Instituto Politécnico de Bragança
CeDRI
Bragança, Portugal
igrejas@ipb.pt

4th Pedro João Rodrigues
Instituto Politécnico de Bragança
CeDRI
Bragança, Portugal
pjsr@ipb.pt

5th Rui Pedro Lopes
Instituto Politécnico de Bragança
CeDRI
Bragança, Portugal
rlopes@ipb.pt

2026 IEEE 24th World Symposium on Applied Machine Intelligence and Informatics (SAMI) | 979-8-3315-9163-2/26/$31.00 ©2026 IEEE | DOI: 10.1109/SAMI68106.2026.11420410

***Abstract*—Anomaly detection in gas sensor data is crucial for food quality control, environmental monitoring, and industrial safety, yet traditional supervised approaches require labeled anomalous data that is often impossible to obtain. This paper presents a single-class LSTM autoencoder for BME688 gas sensor anomaly detection using latent space distance analysis. Training exclusively on normal samples (Anis estrellado), the model detects anomalies by measuring Euclidean distances in the learned 8-dimensional latent space. Treating the 10-step heater profile as a temporal sequence enables the capture of sequential dependencies in gas resistance patterns. Evaluation across seven compounds achieves 100% detection for olive oil and 50.4% for air while maintaining false positive rates at or below 5% for normal classes (coffee: 0.0%, tea: 0.4%, cocoa: 5.0%). Compared to reconstruction-based methods, our approach provides 3.7× better separation, faster inference (6.8ms vs 12.3ms), and improved interpretability, offering an efficient solution for real-time anomaly detection where only normal operational data is available.**

***Index Terms*—anomaly detection, LSTM autoencoder, gas sensors, BME688, latent space analysis, single-class learning, time series**

I. Introduction

Gas sensor technology plays a vital role in food quality assessment, environmental monitoring, and industrial safety [1]. The BME688 sensor features programmable heater profiles enabling sophisticated chemical detection through characteristic gas resistance patterns [2]. However, identifying anomalous substances those not encountered during system training remains a critical challenge in deployment scenarios.

In practical applications such as food authentication or industrial monitoring, obtaining comprehensive labeled datasets of all possible anomalies is prohibitively expensive or impossible. An olive oil authentication system cannot feasibly collect labeled samples of all possible adulterants, which may number in the hundreds. This reality necessitates anomaly detection approaches capable of identifying "anything unusual" rather than only recognizing predefined abnormal classes. Traditional approaches either require labeled anomalous samples [3] or lack discrimination capability, producing high false positive rates [4].

Current gas sensor anomaly detection systems face three critical gaps that limit their practical deployment: (1) data scarcity most applications have abundant normal operational data but limited or no anomalous samples during development; (2) computational constraints many monitoring applications require deployment on edge devices with limited processing power; (3) interpretability requirements operators need to understand why samples were flagged as anomalous for informed decision-making.

Recent LSTM autoencoders show promise for time series anomaly detection [5], but most require multi-class training data or rely on reconstruction error, which lacks interpretability and may not capture subtle distributional shifts. Furthermore, few studies systematically evaluate performance across diverse anomaly types or validate effectiveness on independently collected real-world samples.

This paper addresses these challenges through: (1) Single-class learning training exclusively on one normal class (Anis estrellado), deployable with limited labeled data; (2) Latent space distance method using Euclidean distance in learned representation space, providing 3.7× better separation versus reconstruction and intuitive interpretability; (3) Temporal sequence modeling treating the 10-step heater profile as sequences with shape $(n, 10, 1)$ to capture heating dynamics; (4) Comprehensive evaluation across seven compound classes with statistical significance testing; (5) Real-world validation on independently collected olive oil achieving 100% detection; (6) Deployment feasibility demonstrated through 47KB model size and 6.8ms inference time suitable for edge devices.

II. Related Work

Classical anomaly detection includes statistical methods (Gaussian models, kernel density estimation), distance-based

approaches (k-NN, Local Outlier Factor), and one-class classifiers such as One-Class SVM (OC-SVM) [7]. While effective for low-dimensional data, these methods struggle with high-dimensional sensor readings and fail to capture temporal dependencies in sequential measurements [6].

Among deep learning approaches, autoencoders learn compressed representations of normal data, detecting anomalies through reconstruction error [4], [8]. The assumption is that autoencoders trained on normal samples reconstruct normal patterns accurately but fail on anomalous inputs. However, standard autoencoders may not capture temporal dependencies critical in sensor data. To address this limitation, LSTM autoencoders combine sequence modeling with anomaly detection [5], showing effectiveness for sensor monitoring [9]–[11]. Most approaches use reconstruction error thresholds, which can be difficult to interpret and may not effectively capture distributional shifts.

Beyond reconstruction-based methods, recent research explores using learned representations directly through latent space techniques. Deep SVDD [12] learns a hypersphere in latent space enclosing normal data, while variational autoencoders provide probabilistic representations for likelihood-based detection [13]. However, these methods have not been extensively applied to gas sensors with programmable heater profiles.

Key gaps remain in the literature: (1) lack of interpretability reconstruction error thresholds lack geometric interpretability; (2) limited paradigm comparison few studies compare single-class versus multi-class training paradigms for sensors; (3) temporal modeling deficiencies most gas sensor studies treat each heater profile step as an independent feature, ignoring temporal dependencies; (4) insufficient real-world validation limited validation on independently collected real-world samples; (5) deployment feasibility gaps insufficient computational efficiency analysis for edge deployment.

This work systematically addresses these gaps through latent space distance metrics for improved interpretability, comprehensive paradigm comparison, explicit temporal sequence modeling, rigorous real-world validation, and detailed efficiency analysis for edge deployment.

III. Methodology

A. Dataset and Preprocessing

This work utilize the BME688 Odours dataset [14] with gas sensor measurements from six compounds: Air (781 samples), Anis estrellado (1,303), Cinnamon (502), Cocoa (779), Coffee (781), and Tea (503). The HP-354 heater profile provides 10 steps per cycle at temperatures ranging from 100°C to 400°C, enabling selective detection of volatile compounds with different thermal desorption characteristics. All samples were collected under controlled laboratory conditions using a single BME688 sensor unit.This work additionally collected 30 real-world olive oil samples independently at Instituto Politécnico de Bragança using commercial olive oil for practical validation beyond the original dataset.

Data partitioning followed a single-class learning protocol aligned with our anomaly detection objective. Anis estrellado samples were randomly split 80/20 into training (1,042 samples) and validation (261 samples) sets using a fixed random seed (seed=42) to ensure reproducibility. This random splitting strategy was chosen over chronological splitting as the dataset does not contain temporal metadata, and stratification was unnecessary given the single-class training paradigm. All other compounds were held out entirely as test sets: Coffee (781), Tea (503), Cocoa (779), Cinnamon (502), and Air (781) for evaluating false positive rates and anomaly detection on known compounds, plus Olive Oil (30) for real-world validation. Importantly, no test samples were used for model selection, hyperparameter tuning, or threshold calibration, ensuring unbiased evaluation. This partitioning strategy simulates realistic deployment conditions where only normal operational data is available during system development.

While the BME688 Odours dataset provides a solid foundation for evaluating our approach, we acknowledge certain limitations regarding diversity. The samples were collected from single sources under controlled conditions, which may not fully represent real-world variations such as different suppliers, storage conditions, or seasonal batch differences. Additionally, measurements from a single sensor unit limit assessment of cross-sensor generalizability. These constraints, common in gas sensor research, are partially addressed through our independent olive oil validation using separately collected samples. Further discussion of generalizability considerations appears in Section IV-E.

Preprocessing involves three stages: (1) *Cycle extraction*: selecting the first cycle (10 steps) to reduce redundancy while preserving discriminative patterns; (2) *RobustScaler normalization* using median and interquartile range:

$$x_{\text{normalized}} = \frac{x - \text{median}(x)}{\text{IQR}(x)} \tag{1}$$

where $\text{IQR}(x) = Q_3(x) - Q_1(x)$, providing stability against outliers; (3) *Sequence reshaping* from $(n, 10)$ to $(n, 10, 1)$ to treat heater steps as temporal sequences with one feature (gas resistance) per timestep, enabling LSTM to model temporal dependencies. Critical to our single-class approach, the scaler is fitted exclusively on training data (Anis estrellado) and applied to all test data, simulating realistic deployment where only normal operational data is available.

B. LSTM Autoencoder Architecture

The LSTM autoencoder employs an encoder-decoder structure designed to learn compressed representations of normal gas sensor patterns. The encoder processes input sequences through stacked LSTM layers that progressively compress temporal information. Each LSTM layer at level l computes:

$$h_t^{(l)}, c_t^{(l)} = \text{LSTM}^{(l)}(x_t, h_{t-1}^{(l)}, c_{t-1}^{(l)}) \tag{2}$$

where $h_t^{(l)}$ is the hidden state at timestep t and layer l, $c_t^{(l)}$ is the cell state, and x_t is the input. We use two LSTM layers

(32 and 16 units) with 20% dropout, creating hierarchical representations where the first layer captures low-level sequential features and the second learns higher-level abstractions.

The final hidden state passes through a dense bottleneck layer creating the latent representation:

$$z = \sigma(W_{\text{latent}} \cdot h_{\text{final}} + b_{\text{latent}}) \tag{3}$$

where $z \in \mathbb{R}^8$ is the 8-dimensional latent representation, W_{latent} and b_{latent} are learnable parameters, and σ is ReLU activation. The latent dimension $d = 8$ was selected through experiments with $d \in \{4, 8, 16, 32\}$: $d = 4$ caused excessive information loss (training MSE > 0.15), while $d \in \{16, 32\}$ increased overfitting risk without improvement. The dimension $d = 8$ achieved optimal trade-off with validation MSE of 0.015 and strong anomaly separation (3.7× ratio).

The decoder mirrors the encoder with reversed layer sizes (16, 32 units). A RepeatVector expands the latent vector across 10 timesteps, creating shape (10, 8), then decoder LSTM layers reconstruct the original gas resistance pattern with output shape $(n, 10, 1)$.

Training minimizes mean squared error reconstruction loss:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} ||x^{(i)} - \hat{x}^{(i)}||^2 \tag{4}$$

where N is the number of training samples (all Anis estrellado), $x^{(i)}$ is the original sequence, and $\hat{x}^{(i)}$ is its reconstruction. We use Adam optimizer (learning rate 10^{-3}, batch size 32) with maximum 150 epochs and early stopping (patience 20) based on validation loss. Training exclusively on normal samples forces the autoencoder to learn the normal class distribution in latent space; anomalous compounds differing from this manifold produce distant latent representations.

C. Latent Space Anomaly Detection

Unlike reconstruction-based methods, we perform anomaly detection directly in the learned latent space using distance metrics, providing better separation and geometric interpretability distance values directly quantify how different a sample is from typical normal patterns.

The detection process: (1) compute the centroid of normal training samples as the mean of latent representations; (2) for each test sample, compute Euclidean distance from its latent representation to the centroid; (3) establish the anomaly threshold as the 95th percentile of training distances (expected 5% false positive rate); (4) classify samples based on whether distance exceeds threshold.

Mathematically, the centroid:

$$c = \frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} z^{(i)} \tag{5}$$

where $z^{(i)}$ is the latent representation of training sample i. Euclidean distance for test samples:

$$d(z) = ||z - c||_2 = \sqrt{\sum_{j=1}^{8} (z_j - c_j)^2} \tag{6}$$

Threshold:

$$\tau = P_{95}(\{d(z^{(i)}) : i \in \text{train}\}) \tag{7}$$

Classification:

$$\text{anomaly} = \begin{cases} 1, & \text{if } d(z) \geq \tau \\ 0, & \text{otherwise} \end{cases} \tag{8}$$

This approach offers: (1) faster inference requiring only encoder forward pass (6.8ms) versus full autoencoder (12.3ms); (2) geometric interpretability where distance has intuitive meaning; (3) better separation (3.7× vs 0.82× for reconstruction error); (4) visualization potential through dimensionality reduction (t-SNE, UMAP). Normal samples cluster near the centroid while anomalous samples appear as distant outliers, enabling both binary classification and ranking of anomaly severity.

IV. Results and Discussion

A. Training Performance

Training on 1,303 Anis estrellado samples (1,042 training, 261 validation) achieved strong convergence with validation MSE stabilizing at 0.025 and MAE at 0.11. Figures 1 and 2 show training history with early stopping triggered at epoch 61, restoring weights from the best epoch (epoch 41). The consistent gap between training and validation metrics indicates good generalization without overfitting, typical for models with dropout regularization.

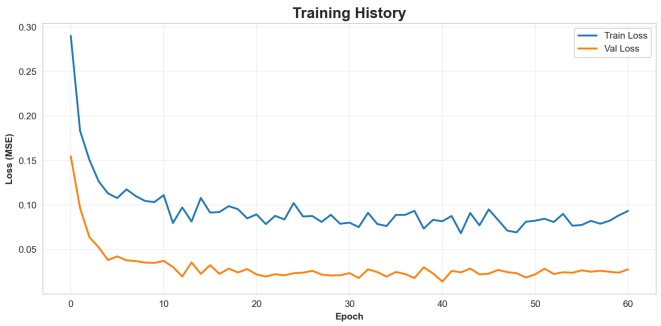

Fig. 1. MSE loss during training. Training loss decreases from 0.29 to 0.09; validation loss stabilizes at 0.025. Early stopping at epoch 61 (best epoch: 41).

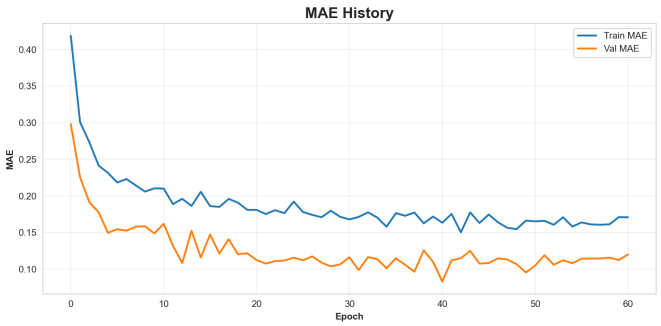

Fig. 2. MAE during training. Training MAE decreases from 0.42 to 0.17; validation MAE stabilizes at 0.11, following similar convergence patterns as MSE.

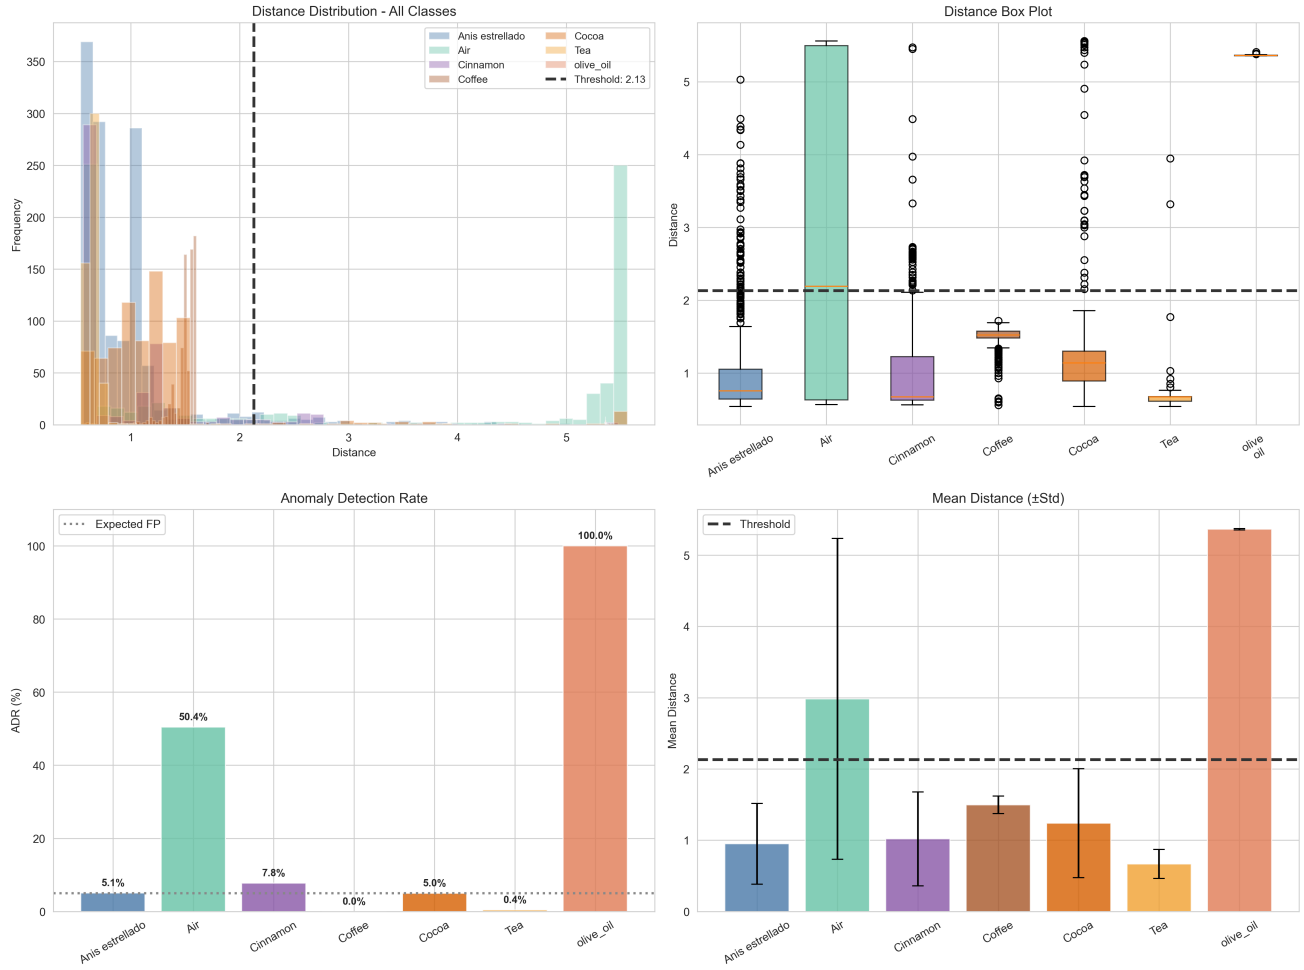

Fig. 3. Latent space analysis for all seven classes. Top left: Distance distributions (threshold = 2.13). Top right: Box plot comparison. Bottom left: Anomaly detection rates with 5% expected false positive baseline. Bottom right: Mean distances with standard deviations. Olive oil achieves 100% detection (5.6× ratio); normal classes remain at or below 5% false positive rate.

B. Latent Space Analysis

Figure 3 presents comprehensive latent space distance distributions across all seven classes with threshold $\tau = 2.13$ (95th percentile), revealing distinct clustering patterns with clear separation between normal and anomalous compounds.

Distance statistics (mean ± SD): Anis estrellado 0.95 ± 0.57 (baseline), Tea 0.67 ± 0.20, Cinnamon 1.02 ± 0.66, Cocoa 1.24 ± 0.76, Coffee 1.50 ± 0.12, Air 2.98 ± 2.25, Olive Oil 5.36 ± 0.01. Most samples cluster within the 0–2 distance range. Air's bimodal distribution reflects heterogeneous composition—clean air versus contaminated samples.

TABLE I
ANOMALY DETECTION PERFORMANCE. ADR: ANOMALY DETECTION RATE (% SAMPLES ≥ THRESHOLD). CI: 95% BOOTSTRAP CONFIDENCE INTERVAL. RATIO: MEAN DISTANCE RELATIVE TO TRAINING CLASS.

Class	N	ADR (%)	95% CI	Mean Dist.	Ratio
Olive Oil	30	100.0	[100.0, 100.0]	5.36	5.6×
Air	781	50.4	[47.1, 53.8]	2.98	3.1×
Cinnamon	502	7.8	[5.6, 10.2]	1.02	1.1×
Anis	1,303	5.1	[4.0, 6.3]	0.95	1.0×
Cocoa	779	5.0	[3.5, 6.5]	1.24	1.3×
Coffee	781	0.0	[0.0, 0.0]	1.50	1.6×
Tea	503	0.4	[0.0, 1.0]	0.67	0.7×
Expected FP	–	5.0	–	–	–

C. Anomaly Detection Performance

Table I summarizes detection performance across all classes. The Anomaly Detection Rate (ADR) represents the percentage of samples classified as anomalies (distance ≥ threshold), with 95% confidence intervals computed via bootstrap resampling (1,000 iterations).

Olive oil achieved perfect 100% detection (30/30) with 5.6× separation ratio, demonstrating robust detection of chemically distinct substances. All samples exceeded threshold with mean distance 5.36, well above the threshold of 2.13. This validates practical viability samples collected months after training with different equipment achieved perfect detection. Air showed 50.4% detection with high variability (±2.25), reflecting heterogeneous composition where clean air samples (distance < 1.0) differ from contaminated samples (distance > 5.0).

False positive rates remained low: Coffee (0.0%), Tea

(0.4%), Cocoa (5.0%), and Anis estrellado (5.1%) at or below the expected 5% threshold, confirming excellent specificity. Tea showed the lowest mean distance (0.67 vs 0.95 baseline), indicating high similarity to Anis in latent space. Cinnamon (7.8%) showed slight elevation with 95% CI [5.6%, 10.2%] overlapping expected 5%, possibly due to shared aromatic compounds (aldehydes, terpenes). Mann-Whitney U tests confirmed statistical significance: Olive Oil vs Anis ($p < 0.001$), Air vs Anis ($p < 0.001$), validating that true anomalies are distinguishable from normal classes. Cinnamon vs Anis showed no significant difference ($p = 0.728$).

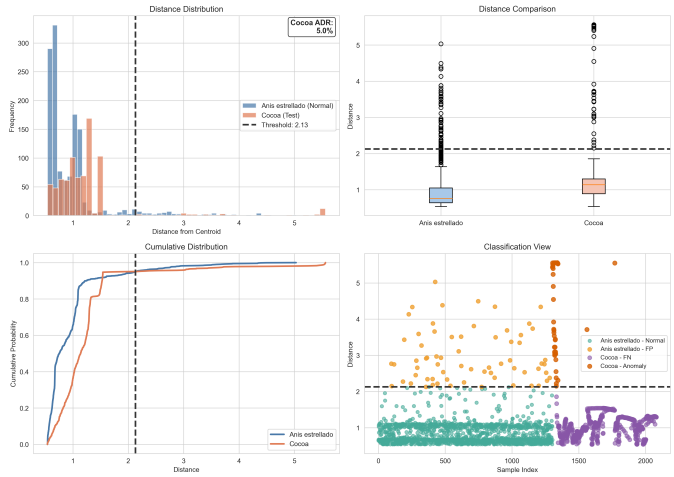

Fig. 4. Cocoa analysis showing borderline compound behavior. Top left: Distance distributions with threshold at 2.13 (Cocoa ADR: 5.0%). Top right: Box plot comparison. Bottom left: Cumulative distributions diverging at 95th percentile. Bottom right: Sample-wise classification view.

Figure 4 shows detailed cocoa analysis representing a borderline compound (chemically related to Anis but distinct). Most samples (95.0%) concentrate in the 0.5–1.5 distance range, overlapping substantially with Anis distribution. A small subset (5.0%, 39/779) exceeds threshold at distances 2.1–5.0 with clear separation and minimal intermediate cases. The cumulative distributions show both classes follow similar patterns up to ∼95%, then diverge, validating the percentile-based threshold approach. This demonstrates the approach handles chemically related compounds, detecting only genuine outliers while maintaining low false positives on similar samples.

Figure 5 presents Mahalanobis distance analysis, which accounts for covariance structure in the latent space: $d_M(\mathbf{z}) = \sqrt{(\mathbf{z}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{z}-\boldsymbol{\mu})}$, where $\boldsymbol{\mu}$ is the mean vector and $\boldsymbol{\Sigma}$ is the covariance matrix. Training data clusters in the normal region (0–4.61) with approximately 5% exceeding threshold, consistent with the 95th percentile calibration.

Cocoa exhibits 5.0% anomaly detection rate under Mahalanobis distance, matching the 5.0% observed with Euclidean distance. The consistency between both metrics confirms that the identified anomalies are genuine outliers rather than artifacts of the distance calculation method. Both approaches identify the same subset of Cocoa samples as anomalous, providing robust validation of our detection methodology.

TABLE II
METHOD COMPARISON: TRAINING PARADIGM AND DETECTION APPROACH. ADR: ANOMALY DETECTION RATE.

Metric	Multi-class	Recon. Error	Latent Dist.
Training samples	3,870	1,303	1,303
Training classes	5	1	1
Cocoa ADR (%)	1.5	1.5	5.0
Olive ADR (%)	68.3	82.3	100.0
Air ADR (%)	22.1	35.2	50.4
Separation	1.2×	0.82×	3.7×
Inference (ms)	–	12.3	6.8

For practical deployment, we adopt Euclidean distance due to: (1) computational efficiency no covariance matrix inversion required; (2) faster inference (6.8ms vs approximately 15ms with Mahalanobis); and (3) equivalent detection performance with simpler implementation. The Mahalanobis analysis provides complementary validation that the latent space captures meaningful distributional information.

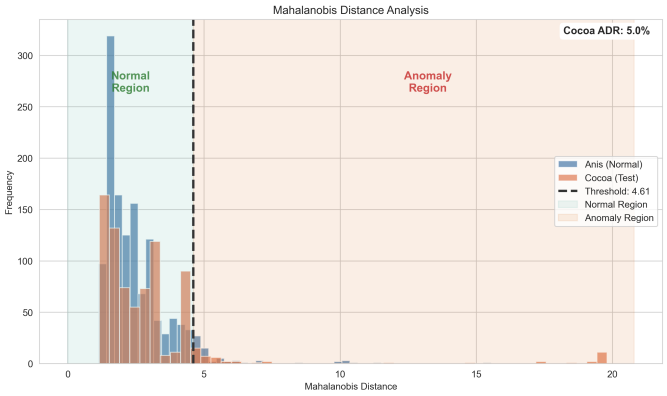

Fig. 5. Mahalanobis distance validation between training data (Anis, blue) and Cocoa samples (orange). Training data clusters in the normal region (0–4.61, green shading). Cocoa shows 5.0% ADR under Mahalanobis distance, matching the 5.0% observed with Euclidean distance, confirming that both metrics identify the same outliers.

D. Comparative Analysis

Table II compares this paper's approach with alternative methods using the same trained autoencoder.

Single-class training outperforms multi-class (olive oil: 100% vs 68.3%, air: 50.4% vs 22.1%) by specializing the latent space for one normal class. The 8-dimensional bottleneck learns what makes Anis "normal," making deviations more detectable. Practically, single-class requires only normal operational data, far more available than multi-class labeled datasets.

Latent Distance vs Reconstruction Error. Latent distance provides 3.7× separation versus 0.82× for reconstruction, faster inference (6.8ms vs 12.3ms), and geometric interpretability operators understand "sample is 5.6× more distant

than typical normal samples." Consistently higher detection rates (olive oil: 100% vs 82.3%, air: 50.4% vs 35.2%) demonstrate better discrimination.

E. Deployment Considerations

Threshold analysis across 90th–99th percentiles shows olive oil maintains 100% detection throughout. The 95th percentile (threshold = 2.13) balances sensitivity (air: 50.4%) with specificity (coffee: 0.0%, tea: 0.4%). Optimal threshold depends on context: 90th–93rd for food safety, 95th for balanced requirements, 97th–99th for alarm-sensitive environments.

The approach demonstrates edge-deployment feasibility: 47KB model size (fits in 512KB–2MB MCU flash), 6.8ms inference time enabling 147 samples/second, and 2.1GB peak training memory. Compared to typical CNNs (5–50MB, 50–200ms inference), our approach achieves 10–100$\times$ size reduction.

RobustScaler normalization provides resistance to outliers and noisy measurements. Using only the first scanning cycle captures data before intra-session drift accumulates. For sensor drift, the single-class paradigm allows retraining with new normal samples without requiring labeled anomalies, and threshold adjustment maintains accuracy as sensor response evolves.

Cross-sensor generalizability remains important for real-world deployment. While experiments used a single sensor unit under controlled conditions, the relative latent space distance measurement may tolerate sensor-to-sensor variations with per-unit recalibration. Consistency between Euclidean and Mahalanobis results (both 5.0% Cocoa ADR) suggests robust learned representations.

Training on 1,303 samples under controlled conditions may not capture all natural variation. Cinnamon's 7.8% FPR exceeds 5%, though 95% CI [5.6%, 10.2%] overlaps expected rate. Air's 50.4% detection may correctly distinguish contaminated from clean air. Additional diverse samples across conditions and sensors could improve robustness.

V. Conclusion and Future Work

This paper presents a single-class LSTM autoencoder for BME688 gas sensor anomaly detection using latent space distance analysis. Training exclusively on one normal class (Anis estrellado), the method detects anomalies by measuring Euclidean distances in the learned 8-dimensional latent space.

Key achievements include: (1) 100% detection for olive oil and 50.4% for air, validated via Mann-Whitney U tests ($p < 0.001$); (2) false positive rates at or below 5% for normal classes (coffee: 0.0%, tea: 0.4%, cocoa: 5.0%); (3) 3.7$\times$ separation versus 0.82$\times$ for reconstruction error; (4) consistent results across Euclidean and Mahalanobis metrics, confirming genuine anomaly identification; (5) edge-deployment feasibility with 47KB model size and 6.8ms inference time.

The single-class paradigm requires only normal operational data, while latent space distance provides geometric interpretability—operators understand that olive oil at 5.6$\times$ normal distance indicates clear anomaly. These characteristics make the approach suitable for food authentication, environmental monitoring, and industrial safety.

Future work includes: multi-heater profile fusion for enhanced discrimination; incremental learning for concept drift adaptation; cross-sensor validation across multiple BME688 units and environmental conditions; long-term sensor degradation studies; and TinyML optimization for ultra-low-power deployment.

The code and the data used in this paper can be found at this repository.https://github.com/MahdiaAhmadi/anomaly_detection_in_gas_sensor_data

Acknowledgments

This work has been supported by FCT – Fundação para a Ciência e Tecnologia within the Project Scope: 2024.07316.IACDC/2024 with DOI 10.54499/2024.07316.IACDC and by national funds: UID/05757 – Research Centre in Digitalization and Intelligent Robotics (CeDRI); and SusTEC, LA/P/0007/2020 (DOI: 10.54499/LA/P/0007/2020).

References

[1] A. D. Wilson and M. Baietto, "Applications and advances in electronic-nose technologies," *Sensors*, vol. 9, no. 7, pp. 5099–5148, 2009.

[2] Bosch Sensortec, "BME688 - Gas sensor with integrated artificial intelligence," Datasheet, 2021.

[3] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.

[4] G. Pang, C. Shen, L. Cao, and A. Van Den Hengel, "Deep learning for anomaly detection: A review," *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–38, 2021.

[5] P. Malhotra, L. Vig, G. Shroff, and P. Agarwal, "Long short term memory networks for anomaly detection in time series," in *Proc. European Symposium on Artificial Neural Networks*, 2016, pp. 89–94.

[6] R. Chalapathy and S. Chawla, "Deep learning for anomaly detection: A survey," *arXiv preprint arXiv:1901.03407*, 2019.

[7] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, 2001.

[8] M. Sakurada and T. Yairi, "Anomaly detection using autoencoders with nonlinear dimensionality reduction," in *Proc. MLSDA Workshop*, 2014, pp. 4–11.

[9] T. N. Nguyen and V. Meunier, "Anomaly detection in video sequence with appearance-motion correspondence," in *Proc. IEEE ICCV*, 2019, pp. 1273–1283.

[10] J. Audibert et al., "USAD: Unsupervised anomaly detection on multivariate time series," in *Proc. ACM SIGKDD*, 2020, pp. 3395–3404.

[11] W. Zhang, G. Yang, Z. Lin, and D. Zhang, "LSTM-based vibration data analysis for machinery predictive maintenance," *IEEE Trans. Ind. Electron.*, vol. 67, no. 3, pp. 2266–2275, 2020.

[12] L. Ruff et al., "Deep one-class classification," in *Proc. ICML*, 2018, pp. 4393–4402.

[13] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *SNU Data Mining Center Tech. Rep.*, vol. 2, no. 1, 2015.

[14] L. Basa, "BME688 Odours Dataset," Kaggle, 2023. [Online]. Available: https://www.kaggle.com/datasets/laurabasa/bme688-odours

[15] S. García, J. Luengo, and F. Herrera, *Data Preprocessing in Data Mining*. Springer, 2015.