

Article

Not peer-reviewed version

Confounder-Invariant Representation Learning (CIRL) for Robust Olfaction with Scarce Aroma Sensor Data: Mitigating Humidity Effects in Breath Analysis

[Md Hafizur Rahman](#)^{*}, [Jayden K Hooper](#), [Alaa Wardeh](#), [Hélène Yockell-Lelièvre](#), Jayan Ozhi Kandathil, [Ashok Prabhu Masilamani](#), [Mojtaba Khomami Abadi](#)

Posted Date: 3 November 2025

doi: 10.20944/preprints202509.0988.v2

Keywords: aroma sensors; aroma data; confounder invariant learning; representation learning; scarce data; relative humidity; deep learning; autoencoders; generalizability



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Confounder-Invariant Representation Learning (CIRL) for Robust Olfaction with Scarce Aroma Sensor Data: Mitigating Humidity Effects in Breath Analysis

Md Hafizur Rahman *, Jayden K. Hooper, Alaa Wardeh, H  l  ne Yockell-Leli  vre, Jayan Ozhi Kandathil, Ashok Prabhu Masilamani, and Mojtaba Khomami Abadi

Noze, 4920 Pl. Olivia, Saint-Laurent, QC H4R 2Z8, Canada

* Correspondence: hrahman@noze.ca

Abstract

Confounding factors in olfactory aroma data, such as high humidity levels, substantially affect sensor outputs, masking subtle volatile organic compound (VOC) patterns and hindering generalizable machine learning models. Traditional representation learning methods often require large datasets to mitigate confounder-induced variance, a resource unavailable in specialized sensor applications with limited data. This study presents Confounder-Invariant Representation Learning (CIRL), a method designed to mitigate confounding influences in data-scarce settings by leveraging explicit confounder information, such as relative humidity. CIRL enhances learned representations by reducing confounder effects, improving data purity and model robustness. Applied to three breath aroma datasets—acetone, ketosis, and peppermint-oil breath, all affected by high humidity—CIRL was integrated with standard autoencoder models. Evaluated within the same framework, CIRL improved generalization performance by 10–15% in classification accuracy across all three datasets. These results demonstrate CIRL’s potential to advance reliable artificial olfaction for applications like breath-based diagnostics in challenging real-world conditions.

Keywords: aroma sensors; aroma data; confounder invariant learning; representation learning; scarce data; relative humidity; deep learning; autoencoders; generalizability

1. Introduction

Digital olfaction systems, or electronic noses (e-noses), have emerged as a promising technology for non-invasive diagnostics by mimicking the human sense of smell to detect patterns of volatile organic compounds (VOCs)[1]. E-noses are based on an array of partially selective sensors with cross-reactivity and use pattern-recognition methods to interpret complex scent mixtures. Several e-nose architectures have been developed, differentiated primarily by their transduction mechanisms. The main types currently used include chemiresistive [2], piezoelectric [3], optical [4], electrochemical[5], bioelectronic [6] or hybrid systems. Chemiresistive e-noses are the most widely employed due to their simplicity, scalability, and compatibility with portable designs. These systems rely on materials whose electrical resistance changes upon VOC adsorption, such as conductive polymers [7], polymer–carbon black composites [8,9], metal-oxide semiconductors [10], or carbon nanostructures like CNTs [11] and graphene [12].

E-nose technology presents the potential to analyze exhaled breath for early detection of diseases such as cancer [13,14] and metabolic disorders [15,16], offering a rapid, cost-effective, and patient-friendly alternative to invasive methods. These systems operate by capturing a characteristic “scentprint” from a sample’s VOC profile that can reflect underlying physiology. However, clinical translation has been limited by the poor reliability of sensor data in real-world conditions. Chemical

sensor arrays are highly susceptible to confounding factors such as humidity, temperature gradients, ambient air dilution, and inter-individual variability which distort the VOC scentprint and can degrade classification accuracy. Among these, humidity is dominant in breath analysis because exhaled samples are naturally near saturation ($\approx 95\text{--}99\%$ RH at 37°C). Water molecules compete for adsorption sites, shift surface charge, and alter the dielectric environment of polymer and metal-oxide films, leading to baseline drift, hysteresis, and non-linear responses [17]. The combination of high humidity and chemically complex breath matrices makes it difficult to separate disease-specific VOC signals from background physiological noise.

Addressing these issues requires a multipronged strategy that prioritizes sensor-level robustness and measurement discipline: (i) sensing materials with intrinsic hydrophobicity or selective permeability; (ii) coatings and architectures that mitigate water uptake; (iii) controlled sampling with preconditioning (dew-point control and humidity-filtering materials); (iv) on-board environmental metrology (dedicated humidity and temperature channels) with baseline management; and (v) data-driven compensation suited for clinical workflows. Historically, humidity compensation has progressed from purely hardware solutions (preconditioning chambers, flow/temperature control) to hybrid sensor-algorithmic approaches that combine reference channels, signal normalization, and modest regression models for temperature–humidity correction [17–25]. For example, rapid detection systems using baseline manipulation or orthogonal signal decomposition have been proposed to mitigate humidity drift in MOS arrays, improving VOC selectivity in breath samples [18,23]. Algorithmic interference suppression, such as temperature–humidity compensation via regression models, has suppressed environmental noise in e-noses for aroma quality assessment, but these often require large calibration datasets unavailable in scarce breath enose dataset scenarios [19,24]. In breath analysis, studies on exhaled VOCs for disease detection (e.g., lung cancer) highlight humidity's role in masking biomarkers, with compensation via moisture filters or signal normalization achieving partial success but failing under variable humidity levels [20,25]. These methods can substantially reduce environmental variance, but performance can still degrade under rapidly changing humidity, across devices, or when calibration data are scarce.

Transitioning to machine learning, representation learning methods have been applied to e-nose data to learn robust features invariant to confounders, bridging sensor materials and data processing [26]. Domain adversarial networks, inspired by domain adaptation, train models to minimize confounder influence by aligning distributions across humidity domains, as seen in breath VOC classification where adversarial losses suppress humidity-induced variance [27,28]. Variational autoencoders (VAEs) enforce probabilistic disentanglement of latent factors, enabling separation of VOC signals from environmental noise in sensor time-series data [29,30]. However, these methods often assume large datasets for variance capture, limiting applicability to scarce breath aroma data where overfitting exacerbates confounder entanglement [31,32].

Disentangled representation learning has emerged as a promising extension for explicit confounder isolation in sensor applications, though challenges persist in non-linear, materials-specific interactions [26,33]. Hamaguchi et al. [34] proposed a similarity loss-based framework for disentangling factors in image data, but its reliance on pairwise constraints fails to enforce semantic separation in sensor signals with humidity confounders, as critiqued by Locatello et al. [32] for lacking inductive biases in non-linear spaces. Sanchez et al. [35] introduced mutual information maximization for paired data, yet without adversarial mechanisms, it struggles with entangled confounders like humidity that non-linearly interact with VOC responses [36,37]. Denton and Birodkar [38] developed DRNET for video disentanglement using adversarial training to separate content from pose, but its focus on temporal factors overlooks materials-related confounders such as sensor drift in e-noses [39]. Wu et al. [40] used orthogonality constraints in Vector-Decomposed Disentanglement (VDD) to split domain-invariant and domain-specific representations, effective for domain shifts but insufficient for humidity's subtle, non-linear effects on sensor materials without adversarial debiasing [41]. Cheng et al. [42] proposed Disentangled Feature Representation (DFR) for

decoupling class-specific features from variations, showing promise in few-shot tasks but assuming separable confounders, which fails in breath data where humidity tightly entangles with VOCs [27].

Recent surveys underscore these limitations in olfaction contexts, noting the need for confounder-robust methods in scarce data regimes [26,43]. For breath analysis, disentangled adversarial autoencoders have been explored for subject-invariant features in physiological signals, but few target measurable confounders like humidity in aroma sensors [44,45]. Invariant Risk Minimization (IRM) and Domain Separation Networks focus on domain shifts, while Variational Fair Autoencoders address fairness, yet none fully disentangle supervised confounders in materials-constrained e-nose data [28,31].

To address this critical gap, we propose Confounder-Invariant Representation Learning (CIRL), a disentangled autoencoder framework. CIRL is designed to adversarially separate task-relevant VOC features from humidity-related confounders, learning a purified, humidity-invariant representation of the aroma data. CIRL integrates adversarial disentanglement with explicit humidity prediction, purifying VOC latents for robust classification in scarce breath datasets, advancing beyond these approaches by leveraging sensor-specific confounder information [46]. In this paper, the approach was validated on three challenging real-world datasets involving acetone quantification and breath analysis. This demonstrates that CIRL not only successfully isolates humidity information but also significantly improves classification accuracy and robustness, particularly in the data-scarce and high-humidity conditions characteristic of clinical breath analysis. This work represents a significant step toward developing reliable and deployable digital olfaction systems for medical diagnostics.

2. Materials and Methods

2.1. E-Nose Devices

2.1.1. Chemiresistive Sensing Array Chip

The e-nose devices used in this experiment utilizes the Noze aroma sensor chip (Figure 1a), featuring an array of 32 distinct chemiresistive sensing elements. The substrate for the aroma sensor is fabricated using Electroless Nickel Immersion Gold (ENIG)-plated gold interdigitated electrodes (IDEs) patterned on copper traces over a 1 mm Rogers substrate. Each IDE finger has a width of 100 μm with an inter-electrode gap of 100 μm , and the electrode height is ~ 12 μm . The chemiresistive sensor thin films are based on proprietary polymer-carbon black (CB) nanocomposites [47], approximately 1 μm thick. The polymer selection follows the functional diversity typically adopted in e-nose designs, covering a wide range of chemical functionalities to generate distinct and complementary sorption patterns, as explained in detail elsewhere [48,49]. A readout circuit measures resistance changes at a sampling rate of 1 Hz, generating a 32-dimensional time-series that serves as the aroma's unique "scentprint". To monitor environmental conditions, an off-the-shelf BME688 sensor was integrated to provide real-time temperature and humidity levels in the aroma sample. The aroma sensor chip is enclosed in a specially designed headspace which is filled through a pump for active aroma sampling (Figure 1b).




Figure 1. (a) The aroma sensor chip used for the method validation experiments, featuring an array of 32 chemiresistive sensing thin films. (b) the headspace setup holding the aroma sensor chip, with the piezoelectric pump ensuring control of incoming airflow.

2.1.2. Vial-Based Aroma Sampler Setup

In the first set of experiments, a prototype was designed (Figure 2) in order to expose the e-nose to the headspace of a 20 mL glass vial filled with 10 mL of DI water spiked with small volumes of acetone. The vial headspace is connected to the e-nose headspace via a 5mm PTFE tube. The pump is operated continuously at a flow rate of 10 mL/min during all the phases of the experiments, while the tube is manually switched between the ambient air and the vial headspace.

The standardized, 3-phase aroma digitization protocol consists of (i) ambient sampling phase: 30 seconds of sampling air from the ambient as a reference, then (ii) aroma sampling phase: 30 seconds of sampling the vial’s headspace, followed by (iii) sensor recovery phase: 50 seconds of sampling from the ambient air again. During the sensor recovery phase, the tube is switched to ambient air in order to clean the sensors by flushing away VOCs from the chip surface; vial pressure changes are negligible due to low flow and large headspace, preventing sample distortion.




Figure 2. The sampling setup used for the vial-based acetone headspace experiment.

2.1.3. Breathalyzer Device Setup

For the two sets of experiments involving breath sampling, samples were collected from participants using the DiagNoze Breathalyzer. This setup (Figure 3) integrates the digital nose unit in a breath sampling module to detect VOCs in exhaled breath. Participants exhale into a detachable, single-patient-use mouthpiece fitted with a microbial filter, humidity filter, and backflow prevention valve for safety. The breath sampling module has a capnography valve which allows only the alveolar portion, enriched with metabolic VOCs, into an internal buffer chamber which is then directed to the digital nose unit for digitization into an aroma scentprint. The pump operates similarly for controlled flow during sampling phases. The same sampling protocol listed in section 2.3 was used for all three sets of experiments, in order to ensure consistency.




Figure 3. The DiagNoze breathalyzer setup used for the experiments conducted on human breath.

2.2. Description of the Experiments

The CIRL humidity-leveraging approach was experimentally validated in three sets of experiments, using a chemiresistive sensor array deployed in two different setups (Table 1).

Table 1. This is a table. Tables should be placed in the main text near to the first time they are cited.

Dataset	Total Samples	Classes	Key Challenge	Source Device
Acetone Headspace	385	6 levels (0-100 uL acetone)	humidity confounds acetone signal	vial-based headspace sampler
Ketogenic Breath	168	low-ketones (112); high-ketones (56)	humidity confounds acetone signals; imbalance	DiagNoze breathalyzer
Peppermint Breath	361	pre-injection (191); post-injection (170)	trace VOC detection amid high humidity; variability	DiagNoze breathalyzer

2.2.1. Acetone Headspace

In the first set of experiments, the e-nose was exposed to the headspace of a vial of DI water containing varying, precisely quantified amounts of acetone. Six different acetone solutions in 10 mL of water in 20mL glass lab vials were tested: 0 μ L (pure water), 5 μ L, 10 μ L, 20 μ L, 50 μ L, and 100 μ L, measured using a microsyringe (Hamilton TLC 25 μ L).

2.2.2. Ketogenic Breath

In the second set of experiments, the breath of a volunteer (male, age 37) was sampled as they went through four ketogenic diet cycles. Each of the four cycles included two weeks of being on the ketosis diet with a maximum carbohydrates daily intake of 20g total, and one week of normal carbohydrates-rich diet. A ketosis monitoring breathalyzer (Biosense) was used to measure the ketosis state of the volunteer.

2.2.3. Peppermint Breath

The third set of experiments was conducted on volunteers (n=19; 14 male, 5 female; age 24-45) and is based on a benchmarking protocol for breath analysis developed by Henderson et al. [50,51] Breath samples from participants were recorded before and after they ingested a capsule of essential oil of peppermint (Nature’s Way Pepogest).

2.3. *Confounder-Invariant Representation Learning (CIRL) Method*

The CIRL, a deep learning framework designed to explicitly disentangle task-relevant VOC features from confounder-related signal variations, was developed to address the challenge of humidity confounding in breath VOC analysis by e-noses.

2.3.1. Conceptual Framework

The core idea behind CIRL is to learn a latent representation of the sensor data that is split into two distinct parts: one that contains only the information needed to perform the classification task (the VOC fingerprint) and another that captures the information related to the confounder (humidity). To achieve this, the model is trained in an adversarial manner, akin to a two-player game. The main model (the encoder) tries to create a "purified" VOC representation that is completely free of any humidity information. Simultaneously, a second part of the model (the confounder predictor) acts as an adversary, trying its best to predict the humidity level from this "purified" representation. By training the encoder to fool the adversary, it learns to systematically strip out humidity-related features, resulting in a latent space that is invariant to the confounder.

2.3.2. Model Architecture

The CIRL framework is implemented as a disentangled autoencoder with three key components (Figure 4):

1. Encoder (f_{enc}): A series of temporal convolutional layers that maps the input sensor data X into two separate latent spaces: a task-relevant space z_{task} and a confounder space $z_{confounder}$.
2. Decoder (f_{dec}): A series of transposed convolutional layers that reconstructs the original input signal from both latent spaces, forcing the model to learn a complete representation.
3. Classifier and Confounder Predictor: The task classifier (c) uses only the purified z_{task} to predict the task label. The confounder predictor (h) attempts to predict the humidity signal from z_{task} .




Figure 4. (a) Overview of the proposed disentangled autoencoder (CIRL) architecture. The input data (X) is passed through the encoder f_{enc} , which generates two disentangled latent representations: z_{task} and $z_{confounder}$. These representations are used by the decoder f_{dec} for reconstruction, on the other hand z_{task} used by the task classifier and confounder predictor for their respective outputs. The overall training is governed by the loss functions $\mathcal{L}_{rec}$, $\mathcal{L}_{task}$ and $\mathcal{L}_{confounder}$, combined into $\mathcal{L}_{total}$. (b) The architecture of the baseline model. The baseline model was designed to isolate the effect of CIRL's architecture. It uses an encoder and task classifier identical to CIRL's but lacks the decoder, the split latent space, and the adversarial confounder predictor, mapping all features to a single latent space. This ensures a fair comparison where performance gains can be directly attributed to the proposed disentanglement mechanism.

2.4. Training and Optimization

The model is trained by optimizing a composite loss function that balances three objectives:

- Reconstruction Loss ($\mathcal{L}_{rec}$): Ensures the decoded signal accurately reconstructs the original input.
- Task Loss ($\mathcal{L}_{task}$): Ensures the task-relevant latent space z_{task} is predictive of the target label.
- Confounder Loss ($\mathcal{L}_{confounder}$): Used adversarially. While the confounder predictor *minimizes* this loss to find humidity information, the encoder is trained to *maximize* it, forcing the encoder to make z_{task} invariant to humidity.

The total loss is formulated as: $\mathcal{L}_{total} = \lambda_{rec}\mathcal{L}_{rec} + \lambda_{task}\mathcal{L}_{task} - \lambda_{confounder}\mathcal{L}_{confounder}$

The hyperparameters λ_{rec} , λ_{task} , and $\lambda_{confounder}$ govern the balance between reconstruction fidelity, task performance, and confounder invariance in the CIRL framework:

- The parameter λ_{rec} emphasizes reconstruction fidelity, where a higher weight (e.g., 1.0–2.0) ensures accurate reconstruction of sensor signals. However overemphasis risks retaining humidity information in z_{task} , reducing humidity-invariance representation.
- The parameter λ_{task} controls the importance of learning task-relevant attributes, and hence underweighting it can lead to poor task performance.

- The parameter $\lambda_{\text{confounder}}$ encourages learning humidity-invariant attributes alongside retaining task-relevant information, however, setting it to an excessive weight (e.g. >0.5) may disrupt task-relevant attribute encoding.

The training and optimization pseudocode as follows:

Input: Data $\mathcal{X}$, confounders $\mathcal{C}$, labels $\mathcal{Y}$, initial λ_{rec} , λ_{task} , and $\lambda_{\text{confounder}}$

Initialize f_{enc} , f_{dec} , h and c with random weights

Initialize an optimization method with a suitable learning rate for f_{enc} , f_{dec} and c

Initialize a different optimization method with an appropriate learning rate for h

For each epoch:

$(z_{\text{task}}, z_{\text{confounder}}) \leftarrow f_{\text{enc}}(\mathcal{X})$

$\hat{\mathcal{X}} \leftarrow f_{\text{dec}}(z_{\text{task}}, z_{\text{confounder}})$

$\hat{\mathcal{C}} \leftarrow h(z_{\text{task}})$, $\hat{\mathcal{Y}} \leftarrow c(z_{\text{task}})$

Compute $\mathcal{L}_{\text{rec}}$ as a chosen distance metric between $\mathcal{X}$ and $\hat{\mathcal{X}}$

Compute $\mathcal{L}_{\text{task}}$ as a selected error measure between $\mathcal{Y}$ and $\hat{\mathcal{Y}}$

Compute $\mathcal{L}_{\text{confounder}}$ as a chosen error measure between $\mathcal{C}$ and $\hat{\mathcal{C}}$

$\mathcal{L}_{\text{total}} \leftarrow \lambda_{\text{rec}}\mathcal{L}_{\text{rec}} + \lambda_{\text{task}}\mathcal{L}_{\text{task}} - \lambda_{\text{confounder}}\mathcal{L}_{\text{confounder}}$

Update f_{enc} , f_{dec} , c to minimize $\mathcal{L}_{\text{total}}$ with their optimization method

Update h to maximize $\mathcal{L}_{\text{confounder}}$ with its optimization method and gradient reversal

Optionally adjust λ_i using $\lambda_i(t) = \frac{\|\nabla \mathcal{L}_i\|}{\|\nabla \mathcal{L}_{\text{total}}\|}$

End For

2.5. Data Preprocessing

Before model training, all sensor signals underwent a standardized preprocessing pipeline:

1. **Ambient Normalization:** Each sensor's response was normalized using the formula: $x'_i = (x_i / \text{median}(x_{\text{baseline}})) - 1$. This normalization strategy preserves the relative magnitude of sensor responses while compensating for inter-sensor variability and baseline drift.
2. **Temporal Sequence Truncation:** Input sequences are truncated at the recovery phase terminus plus 60 seconds, capturing the complete VOC desorption dynamics while eliminating uninformative tail regions. This fixed-window approach ensures consistent temporal context across samples, encompassing baseline (30s), exposure (30s), and recovery phases (50s + 60s buffer), totaling approximately 170 timesteps at 1Hz sampling.
3. **Zero-Padding:** To maintain uniform tensor dimensions required for batch processing in convolutional architectures, sequences shorter than the maximum length are right-padded with zeros post-recovery phase. This post-sequence padding strategy preserves temporal causality and prevents the introduction of artificial signal artifacts during convolution operations, as the padded regions are effectively masked by the learned kernels' receptive fields.

To address the class imbalance in the ketosis dataset, we used class weighting during training, which adjusts the loss function to give more emphasis to the minority (high-ketone) class.

2.6. Experimental Setup and Evaluation

To validate CIRL, we compared its performance against a baseline model. The baseline model consisted of an identical encoder architecture and task classifier but used a single, unified latent space and lacked the decoder and adversarial components (Figure 4b). This ensures that any performance gains are attributable to the disentanglement mechanism.

All experiments were conducted using a 5-fold stratified cross-validation protocol (60% train, 20% validation, 20% test). Hyperparameters for both models (Table 2) were optimized using Bayesian optimization [52] over 50 trials, with the validation macro F1-score as the objective metric.

Table 2. Hyperparameter configuration determined through Bayesian optimization.

Parameter	Search Range	Baseline	CIRL
Learning Rate	[1e-4, 1e-2]	1×10^{-3}	3×10^{-4}
Batch Size		32	32
λ_{rec}	[0.5, 2.0]	–	1.0
λ_{task}	[0.5, 2.0]	–	1.5
λ_{conf}	[0.1, 0.5]	–	0.3

Models were implemented in TensorFlow 2.18 and trained on NVIDIA A100 GPUs.
Detailed model architectures are provided in Table 3.

Table 3. An example architectural specifications of the experiment performed on Acetone Headspace Dataset.

Component	Baseline Model	CIRL Model
Encoder	3 Conv1D layers,	3 Conv1D layers,
	Filters:[256, 128, 64]	Filters:[256, 128, 64]
	Kernel Size:3, Stride:2	Kernel Size:3, Stride:2
	BatchNorm+LeakyReLU(0.2)	BatchNorm + LeakyReLU(0.2)
Latent Space	52-dim (unified)	32-dim z_{task} + 20-dim $z_{confounder}$
Decoder		Mirror of the encoder
Task Classifier	2 FC [256, 128]	2 FC [256, 128]
	Dropout: 0.3	Dropout: 0.3
	Input: Full Latent	Input: z_{task}
Humidity Predictor		2 FC [128, 256] + 1D TransposedConv
		Input: z_{task} Output: Humidity Signal

Performance was evaluated using F1-score, precision, recall, and Area Under the Curve (AUC). To directly quantify disentanglement, we also measured the MSE of humidity signal reconstruction from the latent spaces.

3. Results

3.1. Training Dynamics and Model Convergence

Analysis of the training dynamics provides insight into CIRL's adversarial learning process. Figure 5 illustrates the evolution of the loss components and performance metrics for the acetone dataset. The total loss decomposition (Figure 5a) shows balanced optimization with all component losses converging smoothly. While the confounder loss term decreases as part of the overall optimization, the adversarial dynamics are better understood through the disentanglement metrics in Figure 5c.

The most compelling evidence of disentanglement is shown in Figure 5c. As training progresses, the ability of the adversary to reconstruct the humidity signal from the task space z_{task} deteriorates significantly (MSE increases). Conversely, its ability to reconstruct humidity from the dedicated confounder space $z_{confounder}$ improves (MSE decreases). This clear divergence confirms that humidity-related information is being systematically purged from the task-relevant features and isolated in the confounder space.




Figure 5. Training dynamics for CIRL on acetone dataset: (a) Total loss decomposition showing balanced optimization of reconstruction, task, and adversarial objectives. (b) Task classification F1-score demonstrating CIRL's superiority over baseline. (c) Adversarial humidity disentanglement - increasing MSE from z_{task} confirms successful invariance learning. (d) Reconstruction quality convergence on log scale.

3.2. Quantitative Evaluation of Disentanglement

To formally validate the disentanglement, we measured the final humidity signal reconstruction error from each latent space after training (Table 4). The results show a stark contrast: the humidity signal could not be accurately reconstructed from z_{task} ($MSE > 0.89$), confirming that humidity information has been successfully removed. In contrast, the signal could be reconstructed with very high fidelity from $z_{confounder}$ ($MSE < 0.05$), proving that this information was isolated.

Table 4. Humidity signal reconstruction MSE from the task-relevant vs. confounder-specific latent spaces.

Dataset	MSE from z_{task}	MSE from z_{conf}
Acetone Headspace	0.89 ± 0.12	0.03 ± 0.01
Ketogenic Breath	1.23 ± 0.15	0.05 ± 0.02
Peppermint Breath	1.15 ± 0.18	0.04 ± 0.01

3.3. Classification Performance

With disentanglement validated, we assessed its impact on downstream classification performance. As shown in Table 5 (Acetone) and Table 6 (Breath datasets), CIRL consistently and significantly outperformed the baseline model across all tasks.

On the 6-class acetone classification task, CIRL achieved a macro F1-score of 0.75 ± 0.03 , a 16% absolute improvement over the baseline's 0.59 ± 0.04 ; this improvement was found to be statistically significant (paired t -test across CV folds, $p < 0.01$). The gains were most pronounced for the challenging low-concentration classes, where humidity interference is most likely to obscure the weak acetone signal.

On the breath datasets, the improvements were even more dramatic. For the highly imbalanced ketosis dataset, the performance gain was even more pronounced. CIRL elevated the F1-score for high-ketosis detection from a near-random 0.42 (baseline) to a robust 0.88 (paired *t*-test, *p* < 0.001). Similarly, for the peppermint-oil task, CIRL achieved a stable F1-score of 0.74, whereas the baseline model failed to learn a meaningful pattern for the post-ingestion class (0.38 F1) (paired *t*-test, *p* < 0.01).

These results confirm that the gains from disentanglement are not just large, but statistically robust.

Table 5. Acetone concentration classification results (6-class).

Concentration	Baseline			CIRL		
	F1-Score	Precision	Recall	F1-Score	Precision	Recall
C0: 0 µL (water)	0.62 ± 0.04	0.65 ± 0.03	0.60 ± 0.05	0.86 ± 0.02	0.88 ± 0.02	0.84 ± 0.03
C1: 5 µL	0.55 ± 0.05	0.58 ± 0.04	0.52 ± 0.06	0.67 ± 0.03	0.69 ± 0.03	0.65 ± 0.04
C2: 10 µL	0.47 ± 0.06	0.50 ± 0.05	0.45 ± 0.07	0.63 ± 0.03	0.65 ± 0.03	0.61 ± 0.05
C3: 20 µL	0.68 ± 0.03	0.70 ± 0.03	0.66 ± 0.04	0.73 ± 0.02	0.75 ± 0.02	0.71 ± 0.03
C4: 50 µL	0.64 ± 0.04	0.66 ± 0.03	0.62 ± 0.05	0.80 ± 0.02	0.82 ± 0.02	0.78 ± 0.03
C5: 100 µL	0.58 ± 0.05	0.61 ± 0.04	0.55 ± 0.06	0.82 ± 0.03	0.84 ± 0.02	0.80 ± 0.03
Macro Average	0.59 ± 0.04	0.62 ± 0.03	0.57 ± 0.05	0.75 ± 0.03	0.77 ± 0.02	0.73 ± 0.03

Table 6. Breath VOC detection performance.

Dataset	Model	F1-Score	Precision	Recall	AUC
Peppermint Pre-ingestion	Baseline	0.51 ± 0.05	0.54 ± 0.04	0.48 ± 0.06	0.52 ± 0.04
	CIRL	0.74 ± 0.03	0.76 ± 0.03	0.72 ± 0.04	0.81 ± 0.02
Peppermint Post-ingestion	Baseline	0.38 ± 0.06	0.42 ± 0.05	0.35 ± 0.07	0.46 ± 0.05
	CIRL	0.74 ± 0.03	0.73 ± 0.03	0.73 ± 0.04	0.82 ± 0.02
High Ketosis	Baseline	0.42 ± 0.07	0.45 ± 0.06	0.39 ± 0.08	0.48 ± 0.06
	CIRL	0.88 ± 0.03	0.89 ± 0.02	0.87 ± 0.03	0.93 ± 0.02

3.4. Ablation Study

To isolate the contribution of each component in the CIRL framework, we conducted an ablation study (Table 7). The results show a clear progression: adding the reconstruction objective provides a solid performance boost (+9-16% F1). However, the final and most significant gains (+7-16% additional F1) are delivered by the adversarial disentanglement component. This confirms that adversarial invariance is the key innovation responsible for achieving robustness to confounders.

Table 7. Ablation study demonstrating incremental benefits of CIRL components.

Configuration	Acetone Headspace F1	Ketogenic Breath F1	Peppermint Breath F1
Baseline (single latent)	0.59 ± 0.04	0.60 ± 0.06	0.45 ± 0.05
+ Reconstruction loss	0.68 ± 0.03	0.75 ± 0.04	0.61 ± 0.04
+ Adversarial training (full CIRL)	0.75 ± 0.03	0.91 ± 0.02	0.74 ± 0.03

4. Discussion

The challenge of detecting trace volatile organic compounds (VOCs) in exhaled breath is underscored by the high-humidity, low-concentration conditions of our experiments. In the ketogenic diet study, the breath acetone levels can vary from 0.5 to more than 40 ppm [53], but the high breath humidity presents a challenge for precise monitoring. Similarly, in the peppermint breath experiment, the breath samples that are recorded at different periods of time after ingestion illustrate

the predictable washout profiles of the different peppermint VOCs [50], but the VOCs trace amounts (in the 10s of ppb) can easily be dwarfed by the high breath humidity content, around 50000 ppm. The power of CIRL lies in its explicit, supervised disentanglement mechanism. Rather than implicitly hoping a model learns to ignore confounders, CIRL's adversarial architecture forces the encoder to generate a task-relevant representation, z_{task} , that is probably free of humidity information. This is empirically validated by our key finding: the humidity signal could not be accurately reconstructed from z_{task} (MSE > 0.89), while being preserved with high fidelity in the dedicated confounder space, $z_{confounder}$ (MSE < 0.05). This answers why CIRL succeeds where other models falter: by actively purging the dominant, non-linear distortions caused by water vapor, it allows the downstream classifier to focus on the subtle, low-magnitude VOC patterns that are otherwise obscured. The performance gains—especially the elevation of the high-ketosis F1-score from a near-random 0.42 to a robust 0.88—demonstrate this principle in action. The baseline model was overwhelmed by the humidity signal, whereas CIRL successfully isolated the weak biomarker signature.

An important finding from our experiments is CIRL's robust performance across chemically distinct VOCs. The first two experiments targeted acetone, a small, polar, and hydrophilic molecule that is miscible with water. In contrast, the peppermint-oil experiment targeted menthone and menthol, which are larger, terpene-based compounds that are mostly hydrophobic. The fact that CIRL demonstrated consistent and significant performance gains in both scenarios suggests that its mechanism is largely independent of the target analyte's specific chemical properties. This implies that CIRL works by learning to isolate the signature of water vapor itself, rather than learning the specific, complex interactions between water and a particular VOC. This makes the framework highly generalizable and robust for analyzing diverse chemical mixtures.

The implications of these findings for future clinical e-nose studies are concrete and significant. By effectively neutralizing the impact of humidity, CIRL can:

- **Improve Reproducibility and Standardization:** It reduces a major source of inter-sample and inter-device variability, a critical barrier that has stalled the clinical adoption of e-nose technology.
- **Enhance Feasibility of Point-of-Care Screening:** By making the device more robust to the uncontrolled ambient humidity of a clinical setting, it lowers the need for complex and costly environmental controls, making widespread deployment more practical.
- **Increase Reliability in Longitudinal Studies:** By correcting for short-term humidity-induced drift, CIRL can improve the reliability of monitoring disease progression or treatment response over time, where distinguishing true biological change from instrumental variation is paramount. Furthermore, the CIRL framework is sensor-agnostic. While we used a nanocomposite array, the principle can be applied to any e-nose technology (e.g., MOS, CP, QCM) that is susceptible to humidity that can be measured with a dedicated humidity sensor, and/or other measurable confounding factors.

Despite these strengths, we acknowledge certain limitations that open avenues for future research. The current framework was optimized for a single, measured confounder (humidity). In real-world applications, e-nose data are often affected by multiple confounding variables simultaneously, such as ambient temperature and long-term sensor drift due to material aging. Extending CIRL to handle this is not trivial. A potential architectural adjustment could involve a multi-head confounder predictor, where z_{task} is made invariant to a vector of confounders ($C_1, C_2, \dots, C_n$) by training parallel adversarial discriminators for each. However, this would introduce practical challenges in balancing the multiple adversarial losses and would require a larger $z_{confounder}$ space to capture the joint variance. Future work should therefore focus on extending the CIRL framework to this multi-confounder scenario, perhaps by adding parallel adversarial predictors for temperature. Furthermore, applying this methodology to longitudinal datasets will be crucial for developing models that are robust not only to environmental interference but also to the inevitable effects of long-term sensor drift, thereby enhancing the stability and reliability of e-nose systems in chronic disease monitoring.

In this work, we introduced and validated Confounder-Invariant Representation Learning (CIRL), a deep learning framework designed to address the critical challenge of humidity interference in electronic nose systems. We demonstrated that by using an adversarial training mechanism, CIRL successfully learns to disentangle task-relevant VOC features from humidity-induced signal variations. Our experiments, conducted on three distinct datasets including challenging real-world breath samples classification, showed that this purification of the feature space enabled significant and consistent improvements in classification performance, with absolute F1-score gains of up to 16% over a comparable baseline model. The framework proved particularly effective in data-scarce and class-imbalanced scenarios. Ultimately, this work establishes a viable software-based path toward developing more reliable and deployable digital olfaction systems, helping to bridge the gap between laboratory potential and real-world clinical application.

Conflicts of Interest: Md Hafizur Rahman, Jayden K. Hooper, Alaa Wardeh, Ashok Prabhu Masilamani, H  l  ne Yockell-Leli  vre, Jayan Ozhi Kandathil and Mojtaba Khomami Abadi are employed by Noze.

Definition 1 (Mutual Information): The mutual information $I(A; B)$ (denoted as $I(\cdot; \cdot)$, measuring dependency between variables) between two random variables A and B measures the reduction in uncertainty

about A given knowledge of B , defined as: $I(A; B) = \mathbb{E}_{p(A,B)} \left[\log \frac{p(A,B)}{p(A)p(B)} \right]$, where $\mathbb{E}$ denotes the expectation (average) over the joint probability distribution $p(A, B)$, and $p(\cdot)$ represents the probability density function. or equivalently, $I(A; B) = H(A) - H(A|B)$, where $H(\cdot)$ is the entropy (uncertainty measure) and $H(\cdot|\cdot)$ the conditional entropy.

Definition 2 (Conditional Mutual Information): The conditional mutual information $I(A; B|C)$ (dependency given a third variable) measures the remaining mutual information between A and B given C , defined as: $I(A; B|C) = \mathbb{E}_{p(A,B,C)} \left[\log \frac{p(A,B|C)}{p(A|C)p(B|C)} \right]$, or $I(A; B|C) = H(A|C) - H(A|B, C)$, reflecting the dependence between A and B conditioned on C .

With these tools, we can now define what it means for a representation to be disentangled and identifiable in the context of CIRL:

Definition 3 (Disentangled Representation): A representation $z = (z_{\text{task}}, z_{\text{confounder}})$ is disentangled with respect to the task and confounders if

- (i) z_{task} is sufficient for predicting y , i.e., $I(z_{\text{task}}; y) \approx I(\mathcal{X}; y)$, where $I(\cdot; \cdot)$ denotes mutual information.
- (ii) z_{task} is invariant to c , i.e., $I(z_{\text{task}}; c) \approx 0$.
- (iii) $z_{\text{confounder}}$ captures the information in c , i.e., $I(z_{\text{confounder}}; c) \approx I(\mathcal{X}; c)$.

This ensures VOC features are isolated from humidity in breath data.

Definition 4 (Identifiable Representation): A disentangled representation $z = (z_{\text{task}}, z_{\text{confounder}})$ is identifiable if, given the data distribution $\mathcal{D}$, there exists a unique pair of mappings $f_{\text{enc}} : \mathcal{X} \rightarrow \mathbb{R}^{d_{\text{task}}} \times \mathbb{R}^{d_{\text{confounder}}}$ (encoder function f_{enc}) and $f_{\text{dec}} : \mathbb{R}^{d_{\text{task}}} \times \mathbb{R}^{d_{\text{confounder}}} \rightarrow \mathcal{X}$ (decoder function f_{dec}) (up to permutation and scaling) that satisfy the disentanglement conditions. In practice, this uniqueness supports reliable VOC detection despite humidity variations.

To achieve this disentanglement and identifiability, certain conditions must hold. We build on prior [54] and adapt it to aroma data scenarios:

Assumption 1 (Conditional Independence): The task labels y and confounders c are conditionally independent given the input $\mathcal{X}$, i.e., $p(y, c|\mathcal{X}) = p(y|\mathcal{X})p(c|\mathcal{X})$.

Assumption 2 (Sufficient Encoder): The encoder f_{enc} is sufficiently expressive to capture the information in $\mathcal{X}$, i.e., $I(f_{\text{enc}}(\mathcal{X}); \mathcal{X}) \approx I(\mathcal{X}; \mathcal{X})$ [55].

With these assumptions, we can prove that CIRL learns an identifiable representation:

Theorem 1 (Identifiability of CIRL Representations): Under the above Assumptions (Conditional Independence and Sufficient Encoder), and assuming the confounder predictor $h : \mathbb{R}^{d_{\text{task}}} \rightarrow \mathcal{C}$ is trained to optimality, the CIRL model learns an identifiable representation $z = (z_{\text{task}}, z_{\text{confounder}})$ such that:

- (i) z_{task} is invariant to c , i.e., $I(z_{\text{task}}; c) = 0$.
- (ii) z_{task} is sufficient for y , i.e., $I(z_{\text{task}}; y) = I(\mathcal{X}; y)$.
- (iii) $z_{\text{confounder}}$ captures c , i.e., $I(z_{\text{confounder}}; c) = I(\mathcal{X}; c)$.

Proof. The confounder loss $\mathcal{L}_{\text{confounder}} = \mathbb{E}_{(\mathcal{X}, c) \sim \mathcal{D}} [\ell_{\text{confounder}}(c, h(z_{\text{task}}))]$ is minimized adversarially by maximizing the error of h . At optimality, h cannot predict c from z_{task} , implying $I(z_{\text{task}}; c) = 0$. This follows from the data processing inequality [56], if $h(z_{\text{task}})$ contains no information about c , then z_{task} is independent of c .

The task loss $\mathcal{L}_{\text{task}} = \mathbb{E}_{(\mathcal{X}, y) \sim \mathcal{D}} [\ell_{\text{task}}(y, c(z_{\text{task}}))]$ ensures that z_{task} retains information about y . By the Sufficient Encoder assumption, the encoder captures all relevant information in $\mathcal{X}$. Since $\mathcal{L}_{\text{task}}$ optimizes c to predict y , and y is conditionally independent of c (the Conditional Independence assumption), we have $I(z_{\text{task}}; y) \geq I(\mathcal{X}; y)$. The equality holds when z_{task} is a minimal sufficient statistic for y [54].

The reconstruction loss $\mathcal{L}_{\text{rec}} = \mathbb{E}_{\mathcal{X} \sim \mathcal{D}}[\ell_{\text{rec}}(\mathcal{X}, f_{\text{dec}}(z_{\text{task}}, z_{\text{confounder}}))]$ ensures that z_{task} and $z_{\text{confounder}}$ jointly encode all information in $\mathcal{X}$. Since z_{task} is invariant to $\mathcal{C}$, the remaining information about $\mathcal{C}$ must be encoded in $z_{\text{confounder}}$. Thus, $I(z_{\text{confounder}}; \mathcal{C}) = I(\mathcal{X}; \mathcal{C})$ [54].

Identifiability follows from the uniqueness of the decomposition under the conditional independence assumption, as shown in [54]. The encoder and decoder are unique up to permutation and scaling, as the loss functions enforce distinct roles for z_{task} and $z_{\text{confounder}}$.

Appendix A.3. Optimization as an Information Trade-Off

CIRL achieves this separation through a carefully designed optimization process. The total loss combines multiple objectives: $\mathcal{L}_{\text{total}} = \lambda_{\text{rec}}\mathcal{L}_{\text{rec}} + \lambda_{\text{task}}\mathcal{L}_{\text{task}} - \lambda_{\text{confounder}}\mathcal{L}_{\text{confounder}}$, (where λ_{rec} , λ_{task} , and $\lambda_{\text{confounder}}$ are hyperparameters balancing the losses), where the negative $\mathcal{L}_{\text{confounder}}$ enforces adversarial invariance, minimizing $I(z_{\text{task}}; \mathcal{C})$. This ensures z_{task} focuses on task-relevant features, while $z_{\text{confounder}}$ absorbs confounder effects, aligning with the identifiability proof.

This optimization can be understood through an information-theoretic lens, balancing retention and invariance:

Lemma 1 (Information Trade-Off): The CIRL loss: $\mathcal{L}_{\text{total}} = \lambda_{\text{rec}}\mathcal{L}_{\text{rec}} + \lambda_{\text{task}}\mathcal{L}_{\text{task}} - \lambda_{\text{confounder}}\mathcal{L}_{\text{confounder}}$ implicitly optimizes the objectives,

- (i) Maximizing $I(z_{\text{task}}, z_{\text{confounder}}; \mathcal{X})$ via $\mathcal{L}_{\text{rec}}$.
- (ii) Maximizing $I(z_{\text{task}}; \mathcal{Y})$ via $\mathcal{L}_{\text{task}}$.
- (iii) Minimizing $I(z_{\text{task}}; \mathcal{C})$ via $\mathcal{L}_{\text{confounder}}$. This trade-off is key for scarce breath data, where retaining VOC info while discarding humidity improves generalization.

Proof. The reconstruction loss $\mathcal{L}_{\text{rec}}$ minimizes the divergence between $\mathcal{X}$ and $\hat{\mathcal{X}} = f_{\text{dec}}(z_{\text{task}}, z_{\text{confounder}})$ (where $\hat{\mathcal{X}}$ is the reconstructed input). For Gaussian noise models, this corresponds to maximizing $I(z_{\text{task}}, z_{\text{confounder}}; \mathcal{X})$ [56,57]. The task loss $\mathcal{L}_{\text{task}}$ optimizes the classifier c (task predictor $c: \mathbb{R}^{d_{\text{task}}} \rightarrow \mathcal{Y}$), which maximizes the mutual information $I(z_{\text{task}}; \mathcal{Y})$ by ensuring z_{task} is predictive of $\mathcal{Y}$ [58]. The adversarial confounder loss $\mathcal{L}_{\text{confounder}}$ minimizes $I(z_{\text{task}}; \mathcal{C})$ by training h (confounder predictor $h: \mathbb{R}^{d_{\text{task}}} \rightarrow \mathcal{C}$) to fail at predicting $\mathcal{C}$, as shown in the proof of Theorem 1 [27].

This aligns CIRL with the Information Bottleneck principle [58], where z_{task} serves as a compressed, confounder-free representation.

Appendix A.4. Generalization Bound

Finally, we assess CIRL's generalization to new aroma data, deriving a bound on task error:

Theorem 2 (Generalization Bound): Let $\mathcal{H}_{\text{task}}$ (hypothesis class for the task classifier) be the hypothesis class of the task classifier $c: \mathbb{R}^{d_{\text{task}}} \rightarrow \mathcal{Y}$, and let $\mathcal{D}_n = \{(\mathcal{X}_i, y_i, \mathcal{C}_i)\}_{i=1}^n$ (training set of size n) be a training set of size n . Under the Conditional Independence and Sufficient Encoder assumptions, the expected task error $\mathbb{E}_{(\mathcal{X}, y) \sim \mathcal{D}}[\ell_{\text{task}}(y, c(z_{\text{task}}))]$ (where ℓ_{task} is the task-specific loss) is bounded as:

$$\mathbb{E}[\ell_{\text{task}}] \leq \hat{\mathcal{L}}_{\text{task}} + \mathcal{O}\left(\sqrt{\frac{\text{VC}(\mathcal{H}_{\text{task}}) \log n + \log(1/\delta)}{n}}\right),$$
 (where $\mathcal{O}(\cdot)$ denotes big-O notation, bounding the error term asymptotically as n grows), with probability at least $1 - \delta$ (confidence parameter δ), where $\hat{\mathcal{L}}_{\text{task}}$ is the empirical task loss, and $\text{VC}(\mathcal{H}_{\text{task}})$ (Vapnik-Chervonenkis dimension measuring class complexity) is the VC-dimension of $\mathcal{H}_{\text{task}}$. This bound highlights CIRL's efficiency in small n regimes, common in breath sensor data. Proof: Since z_{task} is invariant to $\mathcal{C}$ (Theorem 1), the task classifier c operates on a confounder-free representation. The expected task error depends only on the complexity of $\mathcal{H}_{\text{task}}$ and the sample size n . Applying standard VC-dimension bounds [59], the generalization gap is:
$$\mathbb{E}[\ell_{\text{task}}] - \hat{\mathcal{L}}_{\text{task}} \leq \mathcal{O}\left(\sqrt{\frac{\text{VC}(\mathcal{H}_{\text{task}}) \log n + \log(1/\delta)}{n}}\right).$$
 The conditional independence of $\mathcal{Y}$ and $\mathcal{C}$ given $\mathcal{X}$ (Assumption 1) ensures that confounder variations do not inflate the generalization error, completing the proof.

This bound confirms that CIRL's invariance to confounders reduces the risk of overfitting to spurious correlations, enhancing robustness across diverse conditions.

References

1. Arshak, K.; Moore, E.; Lyons, G.M.; Harris, J.; Clifford, S. A Review of Gas Sensors Employed in Electronic Nose Applications. *Sens. Rev.* **2004**, *24*, 181–198.
2. Rath, R.J.; Farajikhah, S.; Oveissi, F.; Dehghani, F.; Naficy, S. Chemiresistive Sensor Arrays for Gas/volatile Organic Compounds Monitoring: A Review. *Adv. Eng. Mater.* **2022**, 2200830.
3. Kuchmenko, T.; Lvova, L. A Perspective on Recent Advances in Piezoelectric Chemical Sensors for Environmental Monitoring and Foodstuffs Analysis. *Chemosensors (Basel)* **2019**, doi:10.3390/CHEMOSENSORS7030039.
4. Askim, J.R.; Mahmoudi, M.; Suslick, K.S. Optical Sensor Arrays for Chemical Sensing: The Optoelectronic Nose. *Chem. Soc. Rev.* **2013**.
5. Gebicki, J. Application of Electrochemical Sensors and Sensor Matrixes for Measurement of Odorous Chemical Compounds. *Trends Analyt. Chem.* **2016**, *77*, 1–13.
6. Dung, T.T.; Oh, Y.; Choi, S.-J.; Kim, I.-D.; Oh, M.-K.; Kim, M. Applications and Advances in Bioelectronic Noses for Odour Sensing. *Sensors (Basel)* **2018**, *18*, 103.
7. Wang, Y.; Liu, A.; Han, Y.; Li, T. Sensors Based on Conductive Polymers and Their Composites: A Review. *Polym. Int.* **2020**, *69*, 7–17.
8. Muñoz, B.C.; Steinthal, G.; Sunshine, S. Conductive Polymer-carbon Black Composites-based Sensor Arrays for Use in an Electronic Nose. *Sens. Rev.* **1999**, *19*, 300–305.
9. Lewis, N.S. Comparisons between Mammalian and Artificial Olfaction Based on Arrays of Carbon Black-Polymer Composite Vapor Detectors. *Acc. Chem. Res.* **2004**, *37*, 663–672.
10. Shalini Devi, K.S.; Anantharamakrishnan, A.; Maheswari Krishnan, U. Expanding Horizons of Metal Oxide-based Chemical and Electrochemical Sensors. *Electroanalysis* **2021**, *33*, 1979–1996.
11. Freddi, S.; Sangaletti, L. Trends in the Development of Electronic Noses Based on Carbon Nanotubes Chemiresistors for Breathomics. *Nanomaterials (Basel)* **2022**, *12*, doi:10.3390/nano12172992.
12. Moura, P.C.; Ribeiro, P.A.; Raposo, M.; Vassilenko, V. The State of the Art on Graphene-Based Sensors for Human Health Monitoring through Breath Biomarkers. *Sensors (Basel)* **2023**, *23*, doi:10.3390/s23229271.
13. Baldini, C.; Billeci, L.; Sansone, F.; Conte, R.; Domenici, C.; Tonacci, A. Electronic Nose as a Novel Method for Diagnosing Cancer: A Systematic Review. *Biosensors* **2020**, *10*, doi:10.3390/bios10080084.
14. Yockell-Lelièvre, H.; Philip, R.; Kaushik, P.; Masilamani, A.P.; Meterissian, S.H. Breathomics: A Non-Invasive Approach for the Diagnosis of Breast Cancer. **2025**, *12*, 411.
15. Wilson, A. Advances in Electronic-Nose Technologies for the Detection of Volatile Biomarker Metabolites in the Human Breath. *Metabolites* **2015**, *5*, 140–163.
16. Wilson, A.D. Application of Electronic-Nose Technologies and VOC-Biomarkers for the Noninvasive Early Diagnosis of Gastrointestinal Diseases. *Sensors (Basel)* **2018**, *18*, 2613.
17. 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 (Basel)* **2023**, doi:10.3390/chemosensors11100514.
18. Cai, M.; Xu, S.; Zhou, X.; Lü, H. Electronic Nose Humidity Compensation System Based on Rapid Detection. *Sensors (Basel)* **2024**, *24*, doi:10.3390/s24185881.
19. Liang, Z.; Tian, F.; Yang, S.X.; Zhang, C.; Sun, H.; Liu, T. Study on Interference Suppression Algorithms for Electronic Noses: A Review. *Sensors (Basel)* **2018**, *18*, 1179.
20. Li, Y.; Wei, X.; Zhou, Y.; Wang, J.; You, R. Research Progress of Electronic Nose Technology in Exhaled Breath Disease Analysis. *Microsyst. Nanoeng.* **2023**, *9*, 129.
21. Bax, C.; Robbiani, S.; Zannin, E.; Capelli, L.; Ratti, C.; Bonetti, S.; Novelli, L.; Raimondi, F.; Di Marco, F.; Dellacà, R.L. An Experimental Apparatus for E-Nose Breath Analysis in Respiratory Failure Patients. *Diagnostics (Basel)* **2022**, *12*, 776.
22. Dhanush Gowda, A.M.; Dessai, A.D.; Nayak, U.Y. Electronic-Nose Technology for Lung Cancer Detection: A Non-Invasive Diagnostic Revolution. *Lung* **2025**, *203*, 76.

23. Sanislav, T.; Mois, G.D.; Zeadally, S.; Folea, S.; Radoni, T.C.; Al-Suhaimi, E.A. A Comprehensive Review on Sensor-Based Electronic Nose for Food Quality and Safety. *Sensors (Basel)* **2025**, *25*, 4437.
24. Zhai, Z.; Liu, Y.; Li, C.; Wang, D.; Wu, H. Electronic Noses: From Gas-Sensitive Components and Practical Applications to Data Processing. *Sensors (Basel)* **2024**, *24*, 4806.
25. Shuba, A.; Kuchmenko, T.; Menzhulina, D. Drift Compensation of the Electronic Nose in the Development of Instruments for out-of-Laboratory Analysis. *Chemistry Proceedings* **2022**, *5*, 68.
26. Cheng, D.; Xie, Y.; Xu, Z.; Li, J.; Liu, L.; Liu, J.; Zhang, Y.; Feng, Z. Disentangled Latent Representation Learning for Tackling the Confounding M-Bias Problem in Causal Inference. *arXiv [cs.LG]* 2023.
27. Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; Lempitsky, V. Domain-Adversarial Training of Neural Networks. *J. Mach. Learn. Res.* **2015**, *17*, 189–209.
28. Bousmalis, K.; Trigeorgis, G.; Silberman, N.; Krishnan, D.; Erhan, D. Domain Separation Networks. *Neural Inf Process Syst* **2016**, *29*, 343–351.
29. Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.P.; Glorot, X.; Botvinick, M.; Mohamed, S.; Lerchner, A. Beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. *Int Conf Learn Represent* **2016**.
30. Louizos, C.; Swersky, K.; Li, Y.; Welling, M.; Zemel, R. The Variational Fair Autoencoder. *arXiv [stat.ML]* 2015.
31. Arjovsky, M.; Bottou, L.; Gulrajani, I.; Lopez-Paz, D. Invariant Risk Minimization. *arXiv [stat.ML]* 2019.
32. Locatello, F.; Bauer, S.; Lucic, M.; Gelly, S.; Scholkopf, B.; Bachem, O. Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations. *ICML* **2018**, *97*, 4114–4124.
33. Wang, X.; Chen, H.; Tang, S.; Wu, Z.; Zhu, W. Disentangled Representation Learning. *arXiv [cs.LG]* 2022.
34. Hamaguchi, R.; Sakurada, K.; Nakamura, R. Rare Event Detection Using Disentangled Representation Learning. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR); IEEE, June 2019; pp. 9327–9335.
35. Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; Battaglia, P. Learning to Simulate Complex Physics with Graph Networks. *ICML* **2020**, *abs/2002.09405*, 8459–8468.
36. Hjelm, R.D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; Bengio, Y. Learning Deep Representations by Mutual Information Estimation and Maximization. *arXiv [stat.ML]* 2018.
37. Zhang, B.H.; Lemoine, B.; Mitchell, M. Mitigating Unwanted Biases with Adversarial Learning. In Proceedings of the Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society; ACM: New York, NY, USA, December 27 2018; pp. 335–340.
38. Denton, E.L.; Birodkar, V. Unsupervised Learning of Disentangled Representations from Video. *Neural Inf Process Syst* **2017**, *30*, 4414–4423.
39. Villegas, R.; Yang, J.; Hong, S.; Lin, X.; Lee, H. Decomposing Motion and Content for Natural Video Sequence Prediction. *arXiv [cs.CV]* 2017.
40. Wu, A.; Liu, R.; Han, Y.; Zhu, L.; Yang, Y. Vector-Decomposed Disentanglement for Domain-Invariant Object Detection. In Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV); IEEE, October 2021; pp. 9342–9351.
41. Do, K.; Tran, T. Theory and Evaluation Metrics for Learning Disentangled Representations. *arXiv [cs.LG]* 2019.
42. Cheng, H.; Wang, Y.; Li, H.; Kot, A.C.; Wen, B. Disentangled Feature Representation for Few-Shot Image Classification. *arXiv [cs.CV]* 2021.
43. Ziyatdinov, A.; Marco, S.; Chaudry, A.; Persaud, K.; Caminal, P.; Perera, A. Drift Compensation of Gas Sensor Array Data by Common Principal Component Analysis. *Sens. Actuators B Chem.* **2010**, *146*, 460–465.
44. Han, M.; Ozdenizci, Ö.; Wang, Y.; Koike-Akino, T.; Erdoğan, D. Disentangled Adversarial Autoencoder for Subject-Invariant Physiological Feature Extraction. *IEEE Signal Process. Lett.* **2020**, *27*, 1565–1569.
45. Feng, D.; Li, C.; Dai, W.; Liang, P.P. SMELLNET: A Large-Scale Dataset for Real-World Smell Recognition. *arXiv [cs.AI]* 2025.
46. Ganin, Y.; Lempitsky, V. Unsupervised Domain Adaptation by Backpropagation. In Proceedings of the Proceedings of the 32nd International Conference on Machine Learning; Bach, F., Blei, D., Eds.; PMLR: Lille, France, 07–09 Jul 2015; Vol. 37, pp. 1180–1189.

47. Levaray, N.; Ozhikandathil, J.; Masilamani, A.P.; Panarello, T. U.S. Patent No. 11,788,985. *Patent* **2023**.
48. Ryan, M.A.; Zhou, H.; Buehler, M.G.; Manatt, K.S.; Mowrey, V.S.; Jackson, S.P.; Kisor, A.K.; Shevade, A.V.; Homer, M.L. Monitoring Space Shuttle Air Quality Using the Jet Propulsion Laboratory Electronic Nose. *IEEE Sens. J.* **2004**, *4*, 337–347.
49. Shevade, A.V.; Ryan, M.A.; Homer, M.L.; Manfreda, A.M.; Zhou, H.; Manatt, K.S. Molecular Modeling of Polymer Composite-Analyte Interactions in Electronic Nose Sensors. *Sens. Actuators B Chem.* **2003**, *93*, 84–91.
50. Henderson, B.; Ruszkiewicz, D.M.; Wilkinson, M.; Beauchamp, J.D.; Cristescu, S.M.; Fowler, S.J.; Salman, D.; Francesco, F.D.; Koppen, G.; Langejürgen, J.; et al. A Benchmarking Protocol for Breath Analysis: The Peppermint Experiment. *J. Breath Res.* **2020**, *14*, 046008.
51. Henderson, B.; Slingers, G.; Pedrotti, M.; Pugliese, G.; Malásková, M.; Bryant, L.; Lomonaco, T.; Ghimenti, S.; Moreno, S.; Cordell, R.; et al. The Peppermint Breath Test Benchmark for PTR-MS and SIFT-MS. *J. Breath Res.* **2021**, *15*, 046005.
52. Frazier, P.I. A Tutorial on Bayesian Optimization. *arXiv [stat.ML]* 2018.
53. Anderson, J.C. Measuring Breath Acetone for Monitoring Fat Loss: Review. *Obesity* **2015**, *23*, 2327–2334.
54. Khemakhem, I.; Kingma, D.P.; Hyvärinen, A. Variational Autoencoders and Nonlinear ICA: A Unifying Framework. *AISTATS* **2019**, *108*, 2207–2217.
55. Achille, A.; Soatto, S. Emergence of Invariance and Disentanglement in Deep Representations. *Journal of Machine Learning Research* **2018**, *19*, 1–34.
56. Thomas M. Cover; Thomas, J.A. *Elements of Information Theory: Cover/elements of Information Theory, Second Edition*; 2nd ed.; John Wiley & Sons: Nashville, TN, 2006; ISBN 9780471241959.
57. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. *arXiv [stat.ML]* 2013.
58. Tishby, N.; Pereira, F.C.; Bialek, W. The Information Bottleneck Method. *arXiv [physics.data-an]* 2000.
59. Vapnik, V. An Overview of Statistical Learning Theory. *IEEE transactions on neural networks* **1999**, *10*, 988–999.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.