# SMELLNET: A LARGE-SCALE DATASET FOR REAL-WORLD SMELL RECOGNITION

#### **Anonymous authors**

000

002 003 004

006

008

010

011

012

013

014

015

016

017

018

019

021

023

025

026

027 028

029

031

033

034

037

040 041

042 043

044

046

047

048

050 051

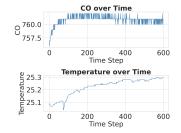
052

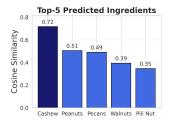
Paper under double-blind review

## **ABSTRACT**

The ability of AI to sense and identify various substances based on their smell alone can have profound impacts on allergen detection (e.g., smelling gluten or peanuts in a cake), monitoring the manufacturing process, and sensing hormones that indicate emotional states, stress levels, and diseases. Despite these broad impacts, there are virtually no large-scale benchmarks, and therefore little progress, for training and evaluating AI systems' ability to smell in the real world. In this paper, we use small gas and chemical sensors to create SMELLNET, the first large-scale database that digitizes a diverse range of smells in the natural world. SMELLNET contains about 828,000 data points across 50 substances, spanning nuts, spices, herbs, fruits, and vegetables, and 43 mixtures among them, with 68 hours of data collected. Using SMELLNET, we developed SCENTFORMER, a Transformer-based architecture combining temporal differencing and sliding-window augmentation for smell data. For the SMELLNET-BASE classification task, SCENTFORMER achieves 58.5% Top-1 accuracy, and for the SMELLNET-MIXTURE distribution prediction task, SCENTFORMER achieves 50.2% Top-1@0.1 on the test-seen split. SCENTFORMER's ability to generalize across conditions and capture transient chemical dynamics demonstrates the promise of temporal modeling in olfactory AI. SMELLNET and SCENTFORMER lay the groundwork for real-world olfactory applications across healthcare, food and beverage, environmental monitoring, manufacturing, and entertainment.







- (a) Sensor setup detecting cashew.
- (b) Time-series signals from CO and temperature sensors.
- (c) Top-5 model predictions using cosine similarity.

Figure 1: Overview of our smell sensing data collection and modeling pipeline. (a) Sensor hardware setup and data capture. (b) Raw sensor readings over time. (c) AI model predictions on the substance.

# 1 Introduction

Advancements in AI have revolutionized how machines perceive and interact with the world. However, most progress has been limited to the text, vision, and audio modalities (Liang et al., 2024; Gan et al., 2022). The human sense of smell is crucial for environmental perception, social interaction, and regulating well-being. Similarly, AI that can recognize smell can revolutionize the entertainment, e-commerce, manufacturing, and food and beverage industries (Deshmukh et al., 2015; Vilela et al., 2019). More ambitiously, smell-based diagnostics can help in early disease detection, (e.g., COVID-19) (Kwiatkowski et al., 2022), and even 'smelling' hormones and indicators of emotional states, stress levels, and for early prognosis of cancer (Tillotson, 2017; Zamkah et al., 2020).

Nevertheless, smell is a new data modality for AI, with little progress compared to computer vision and natural language processing. We believe that large-scale data and real-time AI models are key for learning rich feature representations of smell for accurate sensing, classification, and feature fusion between smell and other human senses. This strategy differs from past research, which has

Table 1: Benchmark datasets for machine learning in olfaction.

Prior work emphasized molecular stimuli and human judgments. SMELLNET adds a large sensor-based set for base odors; our dataset extends this to both base and mixture odors at scale, aiding generalization to everyday odors.

Dataset	Stimulus Type	Data Type	Method	Size
Dravnieks Atlas (Dravnieks, 1985)	Mono-molecular	Semantic descriptor ratings	Human evaluation	21,300 data points
DREAM Challenge (Keller et al., 2017)	Mono-molecular	Semantic descriptor ratings	Human evaluation	9,996 data points
Snitz et al. (2013) (Snitz et al., 2013)	Mixtures of molecules	Perceptual similarity ratings	Human evaluation	191 data points
Olfactory Metamers (Ravia et al., 2020)	Mixtures of molecules	Perceptual similarity ratings	Human evaluation	49,788 data points
SMELLNET (ours)	Natural stimuli	Gas-sensor time series	Metal Oxide gas sensors	828,000 data points

emphasized feature engineering, small datasets, and simple classification models (Achebouche et al., 2022; Guerrini et al., 2017; Lee et al., 2012), and processing pre-recorded smell data collected via large chemistry lab equipment (Tran et al., 2019; Lee et al., 2023), which do not work in real-time. As a step towards real-world and real-time smell sensing, we present SMELLNET, a large-scale database of how food, beverages, and natural objects smell in the natural world. SMELLNET is collected by applying small sensors to 50 substances (nuts, spices, herbs, fruits, and vegetables) and 43 simulated mixtures across 68 hours of data. With 828,000 time step readings across environmental conditions, it is the largest and most diverse open-source smell dataset to date, enabling the training of AI that can classify substances based on their smell alone.

Using SMELLNET, we developed SCENTFORMER, a Transformer-based architecture combining temporal differencing and sliding-window augmentation for smell data. For the SMELLNET-BASE classification task, SCENTFORMER achieves 58.5% Top-1 accuracy, and for the SMELLNET-MIXTURE distribution prediction task, SCENTFORMER achieves 50.2% Top-1@0.1 on the test-seen split. SCENTFORMER's ability to generalize across conditions and capture transient chemical dynamics demonstrates the promise of temporal modeling in olfactory AI. SMELLNET, SCENTFORMER, and the source code are provided in the anonymous supplementary to ensure reproducibility and to facilitate new applications in healthcare, food and beverage, environmental monitoring, manufacturing, and entertainment.

# 2 THEORETICAL BASIS AND RELATED WORK

While large-scale AI for smell is completely unexplored, we are inspired by human smell sensing, the chemistry and biology of smell, and using AI to process small-scale smell data.

**Human smell receptors:** Olfaction, the sense of smell, allows for the detection and discrimination of odors in the environment (Stevenson, 2010). The human nose can detect and discriminate between an estimated trillion different odors, even in minute quantities (Bushdid et al., 2014). This makes the human olfactory system the largest, in terms of the number and diversity of receptors, allowing for the sensing of a vast chemical landscape (Sharma et al., 2019). Human olfactory perception functions through a combinatorial code, where each odorant molecule binds to a specific set of olfactory receptors (ORs) in the nose (Malnic et al., 1999). This binding converts chemical information into electrical signals which are perceived in the brain (Firestein, 2001).

Smell sensors for perceiving smell have been developed, including chemical compositions of gases based on the principles of molecular interaction and chemical potential equilibrium (Brattoli et al., 2011). These sensors can employ different scientific strategies to detect and analyze gas molecules, including those based on semiconducting materials like metal oxides (Nikolic et al., 2020), electrochemical sensors that generate a current proportional to the gas concentration (Bakker and Telting-Diaz, 2002), optical sensors based on different gases absorbing different wavelengths (Hodgkinson and Tatam, 2012), and conductive polymers that change their conductivity when exposed to gas molecules (Miasik et al., 1986). We use gas sensors due to their portability, although all sensors

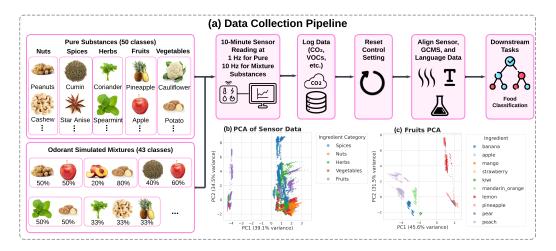


Figure 2: (a) Data collection pipeline. Each ingredient undergoes six 10-minute sensing sessions across different days, using a controlled environment to minimize external noise. During each session, 12-channel gas sensor data is recorded at 1 Hz and labeled with the ingredient identity, collection time, and associated metadata. We further pair each ingredient with high-resolution GC-MS data to enable multimodal learning. This setup enables the creation of a structured and temporally rich dataset for representation learning of smells. (b-c) PCA projections of sensor responses. While broad category separation is evident, clusters remain partially overlapping, underscoring the challenge of distinguishing ingredients and motivating more advanced models.

can suffer challenges due to sensitivity, environmental interference, reproducibility, and device calibration (Sung et al., 2024; Yan et al., 2015).

AI for smell sensing: There has been some work in using technology to process pre-recorded smell data, but these systems are not portable or work in real-time. These include graph neural networks trained to classify smell chemical molecules (especially pre-recorded GC-MS data) (Sanchez-Lengeling et al., 2019; Tran et al., 2019; Lee et al., 2023), but they require data collection via large chemistry lab equipment rather than portable sensors. Past research has also emphasized human domain knowledge, feature engineering, small datasets, and simple classification models rather than large-scale data-driven learning (Achebouche et al., 2022; Guerrini et al., 2017; Lee et al., 2012). Electronic noses have been designed to monitor pollutants and for air quality assessment (Attallah and Morsi, 2022; Payette et al., 2023), but they are not generally applicable for any type of smell. Methods to classify biological olfactory data of mice and humans have also been proposed (Fang et al., 2024; Wang et al., 2021), but do not enable portable real-time sensing.

**Other datasets:** Tab. 1 normalizes prior olfaction datasets by the number of data points. Unlike prior work focused on mono-molecular or pairwise human judgments, SMELLNET provides large-scale *sensor* time series from natural stimuli. To our knowledge, it is the first large-scale, open sensor-based dataset for smell, spanning 50 base substances and 43 mixtures over 68 hours.

# 3 CREATING SMELLNET

# 3.1 A SMALL AND REAL-TIME SMELL SENSOR

We use a series of metal oxide sensors to detect concentrations of various gases. Specifically, we used MQ-3, MQ-5, MQ-9, WSP2110, MP503, and the Grove Multichannel V2 to detect carbon monoxide (CO), nitrogen dioxide (NO), alcohol (CHOH), volatile organic compounds (VOCs), liquefied petroleum gas (LPG), among others. We also used BME680 for pressure, temperature, and humidity atmospheric control. Together, these compounds are present in common odors found in food, drinks, and other common substances. The circuit diagram of our sensor and all the hardware used to create the sensing devices is included in App. A.1.

# 3.2 Smell sensing data collection

SMELLNET comprises two main components: base substances for classification task and mixture substances for distribution prediction task.

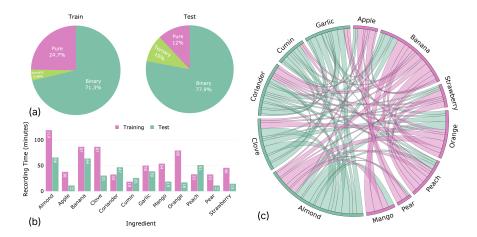


Figure 3: **Composition of SMELLNET-MIXTURE.** (a) Distribution of mixtures across train and test datasets. The test set contains harder samples, including binary (77.9%) and ternary (10%) mixtures. (b) Minutes of valid odor data collected for each ingredient. (c) Ingredient co-occurrence patterns: chords indicate data mixtures, and widths represent the amount of mixture data. Each chord encompasses mixtures of different ratios. Teal chords indicates where the mixture is prominently spices or nuts, while pink indicates mixtures that are prominently fruits on average. The diverse odor mixtures are evenly spread across the entire substance space.

SMELLNET-BASE is collected on 50 substances across 5 classes: nuts, spices, herbs, fruits, and vegetables. The full taxonomy of substances is shown in Fig. 6b. For each, we used the sensor to collect data for a 10-minute session, repeated 6 times across different days. Each session was done in a controlled environment to minimize environmental factors. At the end of each session, we clear out the air in the controlled environment so that the environment returns to atmospheric conditions (see App. A.2 and App. H.3 for analysis using different time frames and detailed procedures for resetting each experiment). Fig. 2 illustrates the overall data collection pipeline. This gives us 1 hour of sensor readings for each substance and 50 hours of data in total. During each session, 12-channel gas sensor data is recorded at 1 Hz, which gives us 180,000 total datapoints. Data is labeled with the name of the substance, its detailed description, and the time and date of collection.

In real world scenarios, mixtures of smells are more common, so we also collect SMELLNET-MIXTURE with mixtures of 12 base odorants. The odorants are selected from various types of fragrance oils, essential oils, and flavor extracts to form our smell palette: almond, apple, banana, clove, coriander, cumin, garlic, mango, orange, peach, pear, and strawberry. Details are included in App. Tab. 14. These odorants span a range of functional groups and chemical families relevant to a broad range of odors, including phenols (clove), aldehydes (almond, cumin), esters (banana, pear, peach, strawberry, apple), terpenes (orange, mango, coriander), and sulfur compounds (garlic). Furthermore, we define two test sets with different generalization challenges: (1) Test-seen contains mixture ratios that appear in the training set but from different recording sessions, testing the model's ability to generalize across temporal and environmental variations. (2) Test-unseen contains novel mixture combinations never encountered during training, challenging the model's compositional generalization in entirely new ingredient pairings, or novel ratios of familiar mixtures.

Due to size restriction of the collection environment, the mixture data is collected with only the Grove Multichannel V2 sensor, with four channels spanning across NO<sub>2</sub>, C<sub>2</sub>H<sub>5</sub>OH, VOC, and CO at 10 Hz. The resulting dataset comprises 18.0 hours of continuous sensor recordings with 648,000 data points across 1,078 distinct measurement sessions, with 679 training sessions (11.32 hours), 215 test-seen odors (3.58 hours), and 184 test-unseen mixtures (3.07 hours). Combining the two subsets together, we have a total of 68 hours of sensor readings for 50 base substances and 43 different mixtures, totaling 828,000 data points, making it the largest sensor based multitask smell dataset up to date.

# 3.3 Pairing with GC-MS

One potential limitation of sensor data is its low resolution, which stems from the quality of the sensors used. We therefore paired this data with preexisting open-source pre-recorded GC-MS data (Kitson et al., 1996). GC-MS devices are large, bulky, and non-portable, but can detect the exact

concentration of various compounds in a substance at a high resolution. As a result, pairing gas sensing data with GC-MS readings provides complementary information for this task and allows for studies of multimodal fusion and cross-modal learning (Liang et al., 2024).

#### 3.4 SMELLNET STATISTICS

We standardize the recorded data into a common format via the concentrations of volatile gases, humidity, barometric pressure, temperature, and other atmospheric conditions over the time period. We plot some examples of sensor readings in Fig. 2(b-c), and App. C. As shown in Fig. 2(b), we apply PCA to the 12-dimensional sensor readings across all time steps to visualize ingredient-level separability. The projection reveals visible clustering according to ingredient categories, particularly for nuts and fruits, which occupy distinct regions in the PCA space. This suggests that sensor responses capture category-specific variance, potentially driven by differences in volatile compound profiles. However, we see that spices and herbs still largely overlap, which makes them hard to distinguish. Future works on better sensor designs and algorithmic advances can potentially provide stronger signals that improves models' performance in these two categories. To further investigate within-category separability, we performed a separate PCA only on fruits. In Fig. 2(c), individual fruits form well-separated clusters, e.g., banana, kiwi, lemon, and pear each exhibit distinct spatial groupings at a much better separation than the global categories. As a result, we expect the model to have a better classification results on fruits categories as compared with spices and herbs.

Per-substance reading statistics are included in App. Tab. 15-16, and a kernel density estimation (KDE) graph of the readings in each category are included in App. Fig. 13-22. Each sensor has a distinct mean and standard deviation, indicating that preprocessing (Sec. 4.2), is necessary for stable predictions. From the KDE distributions, we see that gas resistance sensor is a powerful discriminator for fruits, vegetables and nuts. The  $C_2H_5OH$  sensor, on the other hand, gives a much higher reading for spices, allowing the model to discern spices from other substances. As environmental factors can also affect sensor readings, we include environmental sensors, including temperature, pressure, humidity and altitude sensors, so that the model can discover and disentangle the effect of these factors. For example, the feature correlation matrix in App. Fig. 9 shows that pressure has a correlation of 0.37 with alcohol sensors. By including these environmental sensors, the model can better isolate the effects on readings from the substance being measured.

Fig.3 describes the data distribution of SMELLNET-MIXTURE. In particular, with a binary mixture percentage of 77.9% and a ternary mixture percentage of 10%, the test subset provides a challenging environment for the model to predict the mixture ratios. As shown in Fig.3(c), the mixtures span evenly across the entire substance space, with mixtures both within and across categories.

#### 4 DEVELOPING SCENTFORMER FOR SMELL

Developing AI for smell poses challenges: sensor data is temporal, limited in quantity, and noisy. To address these, we design SCENTFORMER with a Transformer backbone, data-efficient training, and preprocessing to handle noise.

## 4.1 PROBLEM SETUP AND NOTATION

Each example is a multichannel time series  $x=(x_1,\ldots,x_T)\in\mathbb{R}^{T\times d}$ . Our encoder  $f_\theta$  produces an embedding  $\mathbf{h}=f_\theta(x)\in\mathbb{R}^H$ . For classification on SMELLNET-BASE, we map x to a label  $y\in\{1,\ldots,50\}$ . When GC-MS metadata is available, we align  $\mathbf{h}$  with a GC-MS embedding during training (App. E.1). For mixture distribution approximation on SMELLNET-MIXTURE, we map x to a probability vector  $\boldsymbol{\pi}\in[0,1]^{12}$  with  $\sum_{i=1}^{12}\pi_i=1$ . The target  $\boldsymbol{\pi}^*$  encodes the ground-truth mixture fractions (see App. E.2 for label construction).

# 4.2 Sensor Data Preprocessing

Given the temporal nature of the data, we apply the following preprocessing procedures:

**Channel Dropping.** Upon analyzing our SMELLNET in App. D, we decided to keep only 6 channels (NO $_2$ , C $_2$ H $_5$ OH, VOC, CO, Alcohol, LPG). For the other channels, we noticed that the sensor outputs irregular values for a small portion of the data, which could mean some sensors were malfunctioning. The full 12 channel data are still released, and future works could explore utilizing all channels after filtering out abnormal values, and learning additional signals from these partially available channels. Full channel analysis is available in App. H.1.

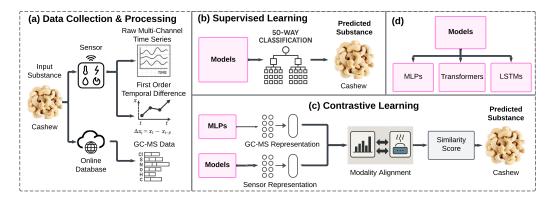


Figure 4: **Overview of the models used in this study.** (a) Raw multi-channel time series data is collected using a portable smell sensor as it samples an input substance (e.g., cashew). The data is optionally transformed using first-order temporal differences to emphasize signal dynamics. In parallel, high-resolution GC-MS data can be retrieved from an online database to provide chemical supervision. (b) In supervised learning, sensor data (raw or preprocessed) is passed through classification models trained to predict the correct substance among 50 classes. (c) In contrastive learning, paired sensor and GC-MS representations are aligned through modality-specific encoders. The resulting similarity scores rank the substance for prediction. (d) Our framework supports multiple model types—MLPs, Transformers, and LSTMs—each capable of ingesting either raw or temporally differenced sensor inputs to perform classification or representation learning.

**Temporal Difference.** Since the sensor outputs are qualitative rather than absolute quantitative measures, relative variations are often more informative than raw values. To capture relative sensor changes, we apply a *temporal difference*. For each sensor channel  $x_t$ , we compute the difference over a fixed lag of p samples:

$$\Delta x_t = x_t - x_{t-p}, \quad \forall t > p.$$

**Sliding Windows.** Due to the limited number of available recordings, we partition each file into smaller windows with size w with a stride of w/2 to increase the effective dataset size. This strategy is popular in time-series domains to improve model generalization (Norwawi et al., 2021).

**Standardization.** All training data are aggregated to compute statistics for standardization, which are then applied to both training and evaluation sets. This ensures comparability across samples and reduces the influence of sensor noise.

#### 4.3 Scentformer Architecture

We employ a pre-norm Transformer encoder over windowed sensor sequences, projected to a latent dimension D and augmented with optional positional encodings and a learnable <code>[CLS]</code> token. Sequences are processed by stacked Transformer layers, with variable lengths supported via key–padding masks. The encoder output is pooled (mean or <code>[CLS]</code>) into a fixed-size vector. We attach a classification head for 50-way odor recognition, and two auxiliary heads for mixture presence and proportion prediction (see App. F.1 for implementation details).

#### 4.4 Training Objectives

We study three objectives: (i) supervised classification from gas sensors, (ii) cross-modal alignment with GC-MS, and (iii) mixture ratio prediction.

**Supervised classification.** We train SCENTFORMER directly on SMELLNET-BASE windows for 50-way classification, using a softmax output and minimizing cross-entropy loss.

**Cross-modal alignment.** To leverage GC-MS information (FooDB Contributors, 2024; Liang et al., 2024), we adopt a symmetric contrastive learning objective (Radford et al., 2021; Socher et al., 2013) that aligns sensor and GC-MS embeddings (see App. F.2 for the full formula).

**Mixture prediction.** For SMELLNET-MIXTURE windows, SCENTFORMER predicts normalized 12-D ratios. We optimize a composite loss combining KL divergence, hinge- $\ell_1$  penalty, and focal BCE (see App. F.3 for details).

Table 2: Single-ingredient odor classification on SMELLNET-BASE (RQ1) and improvements from GC-MS integration (RQ2). We vary preprocessing choices by window size ( $w \in 50, 100$ ) and gradient differencing ( $g \in 0, 25$ ). Gradient features (g = 25) yield large gains over raw signals, longer windows (w = 100) improve stability, and temporal models consistently outperform non-temporal baselines. Adding GC-MS supervision via contrastive learning further boosts weaker models and provides consistent, though smaller, gains for stronger temporal architectures.  $\Delta$  Acc@1 reports the change in accuracy relative to the sensor-only baseline.

Model	Window	Lag	Senso	or-only (RQ	1)	Cross	-modal (RQ	(2)	$\Delta$ Acc@1
1,10001			Acc@1↑	Acc@5↑	<b>F</b> 1↑	Acc@1↑	Acc@5↑	<b>F</b> 1↑	(X-S)
MLP	50	0	21.9	55.8	17.2	23.6	57.1	19.4	+1.7
MLP	50	25	18.2	49.0	17.8	23.8	60.2	23.2	+5.6
MLP	100	0	21.0	54.4	17.4	25.6	56.4	21.0	+4.6
MLP	100	25	26.8	59.7	24.7	28.0	58.7	26.0	+1.2
CNN	50	0	25.5	61.2	23.3	28.4	69.3	26.1	+2.9
CNN	50	25	46.9	81.7	46.1	45.9	84.0	44.6	-1.0
CNN	100	0	29.5	66.6	24.9	31.0	69.3	28.4	+1.5
CNN	100	25	52.7	85.6	50.5	57.1	87.8	55.9	+4.4
LSTM	50	0	29.3	72.2	25.9	29.7	64.8	26.4	+0.4
LSTM	50	25	50.6	84.7	48.8	53.3	81.5	51.3	+2.7
LSTM	100	0	28.8	58.0	27.2	33.7	64.9	30.7	+4.9
LSTM	100	25	57.9	87.0	56.0	56.1	82.8	54.7	-1.8
SCENTF	ORMER (our	s)							
Transf.	50	0	35.1	70.9	33.1	36.2	72.2	33.3	+1.1
Transf.	50	25	50.6	85.0	49.5	50.9	81.5	50.4	+0.3
Transf.	100	0	39.9	74.7	35.7	41.0	72.1	37.9	+1.1
Transf.	100	25	56.1	87.4	55.5	58.5	84.8	58.3	+2.4

#### 5 EXPERIMENTS

We aim to evaluate SCENTFORMER using the SMELLNET dataset for both the classification and distribution tasks. Specifically, we seek to answer the following research questions:

- **RQ1:** To what extent can we classify single substance odors from SMELLNET-BASE alone, and which preprocessing choices contribute most to accuracy?
- **RQ2:** Does cross-modal training with paired GC-MS sensor data improve downstream sensor-only classification, and by how much?
- **RQ3:** Can we predict the composition of mixtures from sensor time-series, and which preprocessing methods work best?

# 5.1 EVALUATION METRICS

For SMELLNET-BASE, we evaluate classification performance using Top-1 and Top-5 accuracy, F1 score, and per-category accuracy. Top-1 and Top-5 measure the proportion of correct top predictions. F1 averages precision and recall, providing a balanced view that mitigates class imbalance.

For SMELLNET-MIXTURE, we evaluate prediction quality using MAE, Top-1@0.1 accuracy, meaning predicted ratio falls within  $\pm 0.1$  of the ground-truth values on non-zero targets. We also report a dynamic Top-k hit rate, where k corresponds to the number of non-zero components in the target distribution prediction accuracy (see App. I.2 for details).

#### 5.2 EXPERIMENTAL SETUP

We evaluate SCENTFORMER on base-substance recognition and distribution prediction.

**SMELLNET-BASE** For each ingredient, we randomly select one of the six days as the held-out test day, and the remaining five days are used for training. We compare SCENTFORMER against non-temporal baseline (MLP) and temporal baselines (CNN and LSTM).

To study temporal dynamics, we vary temporal differencing (lag  $p \in \{0, 25\}$ ) and segment sensor streams into windows of size  $w \in \{50, 100\}$ . All models are trained for 90 epochs with batch size 32, using learning rates  $\{3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}\}$ . We select the best checkpoint per model configuration based on validation Top-1 accuracy. We further investigate cross-modal training variants that incorporate paired GC-MS supervision via contrastive learning.

**SMELLNET-MIXTURE** Mixture experiments use 12 odorants with both *seen* (novel sessions, known ratios) and *unseen* (zero-shot transfer) test splits. Models are trained with window sizes  $w \in \{50, 100\}$ , batch size 64, 60 epochs, same learning rate grid.

Fig. 1 shows the real-world deployment setup. Fig. 4 illustrates the evaluation pipeline, and full implementation details, additional hyperparameters, and reasons behind each choice are provided in App. G.2.

Table 3: **Distribution prediction on seen combinations (RQ3).** SCENTFORMER achieves the best overall performance, with consistently higher Top-1 and Top-*K* accuracy, demonstrating the importance of temporal modeling for resolving overlapping odor signals.

Model	Window	$\mathrm{MAE}\downarrow$	Top-1@0.1↑	Top- <i>K</i> (%)↑
MLP	50	0.0428	44.0	85.0
	100	0.0586	33.7	78.9
CNN	50	0.0404	48.1	86.7
	100	0.0476	36.2	87.0
LSTM	50	0.0399	46.4	89.3
	100	0.0430	46.5	86.3
SCENTFORMER	50	0.0395	50.2	87.9
	100	0.0417	47.9	89.0

# 5.3 RQ1: Preprocessing Choices Evaluation

Based on the results in Tab. 2, we highlight three key findings:

Finding 1.1: Temporal differencing substantially improves accuracy. Adding temporal differencing (lag p = 25) consistently outperforms raw signals (lag p = 0), with an average gain of 16.1% across models and window sizes. This demonstrates that temporal changes in sensor values carry critical discriminative information.

Finding 1.2: Larger windows provide more stable patterns. Window size w = 100 generally yields higher accuracy than w = 50, as longer temporal context captures more stable dynamics, though at the cost of fewer training and test samples. See App. I.1 for window calculation.

**Finding 1.3: Temporal models outperform non-temporal baselines.** CNN, LSTMs, and SCENT-FORMER achieve higher accuracy than MLPs.

# 5.4 RQ2: GC-MS INTEGRATION

Tab. 2 shows the effect of adding GC-MS supervision via contrastive learning.

Finding 2.1: GC-MS supervision strongly boosts weaker models. Raw-signal inputs (p = 0) and non-temporal architectures see the largest gains, showing that GC-MS embeddings provide complementary structure that compensates for limited model capacity.

**Finding 2.2: Gains are smaller but consistent for stronger temporal models.** Architectures like SCENTFORMER already capture much of the discriminative signal, but GC-MS further refines embeddings by grounding them in molecular structure, suggesting the two signals complement rather than replace each other.

Finding 2.3: Effects depend on architecture and preprocessing. In some cases (e.g., CNN at w=50, p=25), alignment brings little or even negative improvement, indicating that GC-MS can conflict with strong short-range features. This underscores that its value depends on how well the base model and preprocessing prepare features for cross-modal alignment.

#### 5.5 RQ3: DISTRIBUTION PREDICTION

Finding 3.1: SCENTFORMER outperforms other architectures on mixture prediction. As shown in Tab. 3, SCENTFORMER achieves the best results across both Top-1@0.1 and Top-K accuracy. This indicates that strong temporal modeling, which benefits single-substance recognition, is equally important for resolving overlapping signals in mixtures.

Finding 3.2: Accuracy drops sharply for unseen mixtures. Tab. 4 shows that performance degrades when evaluating on mixtures not seen

Table 4: **Distribution prediction on unseen combinations (RQ3).** SCENTFORMER achieves the best overall performance, but accuracy drops substantially compared to the seen setting, highlighting the challenge of generalizing to novel odor mixtures.

Model	Top-1@0.1↑	Top- $K$ (%)↑
MLP	11.7	34.0
CNN	12.4	36.4
LSTM	11.8	34.2
SCENTFORMER	16.0	38.9

during training. This suggests limited generalization: the model transfers poorly to novel ratios, even

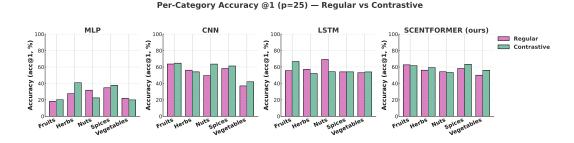


Figure 5: **Per-category accuracy** (acc@1) for four models at lag p = 25. Bars are paired per category (Regular vs. Contrastive). Figure shows the non temporal models suffer from categories like vegetables, but temporal architectures demonstrate stronger robustness across categories.

when all individual components were seen. Accuracy degradation indicates sensitivity to composition shift rather than merely class imbalance or session effects.

Finding 3.3: Top-K performance remains well above chance in unseen setting. Although unseen mixtures are harder, SCENTFORMER still achieves substantially higher Top-K accuracy than random guessing (around 16.7%; see App. I.2). This suggests that the learned representations encode meaningful compositional structure, enabling the model to narrow predictions to a plausible subset of substances even without explicit training on those mixtures.

These results show that SCENTFORMER is effective at mixture prediction, but generalization to unseen mixtures remains challenging. While temporal modeling provides clear benefits, scaling to the combinatorial complexity of real-world odors may require compositional training strategies, data augmentation, or domain adaptation.

#### 5.6 DISCUSSION

To better understand the limitations of our models, we examine the per-category classification accuracy reported in Fig. 5. Several systematic patterns emerge.

First, non-temporal baseline shows pronounced weaknesses on certain categories, particularly vegetables. Vegetables exhibit significant overlapping with other categories in PCA of sensor profiles Fig. 2 making them harder to separate. MLP reaches relatively high accuracy on spices, whose sensor signatures are more distinct. This phenomenon is consistent with the high within-class variance and overlapping volatile compound distributions.

In contrast, temporal architectures demonstrate stronger robustness across categories. By modeling temporal dynamics, these models capture transient variations in gas concentration that help differentiating harder classes. For example, SCENTFORMER maintains relatively balanced performance across all five categories, whereas MLP shows large category-specific disparities.

We also note that contrastive training with GC-MS supervision provides differential benefits. For weaker models, alignment with molecular structure substantially improves recognition across categories, suggesting that the external chemical signal compensates for limited representational capacity. However, for stronger temporal models, the gains are smaller and sometimes negligible, indicating that temporal modeling already captures much of the discriminative signal. This suggests future research on how higher resolution data affects performance.

More ablation studies and analysis for channel activities are in App. H.2.

# 6 CONCLUSION

In this work, we introduced SMELLNET, the first large-scale dataset for real-world smell recognition. Built using portable and low-cost gas sensors, it captures over 828,000 data points across 50 base substances and 43 mixtures, paired with high-resolution GC-MS chemical data. SMELLNET establishes a benchmark for studying olfactory AI at scale, enabling both single-substance recognition and mixture prediction tasks. SMELLNET also inspires the design of SCENTFORMER, a Transformer-based architecture combining temporal differencing and sliding-window augmentation for smell data. We believe SMELLNET will serve as a foundation for future research in AI for smell and various real-world applications.

# 7 ETHICS STATEMENT

This work presents SMELLNET, a large-scale dataset for smell recognition using portable gas sensors. The dataset consists entirely of time-series sensor readings from chemical compounds in food substances and natural objects. No human subjects were involved in data collection, and the dataset contains no personally identifiable information. All data was collected using commercially available sensors measuring volatile organic compounds from common food items purchased from public retailers.

Our collection system and classification models have minimal environmental impacts. Our sensor system is energy efficient, and can be powered by USB cable with 5W input. The models are lightweight; on a single NVIDIA L40S (driver 550.54.14, CUDA 12.4) at batch size 32, SCENTFORMER achieves mean per-window latency of 0.0191-0.0479 s, as shown in App. H.4.

While SMELLNET is designed to advance research in olfactory AI with beneficial applications in food safety, healthcare, and environmental monitoring, we recognize that smell sensing technology could potentially be misused. We encourage responsible use of this dataset and the resulting models, particularly regarding privacy considerations in real-world deployments.

# 8 REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our work, we provide comprehensive materials and documentation. The complete SMELLNET dataset, containing 828,000 time-series data points across 50 base substances and 43 mixtures, is included in the supplementary materials with detailed instructions describing collection protocols and sensor specifications. Our sensor hardware configuration is fully documented in Appendix A.1, including circuit diagrams and component specifications. All preprocessing steps are described in Sec. 4.2, with implementation details in our released codebase. The hyperparameters are listed in App. F.1 App. G Sec:5.2. For GC-MS integration experiments, we provide the molecular descriptor construction process (App. C) and links to the public database used. The mixture label construction methodology is detailed in App. E.2, with the complete list of odorants and their sources in Tab. 14. The dataset follows the hierarchical structure shown in Fig. 12, with CSV files organized by ingredient and recording session.

# REFERENCES

- Rayane Achebouche, Anne Tromelin, Karine Audouze, and Olivier Taboureau. Application of artificial intelligence to decode the relationships between smell, olfactory receptors and small molecules. *Scientific* reports, 12(1):18817, 2022.
- Omneya Attallah and Iman Morsi. An electronic nose for identifying multiple combustible/harmful gases and their concentration levels via artificial intelligence. *Measurement*, 199:111458, 2022.
- Eric Bakker and Martin Telting-Diaz. Electrochemical sensors. Analytical chemistry, 74(12):2781–2800, 2002.
- Bosch Sensortec. Bme680: Low power gas, pressure, temperature & humidity sensor [datasheet]. https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bme680-ds001.pdf, 2023.
- Magda Brattoli, Gianluigi De Gennaro, Valentina De Pinto, Annamaria Demarinis Loiotile, Sara Lovascio, and Michele Penza. Odour detection methods: Olfactometry and chemical sensors. *Sensors*, 11:5290–5322, 2011.
- Caroline Bushdid, Marcelo O Magnasco, Leslie B Vosshall, and Andreas Keller. Humans can discriminate more than 1 trillion olfactory stimuli. *Science*, 343(6177):1370–1372, 2014.
- Shankari Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta: Large-scale self-supervised pretraining for molecular property prediction. *arXiv* preprint arXiv:2010.09885, 2020. URL https://arxiv.org/abs/2010.09885.
- Sharvari Deshmukh, Rajib Bandyopadhyay, Nabarun Bhattacharyya, RA Pandey, and Arun Jana. Application of electronic nose for industrial odors and gaseous emissions measurement and monitoring–an overview. *Talanta*, 144:329–340, 2015.
- Andrew Dravnieks. Atlas of odor character profiles. ASTM International, Philadelphia, PA, 1985.
- Lan Fang, Cuizhu Mao, Haiting Wang, Qian Ding, Wenyao Jiao, Bingshuo Li, Yibo Zhang, and Dunwei Gong. Recent progress of organic artificial synapses in biomimetic sensory neural system. *Journal of Materials Chemistry C*, 2024.

- Stuart Firestein. How the olfactory system makes sense of scents. *Nature*, 413(6852):211–218, 2001.
- FooDB Contributors. Foodb: The food database, 2024. URL https://foodb.ca. Accessed: 2025-05-11.
  - Zhe Gan, Linjie Li, Chunyuan Li, Lijuan Wang, Zicheng Liu, and Jianfeng Gao. Vision-language pre-training: Basics, recent advances, and future trends. *Foundations and Trends® in Computer Graphics and Vision*, 14 (3–4):163–352, 2022.
  - Luca Guerrini, Eduardo Garcia-Rico, Nicolas Pazos-Perez, and Ramon A Alvarez-Puebla. Smelling, seeing, tasting old senses for new sensing. *ACS nano*, 11(6):5217–5222, 2017.
  - Hanwei Electronics Co., Ltd. Technical data mq-3 gas sensor [datasheet]. https://cdn.sparkfun.com/assets/6/a/1/7/b/MQ-3.pdf, n.d.a.
  - Hanwei Electronics Co., Ltd. Technical data mq-5 gas sensor [datasheet]. https://files.seeedstudio.com/wiki/Grove-Gas\_Sensor-MQ5/res/MQ-5.pdf, n.d.b.
  - Hanwei Electronics Co., Ltd. Mq-9 semiconductor sensor for co/combustible gas [datasheet]. https://www.haoyuelectronics.com/Attachment/MQ-9/MQ9.pdf, n.d.c.
  - Jane Hodgkinson and Ralph P Tatam. Optical gas sensing: a review. *Measurement science and technology*, 24 (1):012004, 2012.
  - Sabrina Jaeger, Simone Fulle, and Samo Turk. Mol2vec: Unsupervised machine learning approach with chemical intuition. *Journal of Chemical Information and Modeling*, 58(1):27–35, 2018. doi: 10.1021/acs.jcim.7b00616.
  - Andreas Keller, Richard C. Vosshall, Leslie B. Vosshall, et al. Predicting human olfactory perception from chemical features of odor molecules. *Science*, 355(6327):820–826, 2017. doi: 10.1126/science.aal2014.
  - Fulton G Kitson, Barbara S Larsen, and Charles N McEwen. *Gas chromatography and mass spectrometry: a practical guide.* Academic Press, 1996.
  - Andrzej Kwiatkowski, Sebastian Borys, Katarzyna Sikorska, Katarzyna Drozdowska, and Janusz M Smulko. Clinical studies of detecting covid-19 from exhaled breath with electronic nose. *Scientific reports*, 12(1): 15990, 2022.
  - Brian K Lee, Emily J Mayhew, Benjamin Sanchez-Lengeling, Jennifer N Wei, Wesley W Qian, Kelsie A Little, Matthew Andres, Britney B Nguyen, Theresa Moloy, Jacob Yasonik, et al. A principal odor map unifies diverse tasks in olfactory perception. *Science*, 381(6661):999–1006, 2023.
  - Sang Hun Lee, Oh Seok Kwon, Hyun Seok Song, Seon Joo Park, Jong Hwan Sung, Jyongsik Jang, and Tai Hyun Park. Mimicking the human smell sensing mechanism with an artificial nose platform. *Biomaterials*, 33(6): 1722–1729, 2012.
  - Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations & trends in multimodal machine learning: Principles, challenges, and open questions. *ACM Computing Surveys*, 56(10):1–42, 2024.
  - Bettina Malnic, Junzo Hirono, Takaaki Sato, and Linda B Buck. Combinatorial receptor codes for odors. *Cell*, 96(5):713–723, 1999.
  - Jan J Miasik, Alan Hooper, and Bruce C Tofield. Conducting polymer gas sensors. *Journal of the Chemical Society, Faraday Transactions 1: Physical Chemistry in Condensed Phases*, 82(4):1117–1126, 1986.
  - Maria Vesna Nikolic, Vladimir Milovanovic, Zorka Z Vasiljevic, and Zoran Stamenkovic. Semiconductor gas sensors: Materials, technology, design, and application. *Sensors*, 20(22):6694, 2020.
  - N. M. Norwawi et al. Sliding window time series forecasting with multilayer ... *PMC (NCBI)*, 2021. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC8988917/. "Sliding window time series forecasting" demonstration on COVID-19 data etc.
  - Julie Payette, Fabrice Vaussenat, and Sylvain Cloutier. Deep learning framework for sensor array precision and accuracy enhancement. *Scientific Reports*, 13(1):11237, 2023.
  - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, and Sandhini Agarwal. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021.
  - Aharon Ravia, Kobi Snitz, Danielle Honigstein, Maya Finkel, Rotem Zirler, Ofer Perl, Lavi Secundo, Christophe Laudamiel, David Harel, and Noam Sobel. A measure of smell enables the creation of olfactory metamers. *Nature*, 588(7836):118–123, December 2020. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-020-2891-7.

597

598

600

601

602 603

604

605

607

608

609 610

611

612 613

614

615

616

617

618

619

620

621

622

623 624

625

626

627

628

629

630

631

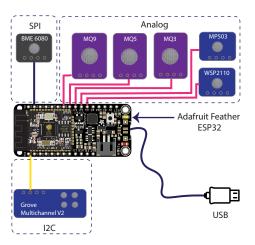
632

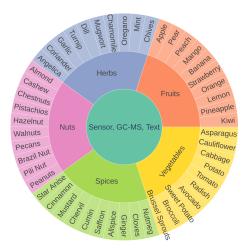
633

- Benjamin Sanchez-Lengeling, Jennifer N Wei, Brian K Lee, Richard C Gerkin, Alán Aspuru-Guzik, and 595 Alexander B Wiltschko. Machine learning for scent: Learning generalizable perceptual representations of 596 small molecules. arXiv preprint arXiv:1910.10685, 2019.
  - Seeed Studio. Grove multichannel gas sensor v2 [datasheet]. https://www.farnell.com/ datasheets/3759010.pdf, n.d.
  - Anju Sharma, Rajnish Kumar, Imlimaong Aier, Rahul Semwal, Pankaj Tyagi, and Pritish Varadwaj. Sense of smell: structural, functional, mechanistic advancements and challenges in human olfactory research. Current neuropharmacology, 17(9):891-911, 2019.
  - Kobi Snitz, Adi Yablonka, Tali Weiss, Idan Frumin, Rehan M. Khan, and Noam Sobel. Predicting Odor Perceptual Similarity from Odor Structure. PLoS Computational Biology, 9(9):e1003184, September 2013. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1003184.
  - Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning with neural tensor networks for knowledge base completion. In Advances in neural information processing systems, pages 926–934, 2013.
  - Richard J Stevenson. An initial evaluation of the functions of human olfaction. Chemical senses, 35(1):3–20,
  - Seung-Hyun Sung, Jun Min Suh, Yun Ji Hwang, Ho Won Jang, Jeon Gue Park, and Seong Chan Jun. Data-centric artificial olfactory system based on the eigengraph. Nature communications, 15(1):1211, 2024.
  - Jenny Tillotson. Emotionally responsive wearable technology and stress detection for affective disorders. Psychiatria Danubina, 29(suppl. 3):604-606, 2017.
  - Ngoc Tran, Daniel Kepple, Sergey Shuvaev, and Alexei Koulakov. Deepnose: Using artificial neural networks to represent the space of odorants. In International Conference on Machine Learning, pages 6305-6314. PMLR, 2019.
  - Alice Vilela, Eunice Bacelar, Teresa Pinto, Rosário Anjos, Elisete Correia, Berta Gonçalves, and Fernanda Cosme. Beverage and food fragrance biotechnology, novel applications, sensory and sensor techniques: An overview. Foods, 8(12):643, 2019.
  - Peter Y Wang, Yi Sun, Richard Axel, LF Abbott, and Guangyu Robert Yang. Evolving the olfactory system with machine learning. Neuron, 109(23):3879-3892, 2021.
  - Jia Yan, Xiuzhen Guo, Shukai Duan, Pengfei Jia, Lidan Wang, Chao Peng, and Songlin Zhang. Electronic nose feature extraction methods: A review. Sensors, 15(11):27804–27831, 2015.
  - Abdulaziz Zamkah, Terence Hui, Simon Andrews, Nilanjan Dey, Fuqian Shi, and R Simon Sherratt. Identification of suitable biomarkers for stress and emotion detection for future personal affective wearable sensors. Biosensors, 10(4):40, 2020.
  - Zhengzhou Winsen Electronics Technology Co., Ltd. Mp503 air-quality gas sensor [datasheet]. https://files.seeedstudio.com/wiki/Grove\_Air\_Quality\_Sensor\_v1.3/res/ Mp503%20English.pdf, n.d.a.
  - Zhengzhou Winsen Electronics Technology Co., Ltd. Wsp2110 voc gas sensor [datasheet]. https://files. seeedstudio.com/wiki/Grove-HCHO\_Sensor/res/Wsp2110-1-.pdf, n.d.b.

# A SENSORS AND DATA COLLECTION ENVIRONMENT

# A.1 SENSOR HARDWARE SPECIFICATIONS





- (a) Circuit diagram of sensor hardware setup.
- (b) Sunburst taxonomy of ingredient categories.

Figure 6: **Overview of the SMELLNET dataset and sensing setup.** (a) Our constructed portable smell sensor detects concentrations of various gases and atmospheric factors through 7 multi-channel gas sensors. (b) SMELLNET includes smell sensor readings of 50 substances spanning nuts, spices, herbs, fruits, and vegetables.

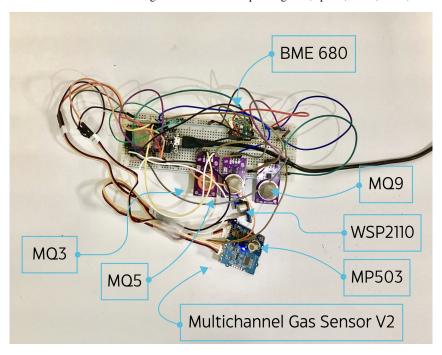


Figure 7: Sensor hardware architecture diagram. The device includes seven gas sensors covering VOCs, alcohol, carbon monoxide, air quality, temperature, and humidity. Components: BME680, MQ-3, MQ-5, MQ-9, WSP2110, MP503, and the Grove Multichannel V2.

Our sensing device was constructed using a suite of commercially available gas sensors, selected for their coverage across a broad range of volatile organic compounds (VOCs), alcohols, and environmental gases. Figure 7 shows a working version of the sensor with all the components labeled, and Table 5 shows all the components we used in making this sensor. We included these information so that readers can recreate the same sensor we used for data collection. These sensors provide complementary chemical sensitivities and were integrated to maximize olfactory coverage in our dataset.

Table 5: Overview of smell sensors used in SMELLNET. The links of all the sensor components are cited for reproducibility.

Sensor	Manufacturer	Description
BME680	Bosch Sensortec Bosch Sensortec (2023)	Low-power sensor measuring gas, atmospheric pressure, temperature, and humidity.
MQ-3	Hanwei Electronics Hanwei Electronics Co., Ltd. (n.d.a)	Metal-oxide sensor optimized for detecting alcohol vapors.
MQ-5	Hanwei Electronics Hanwei Electronics Co., Ltd. (n.d.b)	Detects LPG, natural gas, and coal gas; suitable for combustible gas detection.
MQ-9	Hanwei Electronics Hanwei Electronics Co., Ltd. (n.d.c)	Designed to sense carbon monoxide and combustible gases.
WSP2110	Winsen Electronics Zhengzhou Winsen Electronics Technology Co., Ltd. (n.d.b)	VOC sensor targeting benzene, acetone, toluene, and similar compounds.
MP503	Winsen Electronics Zhengzhou Winsen Electronics Technology Co., Ltd. (n.d.a)	Air-quality sensor responsive to ammonia, hydrogen, and household gases.
Grove Multichannel V2	Seeed Studio Seeed Studio (n.d.)	Modular 4-channel MOX sensor detecting $NO_2$ , $CO$ , $C_2H_5OH$ , and $VOC$ .

# A.2 CONTROLLED ENVIRONMENT

To ensure consistency and minimize external interference during data collection, all sensing sessions were conducted in a controlled environment. During each 10-minute recording interval, we placed both the food sample and the sensor array inside a transparent container. This enclosure prevented environmental factors such as airflow, human movement, or ambient contaminants from affecting the sensor readings. The container allowed gas emitted from the food to accumulate and diffuse evenly, while shielding the sensors from external disturbances such as changes in ambient composition caused by people walking nearby. Between sessions, we ventilated the enclosure to restore ambient conditions and eliminate residual smells from previous trials. After each session, we carefully monitor how the values change overall. Once all the sensor values are stable for 10 minutes, we claim that the environment is stable, and we proceed to the next ingredient for the next session.

Despite these precautions, certain ambient factors, such as temperature, humidity, and background NO<sub>2</sub> levels, varied across different days and could not be entirely eliminated.

#### B SENSOR DATA

# B.1 More PCA

Figure 8 shows PCA projections for nuts, spices, herbs, and vegetables. Across all categories, we observe distinct ingredient-level clustering, indicating that raw sensor signals inherently encode discriminative patterns. Notable examples include *radish* and *sweet potato* among vegetables, *dill* and *angelica* among herbs, and *nutmeg*, *star anise*, and *cumin* among spices. Even in the denser nuts category, ingredients like *pistachios* and *cashew* exhibit identifiable signatures.

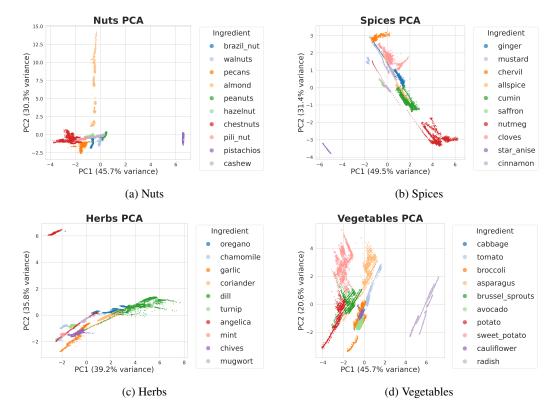


Figure 8: PCA projections of ingredient-level sensor responses for each major category. Each point represents a time step of raw sensor readings, colored by ingredient. Clear clusters are observed, indicating that sensor signals encode discriminative chemical signatures within each food category.

Table 6: Top feature contributions to the first two principal components (PC1 and PC2) of the sensor data. Features are sorted by contribution magnitude.

Feature	PC1	PC2	Magnitude
С2Н5ОН	-0.0015	0.4639	0.4639
VOC	0.0026	0.4620	0.4620
NO2	-0.0002	0.4577	0.4577
Pressure	0.4524	0.0096	0.4525
Altitude	-0.4522	-0.0101	0.4523
Humidity	-0.4521	0.0018	0.4521
Temperature	-0.4500	-0.0102	0.4501
Gas Resistance	0.3334	-0.2416	0.4118
CO	0.0176	0.3967	0.3971
Benzene	0.0565	0.3400	0.3446
Alcohol	0.2295	0.1003	0.2504
LPG	0.1289	0.1420	0.1918

These results support our hypothesis that portable gas sensors capture chemically meaningful variations, both across and within ingredient types, enabling fine-grained classification and motivating representation learning approaches.

Table 6 lists the top feature loadings for PC1 and PC2. PC1 reflects environmental factors (pressure, humidity, temperature, altitude), while PC2 captures volatile compounds (e.g.,  $C_2H_5OH$ , VOC,  $NO_2$ ), illustrating PCA's utility in disentangling physical from chemical influences for downstream interpretation.

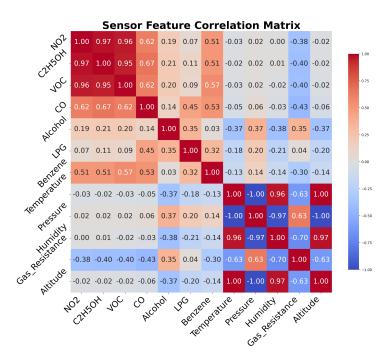


Figure 9: Correlation matrix of all 12 sensor channels, computed using Pearson correlation. Strong correlations among chemical sensors (e.g., NO<sub>2</sub>, C<sub>2</sub>H<sub>5</sub>OH, VOC, and CO) reflect shared response patterns to volatile emissions. Environmental variables (e.g., temperature, pressure, humidity) are also highly interdependent.

#### B.2 FEATURE CORRELATION

Figure 9 shows the Pearson correlation matrix for all 12 sensor channels. Core chemical sensors (NO<sub>2</sub>, C<sub>2</sub>H<sub>5</sub>OH, VOC, CO) exhibit strong positive correlations (r > 0.95), suggesting co-varying responses to shared volatiles. Benzene shows moderate correlation, while LPG and alcohol behave more independently. Environmental features (temperature, pressure, humidity, altitude) are tightly linked, with temperature negatively correlated with pressure and humidity ( $r \approx -0.97$ ). These patterns reveal structured redundancies, highlighting the potential need for feature decorrelation or dimensionality reduction.

#### C GC-MS DATA

#### C.1 GC-MS DATA PROCESSING

To incorporate chemical composition into our framework, we process GC-MS data using compound information from FooDB (FooDB Contributors, 2024), which lists the most abundant volatile compounds for each ingredient. Since these compounds are typically named rather than structured, we use LLM to convert names into molecular formulas. From these, we compute total atomic counts across 18 elements (e.g., C, H, O, N, S, Cl), aggregating the top 10 compounds into a fixed-length vector per ingredient. This GC-MS representation approximates each ingredient's molecular signature and enables downstream tasks such as contrastive learning and multimodal alignment with gas sensor signals (Jaeger et al., 2018; Chithrananda et al., 2020).

# C.2 PCA of GC-MS FEATURES

To better understand the chemical differences between ingredient categories, we performed Principal Component Analysis (PCA) on the raw GC-MS elemental count vectors. As shown in Figure 10, the projection onto the first two principal components reveals distinct clustering patterns. Notably, nuts exhibit a wide spread along PC1, which explains 99.5% of the total variance. Fruits and vegetables show tighter groupings, suggesting more consistent elemental profiles. This clear separation indicates that GC-MS data encodes rich compositional information that differentiates ingredient categories at a molecular level.

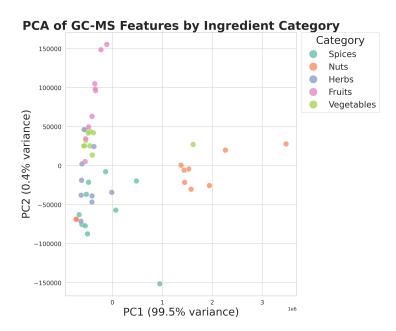


Figure 10: PCA of GC-MS elemental composition across ingredient categories. PC1 accounts for 99.5% of the variance, highlighting dominant compositional differences (e.g., carbon and hydrogen levels).

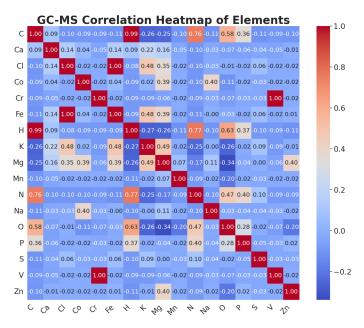


Figure 11: GC-MS correlation heatmap of elemental counts. Strong positive correlations are observed between common organic elements (C, H, N, O), while many trace elements are uncorrelated.

## C.3 CORRELATION STRUCTURE OF GC-MS ELEMENTS

We also analyzed the correlation structure among the 17 elemental features in the GC-MS vectors. Figure 11 shows the resulting correlation heatmap. The strongest positive correlations occur between hydrogen and carbon (r = 0.99), nitrogen and carbon (r = 0.76), and oxygen with both hydrogen and nitrogen. These relationships reflect typical bonding patterns in organic compounds. On the other hand, several elements (e.g., Mn, Zn) are sparsely present and uncorrelated with others, indicating specialized or trace occurrences. This correlation structure highlights both the shared and unique elemental contributions across ingredients.

# D SMELLNET OVERVIEW

Table 7: Descriptive statistics of sensor readings in the training dataset (150,711 samples).

	$NO_2$	C <sub>2</sub> H <sub>5</sub> OH	VOC	CO	Alcohol	LPG	Benzene	Temp. (°C)	Pressure (hPa)	Humidity (%)	Gas Res. $(\Omega)$	Altitude (m)
Mean	97.80	138.33	195.94	792.94	3.41	30.33	1.39e+09	27.14	949.11	46.61	220.89	654.29
Std	118.68	143.40	196.19	61.67	3.28	38.29	2.01e+09	3.28	131.95	27.90	179.99	1274.26
Min	13.00	39.00	26.00	705.00	0.00	2.00	0.00	24.35	688.60	20.15	0.00	-59.47
25%	35.00	65.00	73.00	750.00	1.00	14.00	0.00	25.42	1004.52	27.29	40.65	-28.58
50%	46.00	77.00	106.00	776.00	2.00	23.00	0.00	25.60	1013.99	38.67	220.89	25.00
75%	105.00	140.00	232.50	820.00	5.00	32.00	4.29e+09	25.86	1020.45	45.86	375.42	104.04
Max	753.00	863.00	953.00	1006.00	42.00	507.00	4.29e+09	33.59	1024.19	100.00	704.96	3170.54

Table 8: Descriptive statistics of sensor readings in the testing dataset (29,423 samples).

	NO <sub>2</sub>	C <sub>2</sub> H <sub>5</sub> OH	VOC	CO	Alcohol	LPG	Benzene	Temp. (°C)	Pressure (hPa)	Humidity (%)	Gas Res. $(\Omega)$	Altitude (m)
Mean	92.33	134.08	187.80	791.69	3.47	33.31	1.32e+09	27.02	953.32	45.65	225.79	613.66
Std	113.90	140.97	191.23	60.11	3.14	46.17	1.98e+09	3.21	128.83	27.33	173.61	1244.01
Min	15.00	41.00	26.00	710.00	0.00	3.00	0.00	24.37	688.60	20.32	0.00	-53.70
25%	34.00	65.00	72.00	750.00	1.00	15.00	0.00	25.35	1006.93	27.19	55.64	-30.32
50%	45.00	73.00	93.00	774.00	2.00	23.00	0.00	25.62	1015.15	38.33	226.96	15.36
75%	95.00	135.00	218.00	823.00	6.00	35.00	4.29e+09	25.88	1020.66	45.69	369.13	83.87
Max	775.00	850.00	947.00	1004.00	30.00	444.00	4.29e+09	33.59	1023.49	100.00	685.54	3170.54

#### D.1 SMELLNET SUMMARY

This appendix provides descriptive statistics for all 12 sensor channels in both the training and testing datasets. Tables 7 and 8 summarize key distributional properties, including mean, standard deviation, and range for each feature. The sensor readings exhibit substantial variability across samples, particularly in gas-related channels such as VOC and NO<sub>2</sub>. These statistics highlight the diversity and dynamic range of the collected data, which underpin the challenges of robust model generalization in real-world settings. We also included a text description generated by LLMs of all substances we used for future experiments to enable alignment between text and smell modalities.

#### D.2 DATASET HIERARCHY

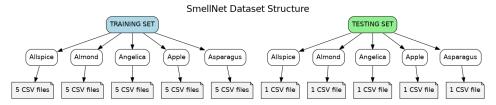


Figure 12: Hierarchical organization of the SmellNet dataset. Each ingredient folder contains multiple CSV files with raw sensor time series data.

Figure 12 illustrates the hierarchical structure of the SMELLNET dataset. Each ingredient is represented as a folder containing multiple time series recordings in CSV format. The training set includes five CSV samples per ingredient to capture variation across trials, while the testing set contains one representative CSV file per ingredient. This structure ensures a consistent, per-ingredient organization and facilitates reproducible supervised learning.

# E ADDITIONAL DETAILS FOR PROBLEM SETUP AND NOTATION

## E.1 GC-MS DESCRIPTOR CONSTRUCTION

When a GC-MS readout is profiled for a substance class, we construct a fixed-length descriptor  $g \in \mathbb{R}^{d'}$  by aggregating element counts over a fixed set  $\mathcal{E}$  (e.g., C, H, O, N, S, Cl). We then apply per-dimension standardization using training-set statistics:

$$\tilde{g}_j = \frac{g_j - \mu_j}{\sigma_j}, \qquad j = 1, \dots, d',$$

and use  $\tilde{g}$  in all GC-MS-aware objectives.

## E.2 MIXTURE-LABEL CONSTRUCTION AND EVALUATION METRICS

**Targets.** Each example in SMELLNET-MIXTURE has a ground-truth composition over K = 12 odorants. Let  $\tilde{\pi} \in \mathbb{R}_{>0}^K$  denote raw proportions; we normalize to the probability:

$$\boldsymbol{\pi}^{\star} = \frac{\tilde{\boldsymbol{\pi}}}{\sum_{i=1}^{K} \tilde{\pi}_{i}}.$$

The model predicts  $\pi = g_{\theta}(x)$  via a softmax head.

#### F BUILDING SCENTFORMER

# F.1 DETAILED SCENTFORMER ARCHITECTURE

**Backbone.** Input windows  $x' \in \mathbb{R}^{w \times d}$  are linearly projected to dimension D. We use sinusoidal positional encodings and prepend a learnable <code>[CLS]</code> token. The model comprises L pre-norm Transformer layers with H attention heads, feed-forward width 4D, and dropout probability p.

**Pooling and heads.** Given encoder output  $H \in \mathbb{R}^{w' \times D}$ , we apply masked mean pooling (default) or [CLS] embedding to obtain  $h \in \mathbb{R}^D$ . A two-layer MLP projects h to 50 logits. Linear heads predict mixture presence  $\hat{u} \in \mathbb{R}^{12}$  and proportions  $\hat{z} = \operatorname{softmax}(W_m h + b_m) \in \Delta^{11}$ .

# F.2 CONTRASTIVE LEARNING LOSS

Given N gas sensor embeddings  $z_i^{(s)}$  and corresponding GC-MS embeddings  $z_i^{(g)}$ , we minimize:

$$\mathcal{L}_{\text{contrastive}} = -\frac{1}{N} \sum_{i=1}^{N} \Bigg[ \log \frac{\exp(\sin(z_i^{(s)}, z_i^{(g)}) / \tau)}{\sum_{j} \exp(\sin(z_i^{(s)}, z_j^{(g)}) / \tau)} + \log \frac{\exp(\sin(z_i^{(g)}, z_i^{(s)}) / \tau)}{\sum_{j} \exp(\sin(z_i^{(g)}, z_j^{(s)}) / \tau)} \Bigg],$$

where sim is cosine similarity and  $\tau$  a temperature.

#### F.3 MIXTURE PREDICTION OBJECTIVE

**Intuition.** Our loss for mixture prediction blends three complementary terms to balance proportion accuracy, robustness, and class imbalance. (i) KL divergence encourages the predicted distribution  $\hat{p} = \operatorname{softmax}(z)$  to match the ground-truth proportions p. (ii) A hinge- $\ell_1$  penalty is applied only to components that are truly present, tightening errors beyond a small tolerance  $\varepsilon$ . (iii) Focal binary cross-entropy (Focal BCE) operates on presence/absence labels and down-weights easy negatives while focusing learning on hard positives. Scalars  $\alpha$  and  $\beta$  balance the second and third terms relative to KL.

**Full objective.** Let  $r_i = \mathbf{1}[p_i > 0]$  indicate presence, and  $S = \{i : r_i = 1\}$  the set of present components. The loss is

$$\mathcal{L} = \mathrm{KL}(p \| \hat{p}) + \alpha \frac{1}{|S|} \sum_{i \in S} \max(|\hat{p}_i - p_i| - \varepsilon, 0) + \beta \cdot \mathrm{FocalBCE}(s, r),$$

with  $\hat{p} = \operatorname{softmax}(z)$ , and FocalBCE using  $(\alpha_f = 0.75, \gamma = 2.0)$  in our experiments. Here,  $\operatorname{KL}(p \| \hat{p})$  promotes globally accurate proportions, the hinge- $\ell_1$  term tightens errors on present components beyond tolerance  $\varepsilon$ , and the focal term addresses class imbalance in presence prediction.

# G ADDITIONAL EXPERIMENTAL SETUP

#### G.1 MODEL DETAILS

This section specifies input/shape conventions, pooling/masking semantics, and per-architecture hyperparameters used in our code. We purposefully omit high-level architecture and objective overviews that appear in the main text and earlier appendices.

**Input & shapes.** Unless noted otherwise, models consume windows  $x \in \mathbb{R}^{B \times T \times F}$  (batch, time, features). For variable-length batches, we pass sequence lengths  $\ell \in \mathbb{N}^B$ . When a layer expects channel-first, we convert to (B, C, T) internally.

**Mask semantics.** Where masks are used, *True* marks *padding*. For mean pooling over valid tokens we use

$$\bar{h}_b = \frac{\sum_t m_{b,t} \, h_{b,t}}{\max \left( \sum_t m_{b,t}, \, 10^{-6} \right)}, \quad m_{b,t} = \mathbf{1} \big[ t < \ell_b \big],$$

and apply the same m to exclude pad tokens from attention or max-pool operations.

# SCENTFORMER

1026

1027

1028

1029

1030

1031 1032

1033

1034

1035 1036

1037

1039

1040

1041

1043

1044 1045

1046

1047

1048

1049 1050

1051

1052

1053

1054

1056

1057

1058

1061

1062

1063 1064

1067

1068

1069 1070

1071 1072

1078

1079

- Input stem: Linear(F → D) then LayerNorm; optional sinusoidal positional encodings added in-place.
- Tokenization: Optional [CLS] (learnable,  $\mathcal{N}(0, 0.02^2)$ ). Pooling is either masked mean (default) or [CLS].
- Encoder: L pre-norm layers with H heads, FFN width 4D, dropout p, activation gelu.
- **Head:** Linear  $(D \rightarrow D/2)$  GELU–Dropout–Linear  $(D/2 \rightarrow C)$ .
- **Defaults:** H=8, L=4, p=0.1, activation=gelu, positional enc.=on, [CLS]=off, pool=mean.

#### LSTMNet

- Core: LSTM $(F \rightarrow H)$  with L layers, bidirectional by default; dropout p only if L > 1.
- Pooling: last (concat final fwd/bwd), masked mean, or masked max.
- Variable length: Uses pack\_padded\_sequence / pad\_packed\_sequence.
- **Projection & head:** Linear( $\cdot \rightarrow D_{\text{emb}}$ ) then Linear( $D_{\text{emb}} \rightarrow C$ ).
- **Defaults:** L=1, bidirectional, p=0.1, pool=mean.

## **CNN1D** classifier

- · Layout: Stack of
  - Convld( $C_{in} \rightarrow C_{out}$ , same padding via k//2)-BatchNorm-ReLU-(Dropout).
- **Head:** Global average pooling over T then Linear  $(C' \rightarrow C)$ .
- Channel order: Accepts (B, T, C) (channel\_last=true) or (B, C, T); we coerce to (B, C, T) internally.
- **Defaults:** channels=(64,128,256), kernel size k=5, dropout 0.2, BatchNorm on, channel\_last=true.

#### MLP classifier

- **Pooling:** If input is (B,T,C) or (B,C,T), pool over T via mean/max (default: mean). Optionally flatten requires fixed T with input dim  $C \times T$ .
- Backbone: Repeated [Linear-(BatchNorm)-ReLU-(Dropout)] blocks; head is Linear  $\rightarrow C$ .
- **Defaults:** hidden sizes (256,256), BatchNorm on, dropout 0.2, pool=mean, channel\_last=true.

#### **GC-MS MLP encoder**

- Stem: Optional LayerNorm on input, then MLP with ReLU and optional Batch-Norm/Dropout, ending in Linear→ D.
- Normalization: Optional ℓ<sub>2</sub> normalization of the final embedding when used in contrastive objectives.
- **Defaults:** hidden (512,256), D=256, dropout 0.1, LayerNorm on, BatchNorm off,  $\ell_2$  off.

## Hyperparameters

Compone	ent I	Key defaults / toggles
Transform	ner <i>1</i>	H=8, $L=4$ , FFN = $4D$ , $p=0.1$ , GELU, PE on, CLS off, pool=mean
LSTMNe	et <i>1</i>	$L=1$ , bi=True, $p=0.1$ , pool $\in$ {mean,last,max}, $D_{emb}$ as set
CNN1D	C	channels= $(64,128,256)$ , $k=5$ , BN on, dropout 0.2, channel_last=true
MLP	h	nidden=(256,256), BN on, dropout 0.2, pool $\in$ {mean,max,flatten}
GC-MS e	enc. h	nidden=(512,256), $D$ =256, LayerNorm on, BN off, dropout 0.1, $\ell_2$ off

#### G.2 Training hyperparameters

We standardize training across baselines and SCENTFORMER to ensure a fair comparison and to avoid overfitting to any one configuration.

Table 9: **Single-ingredient classification with all 12 sensor channels** (no channel dropping). Setup mirrors Table 2: we vary window size  $(w \in \{50, 100\})$  and temporal differencing  $(p \in \{0, 25\})$ . Values are from our reproduction (seed=42); metrics in %.  $\Delta$  reports cross-modal minus sensor-only Acc@1.

Model	Window	Grad.	Sensor-only (RQ1)			Cross-modal (RQ2)			$\Delta$ Acc@1
Model	Williaow	Grau.	Acc@1↑	Acc@5↑	F1↑	Acc@1↑	Acc@5↑	F1↑	(X–S)
Transf. (12ch)	50	0	44.0	79.6	40.6	44.8	77.8	41.2	+0.8
Transf. (12ch)	50	25	65.3	94.6	64.6	65.9	93.4	65.2	+0.6
Transf. (12ch)	100	0	45.0	81.7	40.6	45.4	80.4	42.3	+0.4
Transf. (12ch)	100	25	71.3	96.8	70.6	71.3	93.0	70.1	+0.0

**Temporal differencing.** To quantify the value of short-range dynamics, we evaluate fixed lags  $p \in \{0, 25\}$  when forming first-order temporal differences  $\Delta x_t = x_t - x_{t-p}$  (Sec. 4.2). This follows prior evidence in our setting that differencing can substantially improve discriminability.

**Sliding-window segmentation.** We segment streams into overlapping windows of length  $w \in \{50, 100\}$  (Sec. 4.2). The shorter window (w = 50) increases the number of training/evaluation samples, whereas the longer window (w = 100) trades sample count for more stable temporal context (App. I.1).

**Learning rate selection.** To keep model comparisons robust and reproducible, we tune over a small, fixed grid shared by all methods:  $\{3 \times 10^{-4},\ 10^{-3},\ 3 \times 10^{-3}\}$ . We report the checkpoint with the best validation Top-1. Limiting the grid prevents "hyperparameter fishing" and reduces variance attributable to optimizer settings.

**Epoch budgets and batching.** For SMELLNET-BASE classification we train for 90 epochs with batch size 32; for SMELLNET-MIXTURE distribution prediction we train for 60 epochs with batch size 64. Using fixed epoch budgets across models minimizes variance due to training length; the best checkpoint is chosen by validation Top-1.

**Randomization.** We fix the Python-level random seed to 42 for reproducibility across all experiments.

**GC-MS supervision.** Contrastive alignment with GC-MS is used only for the single-ingredient classification setting, where per-ingredient GC-MS signals are available. We do not apply GC-MS supervision to mixture prediction because reliable GC-MS profiles for arbitrary mixtures are not available at scale.

**Design choice for mixtures (no temporal differencing).** For SMELLNET-MIXTURE, sensor streams are recorded at 10 Hz on a four-channel array. At this sampling rate, small lags yield negligible signal change, whereas larger lags substantially reduce the effective number of windows. We therefore train mixture models on raw (non-differenced) windows.

# H ABLATION STUDY AND EXPERIMENT ANALYSIS

#### H.1 FULL CHANNEL ANALYSIS

Using all 12 channels yields comparable Acc@1 and sometimes higher raw scores, but cross-modal gains are small/inconsistent and top-5 ranking degrades: at w=50, p=25, Acc@5 drops -1.2 points with GC-MS; at w=100, p=25, Acc@5 drops -3.8 points, and  $\Delta$ Acc@1 is  $\approx$  0.0. This pattern suggests that a small portion of irregular values in some channels injects noise that hurts representation alignment and ranking stability, even if the classifier can partially compensate in Acc@1. Consistent with the main text, we therefore retain the six stable channels (NO<sub>2</sub>, C<sub>2</sub>H<sub>5</sub>OH, VOC, CO, Alcohol, LPG) for all reported results, while releasing the full 12-channel data for future filtering and denoising efforts.

# H.2 CHANNEL IMPORTANCE THROUGH MASKING

All reported values are negative  $\Delta Acc@1$ , i.e., masking any single channel reduces Top-1 accuracy. With temporal differencing (p=25; Table 10), the largest drops occur for LPG (-28.9% to -28.1%), followed by VOC and Alcohol, indicating these channels carry the most decisive information under different preprocessing. Without differencing (p=0; Table 11), NO<sub>2</sub> and CO dominate the impact (up to -25.2%), suggesting the raw-signal model leans more on oxidizing/CO-related responses.

Table 10: Single-channel mask ablation (SCENTFORMER, window=50, temporal difference p=25). Entries are *negative* point changes in Acc@1 when masking a channel (masking always reduces accuracy). Baselines: Cross-modal: 50.97%, Sensor Only=0: 50.60%.

Cross-m	odal	Sensor (	Only
Channel (gas)	Δ Acc@1	Channel (gas)	Δ Acc@1
LPG	-28.07	LPG	-28.90
VOC	-24.65	Alcohol	-26.50
Alcohol	-23.92	VOC	-24.10
CO	-18.93	CO	-22.07
$C_2H_5OH$	-17.36	$NO_2$	-19.76
$NO_2$	-15.60	$C_2H_5OH$	-10.25

Table 11: Single-channel mask ablation (SCENTFORMER, window=50, no temporal difference p=0). Entries are negative point changes in Acc@1 when masking a channel (masking always reduces accuracy). Baselines: Cross-modal: 36.28%, Sensor Only: 35.13%.

Cross-m	odal	Sensor Only				
Channel (gas)	Δ Acc@1	Channel (gas)	Δ Acc@1			
$\overline{NO_2}$	-21.27	CO	-25.15			
CO	-18.80	$NO_2$	-22.77			
$C_2H_5OH$	-18.71	$C_2H_5OH$	-21.01			
VOC	-16.15	VOC	-14.03			
Alcohol	-9.36	Alcohol	-12.44			
LPG	-8.91	LPG	-10.41			

Contrastive training slightly reshapes importance but preserves the main ordering within each preprocessing regime. Overall, the strictly negative  $\Delta Acc@1$  across all cells reinforces that each gas channel contributes uniquely; the magnitude pattern shifts with temporal preprocessing and contrastive alignment.

## H.3 TIMESTAMP SIZE ANALYSIS

To justify our 10-minute recording interval, we conducted an ablation that varies the number of initial time steps fed to SCENTFORMER(window size w=50) under differencing lags  $p \in \{0, 25\}$  (Table 12). Using only the first 200 steps yields lower accuracy; as the number of steps increases, performance improves and then plateaus around 400-600 steps. With temporal differencing (p=25), Acc@1 increases from 45.7 at 200 steps to 53.2 at 500 steps and remains essentially unchanged at  $\approx$ 600 steps (50.6/50.9 for sensor/cross-modal). This saturation suggests that a 10-minute window captures sufficient temporal dynamics for robust recognition, while longer acquisitions provide diminishing returns. Based on these preliminary findings and practical data-collection constraints, we therefore adopt 10-minute intervals throughout.

# H.4 RUNTIME AND MEMORY ANALYSIS

All experiments use SCENTFORMER on an NVIDIA L40S (compute capability 8.9; driver 550.54.14; CUDA 12.4; 46,068 MiB VRAM) with FP32, batch size = 32, and a  $\sim$  2.4109M-parameter model (live memory  $\sim$  18.9626 MB, peak GPU  $\leq$  75.4966 MB). Inference latency is extremely low across settings: mean per-window latency ranges from 0.0191-0.0479 ms, and throughput remains high at 20,892-52,371 windows/s. As measurements are forward-only, temporal differencing is irrelevant; and at this batch size, both the presence of contrastive training and the window size have negligible practical impact on latency.

Table 12: Ablation on the number of initial time steps used (SCENTFORMER, window size w=50). p denotes the temporal differencing lag (in samples). Accuracies and F1 are reported in %. Increasing beyond 400-600 steps yields diminishing returns.

# Timestamps	p	Sensor (S	CENTFOR	MER, $w = 50$ )	Cross-N	Modal (GC-	F1 32.6 30.7 32.8 44.5 33.4		
	F	Acc@1	Acc@5	F1	Acc@1	Acc@5	F1		
200	0	33.1	78.2	30.9	34.5	70.2	32.6		
200	25	45.7	80.8	43.1	35.6	66.1	30.7		
300	0	35.2	70.1	31.1	36.2	69.4	32.8		
300	25	48.3	84.3	46.7	46.7	77.0	44.5		
400	0	35.4	73.7	32.0	37.8	68.4	33.4		
400	25	50.5	83.3	49.2	50.5	81.8	48.6		
500	0	34.6	78.2	32.0	43.9	77.1	41.5		
500	25	53.2	87.4	51.2	50.9	83.7	49.4		
≈600	0	35.1	70.9	33.1	36.2	72.2	33.3		
≈600	25	50.6	85.0	49.5	50.9	81.5	50.4		

Table 13: Latency and resource metrics for **Scentformer** (batch=32, FP32) on NVIDIA L40S.

Variant	Win	Mean (s)	WPS	Eval (s)	GPU MB	Live MB	Params (M)
No contrastive	50	0.0191	52371	0.0189	63.3540	18.9626	2.4109
Contrastive	50	0.0203	49237	0.0310	61.0024	18.9626	2.4109
No contrastive	100	0.0479	20892	0.0218	73.6172	18.9626	2.4109
Contrastive	100	0.0420	23821	0.0151	75.4966	18.9626	2.4109

# I MATH

# I.1 Number of Windows

Let T be the number of time steps in a recording, w the window length (in steps), and s the stride. Using valid (no-padding) windows, the number of extracted windows is

$$N(T, w, s) = \begin{cases} \left\lfloor \frac{T - w}{s} \right\rfloor + 1, & T \ge w, \\ 0, & T < w. \end{cases}$$
 (1)

We use 50% overlap, i.e.,  $s = \frac{w}{2}$ . Thus a 10-minute recording at 1 Hz has T = 600 steps and yields

$$N(600, 50, 25) = \left\lfloor \frac{600 - 50}{25} \right\rfloor + 1 = 22 + 1 = 23, \quad N(600, 100, 50) = \left\lfloor \frac{600 - 100}{50} \right\rfloor + 1 = 10 + 1 = 11.$$

With sampling rate  $f_s$  (Hz) and duration L (s),  $T = f_s L$ , and each window spans  $w/f_s$  seconds. If right-padding is used to include a final partial window, replace the floor in (1) with a ceiling.

## I.2 TOP-K PERFORMANCE CALCULATION

For example  $n \in \{1, ..., N\}$  with class probabilities  $\boldsymbol{p}^{(n)} = \operatorname{softmax}(\boldsymbol{s}^{(n)}) \in [0, 1]^C$ , let  $R_n = \{c : y_c^{(n)} > 0\}$  be the (nonempty) set of truly present classes and  $P_n = |R_n|$  its size. Let  $\pi_n(1), ..., \pi_n(C)$  index classes in descending  $p^{(n)}$ , and  $\Pi_n(k) = \{\pi_n(1), ..., \pi_n(k)\}$ . Our metric is the label-recall with a *per-example* cutoff  $K_n = P_n$ :

$$DynTopK = \frac{\sum_{n=1}^{N} |R_n \cap \Pi_n(P_n)|}{\sum_{n=1}^{N} |R_n|}$$
 (×100% for reporting). (2)

With N = 12 examples. Define

$$H = \sum_{n=1}^{12} |R_n \cap \Pi_n(P_n)|, \qquad M = \sum_{n=1}^{12} |R_n|.$$

Then

DynTopK = 
$$\frac{H}{M} \times 100\%$$
.

 Special case (one present class per example). If  $|R_n| = 1$  for all n (single-label data), then M = 12 and

DynTopK =	$\frac{\text{\#hits}}{12}$	×	100%.
-----------	----------------------------	---	-------

Hits out of 12	DynTopK (%)
0	0.0%
1	8.3%
2	16.7%
3	25.0%
4	33.3%
5	41.7%
6	50.0%
7	58.3%
8	66.7%
9	75.0%
10	83.3%
11	91.7%
12	100.0%

#### J DETAILS OF SMELL MIXTURES

For the collection of mixture data, each base odorant was treated as a distinct class, yielding 12 base classes. Beyond these bases, we constructed binary and ternary mixtures of the base odorants at fixed volumetric ratios, resulting in a total of 126 unique mixture combinations (54 in training, 45 in test-seen, 27 in test-unseen). The dataset exhibits the following mixture distribution:

- Base odorants: 24.7% of training sessions (168 sessions), 22.3% of test-seen (48 sessions), 0% of test-unseen
- **Binary mixtures**: 71.3% of training sessions (484 sessions), 73.5% of test-seen (158 sessions), 83.2% of test-unseen (153 sessions)
- **Ternary mixtures**: 4.0% of training sessions (27 sessions), 4.2% of test-seen (9 sessions), 16.8% of test-unseen (31 sessions)

The composition of odor readings are shown in Fig. 3. Binary mixture ratios in the training set span multiple combinations including 20/80 (160 sessions), 50/50 (174 sessions), and 80/20 (120 sessions), with additional ratios spanning from 10/90 to 90/10. Ternary mixtures include both balanced (33/33/33) and asymmetric (10/30/60) distributions. This comprehensive ratio coverage enables the model to learn ratio prediction across the full spectrum of possible combinations.

## K LLM USAGE

We used a large language model (LLM) solely for light copy-editing (grammar, clarity, phrasing). No technical content, experiments, analyses, citations, or claims were generated by the LLM. All text was verified and edited by the authors, who take full responsibility for the content.

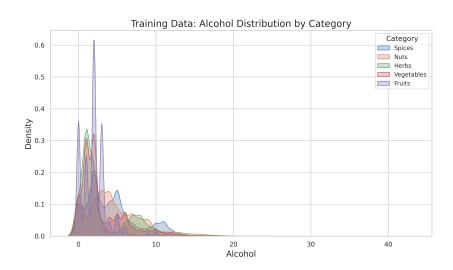


Figure 13: **KDE** distribution of Alcohol Sensor by Category. These numbers represent raw, unfiltered data readings. Normalizations and filtration of abnormal channels are performed as preprocessing steps before modeling. We provide standard kits for preprocessings, but the raw values are provided to preserve as much information as possible.

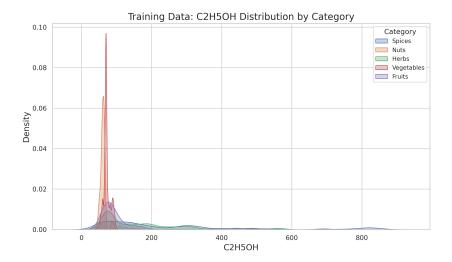


Figure 14: KDE distribution of C2H5OH Sensor by Category.

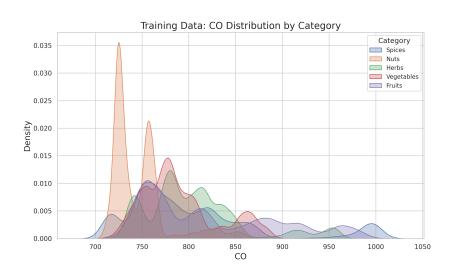


Figure 15: KDE distribution of CO Sensor by Category.

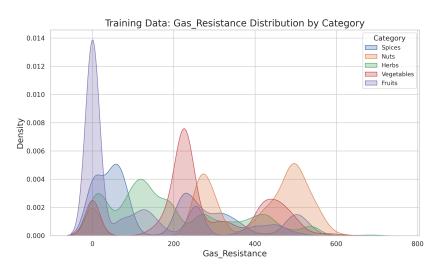


Figure 16: KDE distribution of Gas Resistance Sensor by Category.

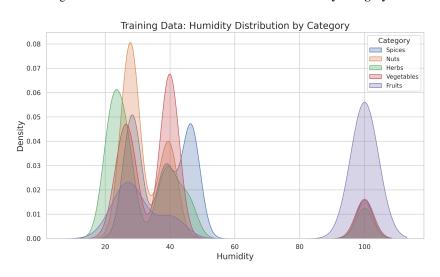


Figure 17: KDE distribution of Humidity Sensor by Category.

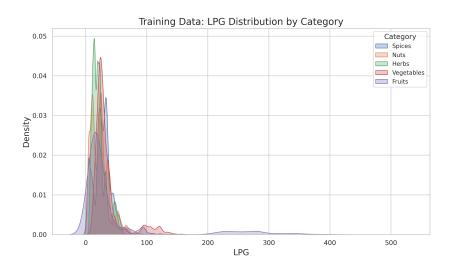


Figure 18: KDE distribution of LPG Sensor by Category.

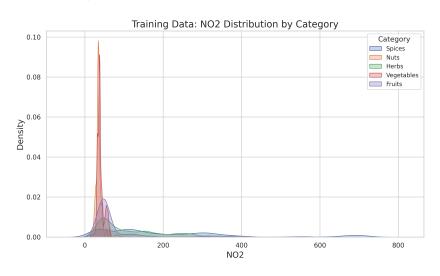


Figure 19: KDE distribution of NO2 Sensor by Category.

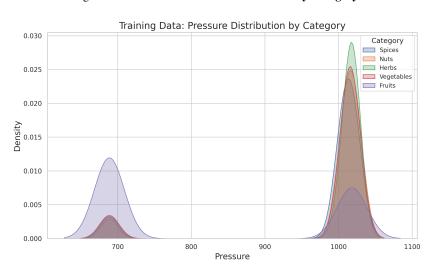


Figure 20: KDE distribution of Pressure Sensor by Category.

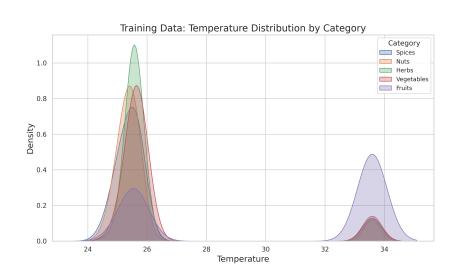


Figure 21: KDE distribution of Temperature Sensor by Category.

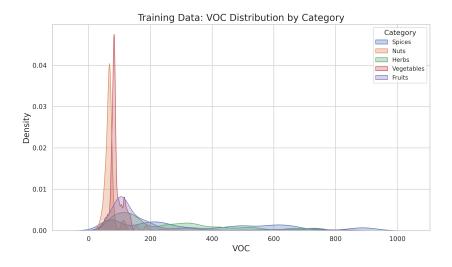


Figure 22: KDE distribution of VOC Sensor by Category.

Table 14: Olfactory materials used in SMELLNET, organized alphabetically by label. We distinguish essential oils (volatile natural extracts with identifiable key odorants), flavor extracts (culinary preparations in oil or alcohol carriers), fragrance oils (synthetic or proprietary blends), and cosmetic oils (carrier oils with an odor).

Label	Туре	Material	Vendor
Almond	Flavor Extract	Pure Almond Extract	CADIA
Apple	Fragrance Oil	Apple Fragrance Oil	P&J Trading
Banana	Cosmetic Oil	Banana Oil	The Aromatherapy Shop Ltd
Clove	Essential Oil	Clove Bud Essential Oil	365 Whole Foods Market
Coriander	Essential Oil	Coriander Essential Oil	Skylara Essentials
Cumin	Essential Oil	Cumin Essential Oil	Silky Scents
Garlic	Essential Oil	Garlic Essential Oil	Skylara Essentials
Mango	Essential Oil	Mango Essential Oil (Egypt)	The Aromatherapy Shop Ltd
Orange	Flavor Extract	Orange Flavor	Frontier Co-op Store
Peach	Fragrance Oil	Peach Fragrance Oil	P&J Trading
Pear	Fragrance Oil	Pear Fragrance Oil	P&J Trading
Strawberry	Fragrance Oil	Strawberry Fragrance Oil	P&J Trading

Table 15: **Statistics of Gas Sensor Measurements - Part 1.** These numbers represent raw, unfiltered data readings. Normalizations and filtration of abnormal channels are performed as preprocessing steps before modeling. We provide standard kits for preprocessings, but the raw values are provided to preserve as much information as possible.

	N(	O2	С2Н5ОН		VC	OC .	C	)	Alcohol		LP	G
Ingredient	mean	std	mean	std	mean	std	mean		mean		mean	std
allspice	257.35	47.64	298.64	46.79	493.35	72.72	835.91	23.98		0.52	34.23	1.75
almond	58.10	15.99	112.52	26.03	164.75	61.27	735.84	12.07	4.38	0.66	20.64	1.09
angelica	143.16	41.41	175.46	32.63	316.82	69.57	769.97	9.79	1.07	0.73	8.91	1.89
apple	45.38	3.91	66.93	2.40	126.27	28.85	750.08	7.97	2.25	0.44	10.58	0.53
asparagus	54.49	6.42	87.34	7.87	117.50	15.95	780.18	14.74	4.26	1.84	18.93	3.44
avocado	31.56	1.11	69.69	1.31	77.40	3.22	860.57	14.95	1.72	0.45	24.55	3.32
banana	58.15	7.60	89.30	7.99	154.88	34.11	804.62	18.02	2.71	0.52	31.82	3.10
brazil nut	32.12	3.84	58.70	3.63	61.27	8.54	727.97			0.59	23.91	1.10
broccoli	35.14	3.87	69.14	5.73	72.25	11.89	855.23		1.05		34.33	4.2
brussel sprouts		5.18	65.20	3.23	84.28	8.67	744.62		3.73		20.03	3.8
cabbage	37.25	1.86	70.06	1.48	82.76	3.74	776.35		1.67	0.48	25.44	1.1
cashew	34.52	1.28	66.05	0.84	67.58	2.70	757.08	1.83	5.59		13.81	2.5
cauliflower	36.75	1.08	69.59	0.90	81.63	1.98	777.96		1.29	0.46	30.06	1.6
chamomile	47.79	5.60	69.86	7.89	104.80	14.10	752.26		1.41	0.73	12.40	1.7
chervil						5.74						
	38.36	3.65	62.39	1.86	77.55	3.74 8.48	718.41				33.60	5.4
chestnuts	34.51	5.07	60.54	2.88	67.40		819.79				54.55	9.1
chives	46.26	4.93	80.56	6.78	114.11	16.46	801.50				22.74	2.9
cinnamon	163.05	37.50	182.66	31.10	374.53	96.46	787.42			1.01	23.16	2.5
cloves	121.25	34.45	105.63	25.27	238.72	69.35	778.28		11.29			4.0
coriander	237.54		275.84		481.83	70.99	830.32		1.42			4.3
cumin	322.36		454.09	62.83	608.59					1.63	31.79	3.0
dill	305.26		472.43	85.59			927.31		9.29	4.52		5.2
garlic	75.73	11.98	111.27	23.07	149.23		826.53	26.66		0.73	15.95	1.9
ginger	281.14	53.82	281.85	45.31	585.09		810.70	17.01		0.50	20.41	0.7
hazelnut	35.27	1.92	62.76	0.71	71.75	2.57	727.42	2.93	3.62	1.15	26.62	1.8
kiwi	30.99	2.33	63.09	2.98	63.09	6.45	771.93	2.86	1.02	0.15	20.80	1.1
lemon	373.19	108.93	497.90	109.86	672.55	120.62	877.10	44.48	2.89	0.50	14.61	3.6
mandarin	111.15	35.02	145.58	33.67	246.67	74.52	872.39	35.47	1.59	0.53	23.36	3.3
mango	49.58	6.39	82.23	8.90	115.03	13.03	757.49	7.69	2.04	0.22	11.91	0.8
mint	60.52	4.89	81.03	6.27	141.44	17.75	745.44	4.56	0.17	0.38	12.38	1.3
mugwort	116.18	19.94	171.32	33.38	264.83	41.12	786.64		1.19	0.43	18.76	1.0
mustard	29.07	3.12	65.48	3.93	63.17	8.33	754.74			0.44	10.43	1.8
nutmeg			789.17	94.66	850.61	99.73	985.72		6.76		66.55	25.
oregano			245.84		344.29	48.64	794.02		7.09	3.08	45.97	7.7
peach	64.74	23.37	153.68	67.61	224.51	86.09	956.07		5.98		262.50	
peanuts	35.84	2.00	64.09	2.16	63.25	5.72	753.14	2.78	1.60	1.18	8.60	1.6
pear	40.39	3.34	67.87	2.88	98.02	9.45	802.24		2.32		18.23	2.9
1	25.14	3.14	49.67	3.58	47.55	8.22	726.83		7.28		30.81	3.5
pecans	33.74	1.63	57.68	0.52	67.48	1.88	726.42				36.68	2.0
pili nut	47.09	5.87					897.35					
pineapple			82.72	6.47	102.16	12.14						19.
pistachios	39.79	4.90	56.79	3.38	69.73	9.04	720.86		0.00		4.78	0.8
potato	32.79	5.49	56.15	4.86	65.47		745.50					4.3
radish	46.52	8.50	82.89	9.08	102.48		787.99					
saffron	110.59	13.60	150.79	12.46						0.12		2.4
star anise	91.89	21.29	115.33				744.94		0.00		4.47	0.6
strawberry	54.26	6.89	91.86	10.47	108.25	12.18	749.20				3.67	0.5
sweet potato	38.58	2.85	73.99	4.61	98.93	9.31	763.77	11.52			107.87	16.
tomato	39.76	10.09	73.82	6.74	90.18	19.60	809.73	13.25	1.60	0.49	22.61	1.9
turnip	30.31	0.79	67.94	1.42	71.80	2.71	821.85				22.17	1.6

Table 16: **Statistics of Gas Sensor Measurements - Part 2.** These numbers represent raw, unfiltered data readings. Normalizations and filtration of abnormal channels are performed as preprocessing steps before modeling. We provide standard kits for preprocessings, but the raw values are provided to preserve as much information as possible.

	Temperature		Press		Humi	-		sistance	Alti	
Ingredient	mean	std	mean	std	mean	std	mean	std	mean	std
allspice	25.60	0.04	1022.69	0.22	28.23	0.10	79.13	54.95	-47.07	1.84
almond	25.24	0.04	1007.47	0.38	28.30	0.47		26.24	79.32	3.20
angelica	31.98	3.21	755.47	133.04	84.79	30.26	39.09	98.19	2525.76	1282.78
apple	33.59	0.00	688.60	0.00	100.00	0.00	0.00	0.00	3170.54	0.00
asparagus	25.66	0.07	1018.41	0.35	26.09	1.49	319.26	63.33	-11.69	2.87
avocado	25.56	0.02	1014.76	0.20	40.16	0.74	237.52	9.22	18.63	1.67
banana	25.58	0.03	1022.44	0.37	29.44	0.77	430.72	39.71	-45.04	3.05
brazil nut	25.47	0.09	1021.20	0.12	28.60	0.22	481.79	15.60	-34.81	1.01
broccoli	25.53	0.15	1011.67	0.21	40.38	0.62	183.54	24.03	44.32	1.77
brussel sprouts	25.40	0.18	1019.15	0.12	24.59	1.19	425.15	55.98	-17.84	0.96
cabbage	25.98	0.05	1011.51	0.08	39.22	0.49	229.04	7.63	45.61	0.68
cashew	25.15	0.11	1008.86	0.17	39.85	0.44	262.52	9.40	67.76	1.39
cauliflower	25.86	0.13	1011.55	0.08	39.68	0.65	215.34	7.51	45.35	0.70
chamomile	25.56	0.13	1015.89	4.93	30.18	7.31	413.34	117.84	9.31	40.95
chervil	25.80	0.06	1021.89	0.14	26.99	0.10	505.58	18.81	-40.49	1.17
chestnuts	25.61	0.03	1023.09	0.25	28.17	0.82	473.14	42.31	-50.35	2.03
chives	25.49	0.05	1018.33	0.04	26.96	0.59	334.81	45.50	-11.02	0.31
cinnamon	25.26	0.38	1005.90	1.33	35.95	7.88	212.44	117.63	92.53	11.13
cloves	25.00	0.21	1021.52	0.24	39.52	0.43	253.76	19.85	-37.40	1.96
coriander	25.57	0.13	1018.10	0.08	21.23	0.69	130.46	32.56	-9.12	0.67
cumin	25.63	0.08	1020.75	0.07	37.84	0.19	55.22	17.99	-31.05	0.55
dill	25.62	0.03	1018.49	0.08	26.59	0.62	108.59	41.22	-12.31	0.65
garlic	25.27	0.30	1018.61	0.45	21.65	0.51	179.96		-13.37	3.71
ginger	25.24	0.08	1008.69	0.31	29.14	0.45	336.80	37.57	69.17	2.62
hazelnut	25.59	0.04	1021.57	0.17	28.11	0.17	498.52	15.33	-37.83	1.43
kiwi	25.30	0.19	1013.12	0.36	40.33	0.81	254.57	9.22	32.26	2.99
lemon	33.59	0.00	688.60	0.00	100.00	0.00	0.00	0.00	3170.54	0.00
mandarin	25.58	0.09	1018.25	0.12	24.76	2.39	103.54	27.71	-10.36	1.02
mango	33.59	0.00	688.60	0.00	100.00	0.00	0.00	0.00	3170.54	0.00
mint	25.48	0.11	1020.56	0.10	22.61	0.35	395.47	43.60	-29.49	0.79
mugwort	25.87	0.07	1020.60	0.12	37.42	0.29	130.43	21.31	-29.79	0.95
mustard	24.98	0.18	1009.82	0.31	47.51	0.79	221.07	7.46	59.76	2.57
nutmeg	25.29	0.21	1004.76	0.25	46.82	0.61	9.71	18.84	102.06	2.07
oregano	25.58	0.10	1004.49	0.07	45.78	0.21	26.47	5.24	104.26	0.55
peach	25.65	0.05	1019.03	0.22	26.42	0.75	140.87	48.48	-16.85	1.78
peanuts	25.35	0.43	1009.95	0.64	37.34	1.71	299.68	11.93	58.67	5.38
pear	33.59	0.00	688.60	0.00	100.00	0.00	0.00	0.00	3170.54	0.00
pecans	25.16	0.27	1023.73	0.20	27.04		552.87	18.62	-55.64	1.66
pili nut	25.50	0.06	1023.73	0.14	25.38	0.13	507.05	12.53	-41.33	1.13
pin nut pineapple	33.59	0.00	688.60	0.00	100.00		0.00	0.00	3170.54	0.00
pistachios	33.59	0.00	688.60	0.00	100.00		0.00	0.00	3170.54	0.00
potato	25.19		1021.77	0.00	27.73		489.65	35.72	-39.49	0.12
radish	33.59		688.60	0.00	100.00	0.00	0.00	0.00	3170.54	0.12
saffron	25.69	0.00	1004.32	0.08	45.41	0.10	52.87	10.69	105.71	0.64
					100.00					
star anise	33.59 33.59	0.00	688.60	0.00			0.00	$0.00 \\ 0.00$	3170.54	0.00
strawberry	25.77	0.00	688.60	0.00	100.00	0.00	0.00		3170.54	
sweet potato		0.04	1019.97	0.29	26.91	0.81		18.80	-24.60	2.44
turnin	25.58	0.05	1014.10	0.22	40.02		225.41	19.23	24.12	1.83
turnip	25.54		1015.79	0.25	40.77		195.49	15.89	10.01	2.09
walnuts	25.16	0.23	1008.69	0.28	39.91	0.74	271.63	10.47	69.17	2.34