

Article

High-Accuracy Detection of Odor Presence from Olfactory Bulb Local Field Potentials via Deep Neural Networks

Matin Hassanloo ^{1,*} , Ali Zareh ¹  and Mehmet Kemal Özdemir ² ¹ Department of Computer Engineering, Istanbul Medipol University, Kavacık Campus, Istanbul 34810, Turkey; ali.zareh@std.medipol.edu.tr² Department of Artificial Intelligence Engineering, Istanbul Medipol University, Kavacık Campus, Istanbul 34810, Turkey; mkozdemir@medipol.edu.tr

* Correspondence: martin.hassanloo@std.medipol.edu.tr

Abstract

Odor detection underpins food safety, environmental monitoring, medical diagnostics, and many more fields. Current artificial sensors developed for odor detection struggle with complex mixtures, while non-invasive recordings lack reliable single-trial fidelity. To develop a general system for odor detection, in this study we present preliminary work where we test two hypotheses: (i) that spectral features of local field potentials (LFPs) are sufficient for robust single-trial odor detection and (ii) that signals from the olfactory bulb alone are adequate. To test these hypotheses, we propose an ensemble of complementary one-dimensional convolutional networks (ResCNN and AttentionCNN) that decodes the presence of odor from multichannel olfactory bulb LFPs. Tested on 2349 trials from seven awake mice, our final ensemble model supports both hypotheses, achieving a mean accuracy of 86.2%, an F1-score of 85.3%, and an AUC of 0.942, substantially outperforming previous benchmarks. The t-SNE visualization confirms that our framework captures biologically significant signatures. These findings establish the feasibility of robust single-trial detection of odor presence from extracellular LFPs and demonstrate the potential of deep learning models to provide deeper understanding of olfactory representations.

Keywords: olfactory neural signals; deep neural networks; odor detection; extracellular recordings; time–frequency analysis; local field potentials; brain–computer interfaces

1. Introduction

The development of advanced sensor technologies for the reliable detection of odorants is a critical challenge in different fields ranging from environmental safety to medical diagnostics [1–5]. Moving beyond the limitations of conventional electronic noses (e-noses), the field of neural sensing aims to emulate the olfactory system, using its rapid and complex processing to create high-performance brain–computer interfaces (BCIs) [6–11]. Central to this approach is the decoding of neural signals, and among these, the LFP offers a uniquely powerful data source. Unlike the sparse firing patterns of individual neurons, the LFP represents the synchronized synaptic activity of thousands of neurons, offering a robust signal for single-trial classification [12–14].

LFP recordings overcome the limitations of non-invasive methods in two critical aspects. First, direct LFP recordings offer superior spatial resolution, eliminating contamination from non-neural sources [13]. Second, these recordings provide a high signal bandwidth that allows analysis of high-frequency oscillations known to be important for



Academic Editor: Anastasios Doulamis

Received: 16 November 2025

Revised: 18 January 2026

Accepted: 21 January 2026

Published: 2 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

olfactory processing [15]. By capturing this clean and full spectrum of neural activity, LFP recordings offer a powerful and decodable feature source for the specific task of binary odor detection, potentially overcoming the limitations of traditional sensors.

Previous studies on the decoding of odor presence from neural signals have generally focused on non-invasive methods, which suffer from a critical loss of high-fidelity spectral information [16–18]. Scalp electroencephalography (EEG), for instance, struggles with a low signal-to-noise ratio (SNR) due to the deep cortical source of olfactory signals [19,20]. Consequently, achieving successful classification with these non-invasive signals requires methods that are impractical for real-world applications. For example, many studies rely on averaging responses across dozens of trials, making real-time detection impossible [21,22]. The challenge of non-invasive decoding is illustrated in a recent study by Rajabi et al., who targeted olfactory bulb (OB) activity non-invasively with an electrobulbogram (EBG) and a 1D CNN. Their model achieved an Area Under the Curve (AUC) of 0.58, indicating a performance close to the binary chance level for single-trial classification [23].

We introduce an ensemble of complementary convolutional neural networks, combining an AttentionCNN and a ResCNN, to determine odor presence from 32-channel extracellular LFP recordings. We propose that using high-fidelity direct LFP recordings instead of low-fidelity non-invasive signals will help close the existing performance gap. Based on the proposed ensemble model, we propose two central hypotheses. First, we hypothesize that the spectro-temporal features within single-trial LFPs provide sufficient information to robustly and accurately classify odor-presence versus odor-absence conditions. Second, we hypothesize that neural activity from the olfactory bulb alone is sufficient for the odor presence detection task, without requiring signals from downstream processing areas like the piriform cortex (PCx). By developing an ensemble of deep convolutional neural networks [24,25], we demonstrate a significant leap in performance, thereby establishing the feasibility of using LFP recordings as a foundation for odor sensing.

2. Materials and Methods

In this section, we detail the pre-processing steps applied to LFP signals, introduce two core architectures (AttentionCNN and ResCNN), and explain the procedures for training and ensembling these models, along with the evaluation metrics. Figure 1 provides an overview of our methodological pipeline.




Figure 1. Workflow of the proposed model for odor presence detection from OB LFPs. Raw recordings undergo bandpass filtering, downsampling, and normalization before training complementary CNNs. Ensemble fusion enables binary odor classification and performance evaluation.

2.1. Data Source and Experimental Context

Our analysis uses the *pcx-1* dataset [26,27] prepared by Bolding and Franks (2018), which is publicly available on CRCNS.org (<https://doi.org/10.6080/K00C4SZB>). It contains simultaneous extracellular recordings of mouse OB and PCx, acquired with 32-channel NeuroNexus Poly3 silicon probes (NeuroNexus, Ann Arbor, MI, USA) at a sampling rate of 30 kHz, together with respiration traces sampled at 2 kHz. Although both OB and PCx signals are provided, in this study we analyze only the OB recordings.

The dataset comprises two complementary recording sets. The main dataset contains 2349 trials from seven awake head-fixed mice, with odor stimuli (ethyl butyrate, isoamyl acetate, 2-hexanone, hexanal, ethyl acetate, and ethyl tiglate), each delivered at 0.3% *v/v* concentration; mineral oil served as the baseline control ($n = 336$ trials). A supplementary concentration dataset contains 2600 trials examining ethyl butyrate across four concentrations (0.03%, 0.1%, 0.3%, and 1.0% *v/v*; $n = 200$ trials per concentration) with matched mineral oil controls ($n = 200$ trials), to enable a systematic analysis of concentration-dependent detection thresholds. All analyses were implemented in Python (version 3.12.12) using NumPy (version 2.0.2) for numerical computations.

2.2. Pre-Processing

2.2.1. Dataset Preparation and Class Balancing

A systematic pre-processing pipeline was implemented to prepare the raw LFP signals for analysis. The ‘Odor-Presence’ class was established from all active odorant trials ($n = 2013$) while the Mineral Oil controls ($n = 336$) defined the ‘Odor-Absence’ class. This grouping resulted in a significant 6:1 data imbalance. To mitigate the risk of classification bias, we applied random under-sampling (without replacement) to the ‘Odor-Presence’ group, selecting 336 trials to match the minority class. This procedure yielded a final balanced dataset composed of 672 trials for the binary odor detection task.

2.2.2. Signal Filtering and Downsampling

Every 32 recording channels were subjected to a fifth-order Butterworth bandpass filter (SciPy version 1.16.3) with cut-off points at 0.5 Hz and 100 Hz. These cut-off points were chosen to allow odor-relevant frequencies in the delta to gamma range to be preserved and to eliminate baseline drift and high-frequency artifacts [14,28,29]. After filtering, the data were downsampled at a rate of 30 to 1 kHz, resulting in 2000 data samples per 2-s trial with no aliasing artifacts (since the cut-off is well below the Nyquist frequency at 500 Hz post-downsampling).

2.2.3. Spectral Feature Extraction and Normalization

Following the pre-processing, spectral features were extracted and normalized. First, power spectral densities (PSDs) were computed for each channel using Welch’s method with a 256-point Hann window and 50.0% overlap [30]. The resulting spectra were then normalized using a RobustScaler (scikit-learn version 1.6.1) by subtracting the median and dividing it by the inter-quartile range to reduce the impact of spectral outliers [31].

2.3. Network Architecture

AttentionCNN: AttentionCNN is a well-known architecture for processing complex feature sets. Given the richness and variability of odor-evoked spectral activity, we chose this model to focus on the most discriminative temporal patterns. Our implementation begins with a stand-alone max-pooling layer (stride 2), which is followed by two identical convolutional blocks. Each block consists of the sequence [Conv1D → BatchNorm → ReLU → MaxPool (stride 4)], and these operations collectively expand the output to

128 channels. We then apply three parallel Conv1D branches (kernel sizes 1, 3, 5; 64 filters each), concatenate to form 192 channels, and then recalibrate via a squeeze-and-excitation channel attention module [32], followed by spatial attention.

For a given feature map $\mathbf{U} \in \mathbb{R}^{C \times T}$, the squeeze-and-excitation operation performs the following:

$$z_c = \frac{1}{T} \sum_{t=1}^T u_c(t) \quad (\text{global average pooling}), \quad (1)$$

$$\mathbf{s} = \sigma(\mathbf{W}_2 \delta(\mathbf{W}_1 \mathbf{z})) \quad (\text{excitation}), \quad (2)$$

$$\tilde{\mathbf{U}}_c = s_c \cdot \mathbf{U}_c \quad (\text{recalibration}), \quad (3)$$

where σ is the sigmoid function, δ is ReLU, and $\mathbf{W}_1 \in \mathbb{R}^{C/r \times C}$, $\mathbf{W}_2 \in \mathbb{R}^{C \times C/r}$ with reduction ratio $r = 16$.

A global average pooling layer reduces the temporal dimension, producing a 192-dimensional vector that passes through dropout ($p = 0.3$), a 256-unit ReLU fully connected layer, dropout ($p = 0.5$), and a final linear classifier for binary odor detection.

ResCNN: ResCNN builds on residual blocks to enable very deep networks, improving the gradient flow and feature reuse. This makes it ideal for capturing hierarchical temporal features in LFP spectra. Our ResCNN adaptation starts with max-pooling (stride 2), a Conv1D layer (kernel 7, stride 2) with BatchNorm and ReLU, and another max-pool (stride 4). Three residual blocks (64 channels each) then refine the features, followed by a Conv1D (kernel 3, stride 1 $\rightarrow$ BatchNorm $\rightarrow$ ReLU) to expand to 128 channels. After a final max-pool (stride 2) and two more 128-channel residual blocks, global average pooling condenses the temporal dimension. A dropout layer ($p = 0.4$) precedes the final linear classifier, producing robust single-trial odor-presence decisions.

2.4. Rationale for Architecture Selection

To justify our architecture choices and ensure reproducibility, we compared eight CNN variants on our dataset: (1) Vanilla CNN (3-layer baseline), (2) Deep CNN (6 layers without skip connections), (3) Dilated CNN (dilated convolutions), (4) Wide CNN (increased channel width), (5) Shallow CNN (2 layers), (6) AttentionCNN (our proposed architecture with CBAM attention), (7) ResCNN (our proposed architecture with squeeze-and-excitation residual blocks), and (8) Ensemble (combination of AttentionCNN and ResCNN). All models were trained using identical protocols: 5-fold cross-validation, 150 epochs, AdamW optimizer with cosine annealing, and fixed random seed (42) to have a fair comparison [33–35]. The comprehensive performance comparison is presented in Section 3.1.

2.5. Training Procedure

To train and evaluate our models, we employed a five-fold cross-validation scheme on the 2349 trials. For each fold, the data were partitioned into a training set (80.0%) and a test set (20.0%), with 10% of the training set being used for the validation. We optimized network parameters using the AdamW optimizer (PyTorch version 2.9.0; initial learning rate 5×10^{-4} , weight decay 1×10^{-4}) [36]. All GPU computations were performed on NVIDIA A100-SXM4-80GB (NVIDIA Corporation, Santa Clara, CA, USA) with CUDA (version 12.6). The learning rate schedule was varied depending on the model being evaluated: for ResCNN and AttentionCNN, the learning rate followed a Cosine Warm Restarts schedule ($T_0 = 10$ epochs, $T_{\text{mult}} = 2$). For the ensemble evaluation, the models within each fold were trained using a One-Cycle policy [37]. Early stopping was applied in each run, with a patience of 15–20 epochs and a minimum required improvement (Δ) of

0.001 in the validation loss. The model checkpoint with the best validation loss was then used for the final testing.

2.6. Ensemble Strategy

To combine the complementary outputs of AttentionCNN and ResCNN without additional training, we employed a late fusion of their softmax probabilities. For a given trial x , let

$$p_{\text{Res}}(x) = \text{softmax}(m_{\text{Res}}(x)), \quad (4)$$

$$p_{\text{Att}}(x) = \text{softmax}(m_{\text{Att}}(x)). \quad (5)$$

We compute the ensemble probability as the arithmetic mean,

$$p_{\text{ens}}(x) = \frac{p_{\text{Res}}(x) + p_{\text{Att}}(x)}{2}, \quad (6)$$

and assign an “odor” label if $p_{\text{ens}}(x) > 0.5$. This simple fusion does not require additional parameters and incurs minimal inference overhead. Under five-fold cross-validation, this probability averaging was performed within each fold: For each of the five splits, an AttentionCNN and a ResCNN were trained on the training portion, and their ensembled predictions were evaluated on the held-out test portion. The final reported metrics are the mean and standard deviation of the performance across these five folds.

2.7. Evaluation Metrics

We evaluated the classification performance using the following metrics:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (7)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (8)$$

$$\text{Recall (Sensitivity)} = \frac{TP}{TP + FN}, \quad (9)$$

$$\text{Specificity} = \frac{TN}{TN + FP}, \quad (10)$$

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}, \quad (11)$$

where TP , TN , FP , and FN denote true positives, true negatives, false positives, and false negatives, respectively. Additionally, we computed the Area Under the Receiver Operating Characteristic Curve (AUC) and used confusion matrices to clarify detailed error distributions. All metrics were calculated within each fold of a five-fold cross-validation and reported as the mean $\pm$ standard deviation. Furthermore, to verify the reliability of the probability estimates, we compared the model output with the actual labels for each test trial to assess the calibration (for example, a predicted odor probability $\sim 80.0\%$ typically indicated actual odor trials $\sim 80.0\%$).

3. Results

3.1. Architecture Comparison and Selection

Table 1 presents the systematic comparison of eight CNN architectures, justifying our selection of AttentionCNN and ResCNN for the ensemble model. The baseline Vanilla CNN (three convolutional layers) achieved 83.9% accuracy, establishing our performance floor. Simply adding depth without architectural innovations proved ineffective—the Deep CNN (six layers without skip connections) matched this performance at 83.9%, likely due

to gradient degradation. Dilated convolutions (Dilated CNN: 83.0%) and increased channel width (Wide CNN: 82.1%) similarly failed to improve upon the baseline, while the Shallow CNN (79.9%) confirmed that insufficient model capacity limits the performance.

In contrast, our proposed architectures demonstrated clear advantages. ResCNN with squeeze-and-excitation residual blocks reached 85.6%, demonstrating the value of skip connections and channel attention for multi-electrode LFP classification. AttentionCNN with CBAM attention achieved 84.7%, learning complementary representations through its multi-scale convolutional branches. Combining both architectures via ensemble averaging yielded the highest performance (86.2% accuracy, 0.942 AUC), a 2.3 percentage point improvement over the baseline, validating our design choices.

Table 1. Performance comparison of CNN architectures for binary odor detection. All models are trained with identical protocols (five-fold CV, 150 epochs, AdamW optimizer, seed = 42).

Architecture	Accuracy (%)	F1 (%)	AUC
Ensemble	86.2 ± 2.8	85.3 ± 3.4	0.942 ± 0.011
ResCNN	85.6 ± 2.8	85.0 ± 3.5	0.931 ± 0.021
AttentionCNN	84.7 ± 2.2	83.6 ± 2.4	0.921 ± 0.016
Vanilla CNN	83.9 ± 2.4	83.2 ± 3.0	0.911 ± 0.021
Deep CNN	83.9 ± 1.4	83.3 ± 3.0	0.922 ± 0.016
Dilated CNN	83.0 ± 4.1	83.1 ± 4.6	0.912 ± 0.034
Wide CNN	82.1 ± 1.8	81.1 ± 2.9	0.906 ± 0.022
Shallow CNN	79.9 ± 2.7	77.4 ± 3.7	0.875 ± 0.015

Bold entries indicate the three primary architectures proposed in this study.

3.2. Performance Metrics Distribution Across Models

Table 2 summarizes the comprehensive evaluation metrics for our three primary models across five-fold cross-validation. AttentionCNN achieved 84.7% accuracy with high specificity (92.2%), while ResCNN reached 85.6% accuracy with more balanced sensitivity (87.0%) and specificity (86.0%). The ensemble combined these complementary strengths, achieving 86.2% accuracy with 84.0% sensitivity and 90.0% specificity. These results indicate that (i) both architectural innovations yield strong single-trial detection, and (ii) ensemble fusion maintains peak accuracy while balancing the sensitivity and specificity for robust odor detection.

Table 2. Comprehensive performance metrics for AttentionCNN, ResCNN, and ensemble models (mean ± SD over five folds). Abbreviations: Acc. = accuracy; F1 = F1 score; AUC = area under the receiver operating characteristic curve; Sens. = sensitivity; Spec. = specificity.

Model	Acc. (%)	F1 (%)	AUC	Sens. (%)	Spec. (%)
AttentionCNN	84.7 ± 2.2	83.6 ± 2.4	0.921 ± 0.016	76.0 ± 2.5	92.2 ± 1.8
ResCNN	85.6 ± 2.8	85.0 ± 3.5	0.931 ± 0.021	87.0 ± 1.9	86.0 ± 2.2
Ensemble	86.2 ± 2.8	85.3 ± 3.4	0.942 ± 0.011	84.0 ± 2.0	90.0 ± 1.7

3.3. Learned Feature Representation

Beyond the quantitative metrics, we validated that our model learned biologically significant features. Figure 2 visualizes the feature space learned by the network using t-SNE [38]. The per-odor colored embedding (Figure 2) achieves a silhouette score of 0.544. Mineral oil trials cluster distinctly from odor trials, while the relative positions of individual odorants correlate with their classification accuracy—hexanal and 2-hexanone form tight clusters far from the control region, whereas ethyl acetate trials appear closer to the boundary, consistent with its lower detectability (detailed in Section 3.4). The 3D projections provide additional perspective, clearly showing the spatial separation between the control and odor-evoked neural patterns.

The partial overlap seen in the t-SNE embedding reflects differences in how well the neural responses to different odorants can be distinguished from the baseline. Our per-odor analysis (Table 3) shows that odorants producing weaker neural signatures, such as ethyl acetate (67.1% accuracy), account for most of the overlap region. In contrast, odorants like hexanal (96.3% accuracy) form tight well-separated clusters. This pattern is consistent with what we know about glomerular activation [39], where some odorant molecules produce more distinct responses than others.




Figure 2. t-SNE visualization of ensemble feature embeddings colored by odor identity. Left: 2D projection with density contours (silhouette score = 0.544). Right: 3D projection showing spatial relationships between odor clusters. Mineral oil (control) forms a distinct cluster separated from the various odorants, with some overlap in the boundary region corresponding to harder-to-classify trials.

Table 3. Per-odor classification performance using the ensemble model. Each row shows binary classification (specific odorant vs. mineral oil) using the same five-fold cross-validation protocol. All odorants tested at 0.3% *v/v* concentration.

Odorant	Trials	Accuracy (%)	Interpretation
Hexanal	337	96.3 ± 1.5	Excellent separability
2-Hexanone	335	95.6 ± 2.6	Excellent separability
Ethyl butyrate	335	91.5 ± 2.7	High separability
Isoamyl acetate	336	90.6 ± 2.8	High separability
Ethyl tiglate	335	85.5 ± 2.4	Moderate separability
Ethyl acetate	335	67.1 ± 5.0	Low separability

3.4. Per-Odor Classification Performance

When we trained separate classifiers for each odorant against the mineral oil control, we found substantial differences in how well each could be detected (Table 3, Figure 3). Hexanal and 2-hexanone both achieved over 95% accuracy, suggesting these compounds produce particularly distinctive neural responses [40]. Ethyl acetate, on the other hand, proved much harder to detect at only 67.1% accuracy. This 29.2 percentage point range tells us that the olfactory bulb does not encode all odors with equal clarity. This pattern indicates that aldehydes (hexanal) and ketones (2-hexanone) show higher detectability compared to esters (ethyl acetate, ethyl tiglate), likely due to their distinct molecular properties. Factors like receptor binding characteristics and the spatial pattern of glomerular activation [41,42] determine how discriminable each odor's neural signature is from the baseline.




Figure 3. Per-odor classification accuracy for binary detection (odor vs. mineral oil) shown as a heatmap. Numbers indicate mean accuracy across five-fold cross-validation. Aldehydes and ketones show higher detectability compared to esters.

3.5. Concentration-Dependent Detection

We tested whether our detection framework has a concentration threshold by analyzing ethyl butyrate across four concentration levels spanning two orders of magnitude (Table 4). At 0.03% *v/v*, the classification was no better than a binary random chance (51.3%), indicating the neural response at this concentration is too weak to detect [43,44]. The performance rose to 60.8% at 0.1% *v/v* and reached 75.2% at 0.3% *v/v*, which we take as the practical detection threshold. At the highest concentration tested (1.0% *v/v*), the accuracy reached 86.4%. This monotonic increase in accuracy with concentration mirrors the dose–response characteristics of sensory neurons [45,46], suggesting our classifier is tracking the strength of the neural response rather than picking up on artifacts.

Table 4. Concentration-dependent classification performance for ethyl butyrate. Statistical significance tested against chance level (50%) using one-sample *t*-test. The detection threshold is defined as the lowest concentration achieving both $p < 0.001$ and accuracy $> 70\%$.

Concentration (% <i>v/v</i>)	Accuracy (%)	vs. Chance	Detection Status
0.03	51.3 ± 5.0	$p > 0.05$	At chance level
0.10	60.8 ± 5.6	$p < 0.05$	Marginal detection
0.30	75.2 ± 2.7	$p < 0.001$	Reliable detection
1.00	86.4 ± 3.4	$p < 0.001$	Robust detection

3.6. Inference Speed Analysis

To assess the real-time capability, we benchmarked the inference speed across a CPU and three GPU platforms (Table 5). The lightweight ensemble architecture (461.4 K parameters) achieves single-sample inference in just 2.58 ms on a CPU—faster than the GPU, due to the kernel launch overhead at small batch sizes. Given the 1000 ms trial duration, this corresponds to a real-time factor of 388×, enabling deployment on standard hardware without GPU acceleration. For batch processing of large datasets, GPU throughput reaches 11,780 samples/second at batch size 64, facilitating rapid offline analysis.

Table 5. Inference speed benchmarking across different hardware platforms. Latency represents single-sample (batch size = 1) processing time. All measurements are mean $\pm$ standard deviation over 200 runs after 50 warm-up iterations. Model sizes: AttentionCNN (155.1 K params, 0.6 MB), ResCNN (306.3 K params, 1.2 MB), ensemble (461.4 K params, 1.8 MB).

Model	Hardware	Latency (ms)	Throughput (Samples/s)	Memory (MB)
AttentionCNN	CPU	1.03 $\pm$ 0.05	970	–
	Tesla T4	1.86 $\pm$ 0.10	539	19
	NVIDIA L4	1.77 $\pm$ 0.02	564	19
	A100-80GB	1.76 $\pm$ 0.04	569	19
ResCNN	CPU	1.75 $\pm$ 0.01	571	–
	Tesla T4	3.86 $\pm$ 0.05	259	16
	NVIDIA L4	4.18 $\pm$ 0.15	239	16
	A100-80GB	3.86 $\pm$ 0.01	259	16
Ensemble	CPU	2.58 $\pm$ 0.01	387	–
	Tesla T4	5.52 $\pm$ 0.04	181	22
	NVIDIA L4	5.89 $\pm$ 0.06	170	22
	A100-80GB	5.56 $\pm$ 0.07	180	22

3.7. Ensemble Prediction Confidence

Figure 4 displays the distribution of the ensemble prediction confidence for correct and incorrect classifications. The horizontal axis shows the confidence level, and the vertical axis indicates the frequency of trials within each confidence bin. The green bars correspond to correct predictions with higher confidence values (mean confidence $\approx$ 0.815), while the red bars (incorrect predictions) are distributed more broadly with lower mean confidence ($\approx$ 0.678). This suggests that the ensemble classifier is well calibrated, as it expresses higher confidence on correct predictions and lower confidence on misclassifications.




Figure 4. Histogram of ensemble prediction confidence for individual test trials, with correct predictions shown in green and incorrect predictions in red.

4. Discussion

4.1. Overcoming Prior Limitations

The findings presented in this study offer two primary contributions to the field of olfactory decoding. The first contribution is related to methodological advance: we establish the feasibility of accurate single-trial odor detection using deep learning on spectral features. This validates our first hypothesis that these signals contain sufficient

information for robust classification without averaging. The second contribution is related to neurological insight: we demonstrate that this performance can be achieved using signals from the olfactory bulb alone. This supports our second hypothesis that the initial stages of olfactory processing are sufficient for the fundamental task of presence detection for seven odors, without requiring contributions from higher cortical regions.

Prior non-invasive work on human olfactory registration has shown lower performance. Rajabi et al. evaluated logistic regression and an end-to-end 1D ResNet on EEG and EBG signals, finding that their linear baseline remained at $AUC \approx 50\%$, and their ResNet-1D could push AUC into the high-50% range (e.g., 56.6% for scalp-EBG, 58.0% for EEG) [23]. The difficulty of cross-subject generalization with EEG is further shown by the work of Ezzatdoost et al., who achieved 64.3% accuracy on a more complex four-odor identification task using handcrafted nonlinear features [19]. Similarly, Kato et al. demonstrated that, while odor representations in human EEG can be decoded within 100 ms post-stimulus onset, achieving reliable classification remains challenging due to signal quality limitations [47]. This demonstrates that current non-invasive recordings lack sufficient SNR for accurate single-trial odor detection [18,48] and highlights the advantage of LFP signals.

Table 6 places our results within the larger context of neural odor decoding. Although direct quantitative comparison is limited by differences in recording modality or experimental procedure, the substantial performance gap corroborates our methodological assumption that high SNR neural recordings are a prerequisite for robust single-trial odor detection. The 22–29 percentage point accuracy gain observed is consistent with the recognized difference in signal fidelity between intracranial LFP recordings versus non-invasive scalp modalities such as EEG or EBG [13], definitively validating our methodological assumptions. Finally, we acknowledge that the performance differential may be influenced by interspecific differences in olfactory processing between mice and humans.

Table 6. Performance comparison across recording modalities and species. Note: Direct comparison is limited due to differences in species (human vs. mouse), recording invasiveness (scalp vs. intracranial), and experimental protocols.

Study	Species	Signal Type	Accuracy (%)	Task
Rajabi et al. [23]	Human	EEG (non-inv.)	58.0	Odor vs. Blank
Rajabi et al. [23]	Human	EBG (non-inv.)	57.0	Odor vs. Blank
Ezzatdoost et al. [19]	Human	EEG (non-inv.)	64.3	4-Odor Identity
Our Work	Mouse	OB-LFP (inv.)	86.2	Odor vs. Blank

Bold entry indicates the results from the present study.

4.2. Methodological Contributions

This work establishes a reliable neural-based binary odor detection system, providing both methodological advances and biological insights. Our ensemble approach effectively combines complementary CNN architectures to achieve reliable performance. The clear separation observed in our t-SNE visualization (Figure 2) and the high confidence levels for correct predictions (Figure 4) further validate that our models have learned biologically meaningful features that capture fundamental differences between odor-presence and odor-absence neural states. Compared to prior EEG-based studies, our LFP approach offers higher spatial resolution and direct access to OB circuitry, enabling the extraction of rich spectro-temporal signatures that were previously inaccessible.

4.3. Limitations and Future Directions

Our concentration analysis (Section 3.5) shows that the detection performance depends on the stimulus strength, with a practical threshold at 0.3% *v/v* for ethyl butyrate. Below this concentration, neural responses become too weak for reliable single-trial classification.

This finding has important implications for real-world applications, where target analyses may be present at trace concentrations. Future work should systematically characterize detection thresholds across all six odorants and investigate whether multi-trial averaging or more sensitive pre-processing methods can lower these limits.

Our evaluation focuses on monomolecular odorants at a fixed concentration (0.3% *v/v*) in head-fixed mice, leaving important questions about generalization to complex mixtures [49], varying concentrations, and naturalistic behavioral conditions. The invasive nature of LFP recordings also limits immediate translational applications, though our findings provide crucial validation of neural-based odor detection principles. Future work will systematically extend our approach to diverse odor mixtures and concentration ranges, validate its performance in freely moving animals, and investigate non-invasive recording modalities to assess translational feasibility.

5. Conclusions

This study demonstrates that deep neural networks can achieve robust single-trial odor detection from olfactory bulb LFPs, establishing two key findings. First, spectral features within single-trial LFPs provide sufficient information for accurate odor presence classification, achieving 86.2% accuracy and 0.942 AUC—substantially outperforming prior non-invasive approaches. Second, signals from the olfactory bulb alone are adequate for this task, without requiring downstream cortical processing. Our ensemble of AttentionCNN and ResCNN architectures leverages complementary feature extraction strategies to capture biologically significant neural signatures, as confirmed by t-SNE visualization and prediction confidence analysis. These findings establish the feasibility of LFP-based odor sensing and provide a foundation for future development of neural-based detection systems.

Author Contributions: Conceptualization, M.K.Ö.; methodology, M.H.; software, M.H.; validation, M.H., A.Z., and M.K.Ö.; formal analysis, M.H.; investigation, M.H.; resources, M.K.Ö.; data curation, M.H.; writing—main text, M.H. and A.Z.; performing simulations, M.H. and A.Z.; idea formation and manuscript editing, M.K.Ö.; visualization, M.H.; supervision, M.K.Ö.; project administration, M.K.Ö.; funding acquisition, M.K.Ö. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Authority TUBITAK, grant number 123E520.

Institutional Review Board Statement: The study was conducted using publicly available data. Ethical review and approval were obtained by the original data collectors as described in [26,27].

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used in this study are publicly available at CRCNS.org (<https://doi.org/10.6080/K00C4SZB>).

Acknowledgments: We thank Kevin A. Bolding and Kevin M. Franks for making their dataset publicly available on CRCNS.org. The simulation of this research was performed over NVIDIA A100 GPUs via NVIDIA's academic grant program.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

- 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* **2025**, *25*, 4437. [\[CrossRef\]](#) [\[PubMed\]](#)
- Kim, C.; Lee, K.K.; Kang, M.S.; Shin, D.M.; Oh, J.W.; Lee, C.S.; Han, D.W. Artificial olfactory sensor technology that mimics the olfactory mechanism: A comprehensive review. *Biomater. Res.* **2022**, *26*, 40. [\[CrossRef\]](#)
- Deng, H.; Chen, Z.; Feng, P.; Tian, L.; Zong, H.; Nakamoto, T. Recent Advances and Applications of Odor Biosensors. *Electronics* **2025**, *14*, 1852. [\[CrossRef\]](#)
- Dennler, N.; Drix, D.; Warner, T.P.A.; Rastogi, S.; Della Casa, C.; Ackels, T.; Schaefer, A.T.; van Schaik, A.; Schmuker, M. High-speed odor sensing using miniaturized electronic nose. *Sci. Adv.* **2024**, *10*, eadp1764. [\[CrossRef\]](#) [\[PubMed\]](#)
- Kim, T.; Kim, Y.; Cho, W.; Kwak, J.-H.; Cho, J.; Pyeon, Y.; Kim, J.J.; Shin, H. Ultralow-power single-sensor-based e-nose system powered by duty cycling and deep learning for real-time gas identification. *ACS Sens.* **2024**, *9*, 3557–3572. [\[CrossRef\]](#) [\[PubMed\]](#)
- Shor, E.; Herrero-Vidal, P.; Dewan, A.; Uguz, I.; Curto, V.F.; Malliaras, G.G.; Savin, C.; Bozza, T.; Rinberg, D. Sensitive and robust chemical detection using an olfactory brain-computer interface. *Biosens. Bioelectron.* **2022**, *195*, 113664. [\[CrossRef\]](#)
- Lu, Q.; Yi, M.; Jiang, J. Bioelectronic nose for ultratrace odor detection via brain-computer interface with olfactory bulb electrode arrays. *Biosens. Bioelectron.* **2025**, *285*, 117585. [\[CrossRef\]](#)
- Qin, C.; Wang, Y.; Hu, J.; Wang, T.; Liu, D.; Dong, J.; Lu, Y. Artificial Olfactory Biohybrid System: An Evolving Sense of Smell. *Adv. Sci.* **2023**, *10*, 2204726. [\[CrossRef\]](#)
- Morozova, M.; Bikbavova, A.; Bulanov, V.; Lebedev, M.A. An olfactory-based brain-computer interface: Electroencephalography changes during odor perception and discrimination. *Front. Behav. Neurosci.* **2023**, *17*, 1122849. [\[CrossRef\]](#)
- Ninenko, I.; Medvedeva, A.; Efimova, V.L.; Kleeva, D.F.; Morozova, M.; Lebedev, M.A. Olfactory neurofeedback: Current state and possibilities for further development. *Front. Hum. Neurosci.* **2024**, *18*, 1419552. [\[CrossRef\]](#)
- Zhu, P.; Liu, S.; Tian, Y.; Chen, Y.; Chen, W.; Wang, P.; Du, L.; Wu, C. In vivo bioelectronic nose based on a bioengineered rat realizes the detection and classification of multiodorants. *ACS Chem. Neurosci.* **2022**, *13*, 1727–1737. [\[CrossRef\]](#)
- Einevoll, G.T.; Kayser, C.; Logothetis, N.K.; Panzeri, S. Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nat. Rev. Neurosci.* **2013**, *14*, 770–785. [\[CrossRef\]](#)
- Buzsáki, G.; Anastassiou, C.A.; Koch, C. The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.* **2012**, *13*, 407–420. [\[CrossRef\]](#)
- Pesaran, B.; Vinck, M.; Einevoll, G.T.; Sirota, A.; Fries, P.; Siegel, M.; Truccolo, W.; Schroeder, C.E.; Srinivasan, R. Investigating large-scale brain dynamics using field potential recordings: Analysis and interpretation. *Nat. Neurosci.* **2018**, *21*, 903–919. [\[CrossRef\]](#)
- Yang, Q.; Zhou, G.; Noto, T.; Templer, J.W.; Schuele, S.U.; Rosenow, J.M.; Lane, G.; Zelano, C. Smell-induced gamma oscillations in human olfactory cortex are required for accurate perception of odor identity. *PLoS Biol.* **2022**, *20*, e3001509. [\[CrossRef\]](#)
- Arpaia, P.; Cataldo, A.; Criscuolo, S.; De Benedetto, E.; Masciullo, A.; Schiavoni, R. Assessment and Scientific Progresses in the Analysis of Olfactory Evoked Potentials. *Bioengineering* **2022**, *9*, 252. [\[CrossRef\]](#)
- Ninenko, I.; Kleeva, D.F.; Bukreev, N.; Lebedev, M.A. An experimental paradigm for studying EEG correlates of olfactory discrimination. *Front. Hum. Neurosci.* **2023**, *17*, 1117801. [\[CrossRef\]](#)
- Iravani, B.; Arshamian, A.; Ohla, K.; Wilson, D.A.; Lundström, J.N. Non-invasive recording from the human olfactory bulb. *Nat. Commun.* **2020**, *11*, 648. [\[CrossRef\]](#)
- Ezzatdoost, K.; Hojjati, H.; Aghajani, H. Decoding olfactory stimuli in EEG data using nonlinear features: A pilot study. *J. Neurosci. Methods* **2020**, *341*, 108780. [\[CrossRef\]](#)
- Abbasi, N.I.; Bose, R.; Bezerianos, A.; Thakor, N.V.; Dragomir, A. EEG-based classification of olfactory response to pleasant stimuli. In Proceedings of the 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Berlin, Germany, 23–27 July 2019; IEEE: New York, NY, USA, 2019; pp. 5160–5163. [\[CrossRef\]](#)
- Hou, H.R.; Han, R.X.; Zhang, X.N.; Meng, Q.H. Pleasantness Recognition Induced by Different Odor Concentrations Using Olfactory Electroencephalogram Signals. *Sensors* **2022**, *22*, 8808. [\[CrossRef\]](#)
- Huart, C.; Legrain, V.; Hummel, T.; Rombaux, P.; Mouraux, A. Time-Frequency Analysis of Chemosensory Event-Related Potentials to Characterize the Cortical Representation of Odors in Humans. *PLoS ONE* **2012**, *7*, e33221. [\[CrossRef\]](#) [\[PubMed\]](#)
- Rajabi, N.; Zanettin, I.; Ribeiro, A.H.; Vasco, M.; Björkman, M.; Lundström, J.N.; Kragic, D. Exploring the feasibility of olfactory brain-computer interfaces. *Sci. Rep.* **2025**, *15*, 18404. [\[CrossRef\]](#) [\[PubMed\]](#)
- Livezey, J.A.; Glaser, J.I. Deep learning approaches for neural decoding across architectures and recording modalities. *Brief. Bioinform.* **2021**, *22*, 1577–1591. [\[CrossRef\]](#) [\[PubMed\]](#)
- Hossain, K.M.; Islam, M.A.; Hossain, S.; Nijholt, A.; Ahad, M.A.R. Status of deep learning for EEG-based brain-computer interface applications. *Front. Comput. Neurosci.* **2022**, *16*, 1006763. [\[CrossRef\]](#)
- Bolding, K.A.; Franks, K.M. Recurrent cortical circuits implement concentration-invariant odor coding. *Science* **2018**, *361*, eaat6904. [\[CrossRef\]](#)

27. Bolding, K.A.; Franks, K.M. Simultaneous Extracellular Recordings from the Mouse Olfactory Bulb and Piriform Cortex in Response to Odor Stimuli. CRCNS.org. 2018. Available online: <https://doi.org/10.6080/K00C4SZB> (accessed on 1 January 2025).
28. Lepousez, G.; Lledo, P.M. Odor Discrimination Requires Proper Olfactory Fast Oscillations in Awake Mice. *Neuron* **2013**, *80*, 1010–1024. [\[CrossRef\]](#)
29. Kay, L.M. Two Species of Gamma Oscillations in the Olfactory Bulb: Dependence on Behavioral State and Synaptic Interactions. *J. Integr. Neurosci.* **2003**, *2*, 31–44. [\[CrossRef\]](#)
30. Welch, P.D. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* **1967**, *15*, 70–73. [\[CrossRef\]](#)
31. Rousseeuw, P.J.; Croux, C. Alternatives to the Median Absolute Deviation. *J. Am. Stat. Assoc.* **1993**, *88*, 1273–1283. [\[CrossRef\]](#)
32. Hu, J.; Shen, L.; Sun, G. Squeeze-and-Excitation Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–23 June 2018; IEEE: New York, NY, USA, 2018; pp. 7132–7141. [\[CrossRef\]](#)
33. Lawhern, V.J.; Solon, A.J.; Waytowich, N.R.; Gordon, S.M.; Hung, C.P.; Lance, B.J. EEGNet: A compact convolutional neural network for EEG-based brain-computer interfaces. *J. Neural Eng.* **2018**, *15*, 056013. [\[CrossRef\]](#)
34. Schirrmester, R.T.; Springenberg, J.T.; Fiederer, L.D.J.; Glasstetter, M.; Eggensperger, K.; Tangermann, M.; Hutter, F.; Burgard, W.; Ball, T. Deep learning with convolutional neural networks for EEG decoding and visualization. *Hum. Brain Mapp.* **2017**, *38*, 5391–5420. [\[CrossRef\]](#) [\[PubMed\]](#)
35. Walther, D.; Viehweg, J.; Haueisen, J.; Mäder, P. A systematic comparison of deep learning methods for EEG time series analysis. *Front. Neuroinform.* **2023**, *17*, 1067095. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the 3rd International Conference on Learning Representations (ICLR), San Diego, CA, USA, 7–9 May 2015. Available online: <https://arxiv.org/abs/1412.6980> (accessed on 1 January 2026).
37. Smith, L.N. A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay. *arXiv* **2018**, arXiv:1803.09820. [\[CrossRef\]](#)
38. van der Maaten, L.; Hinton, G. Visualizing Data using t-SNE. *J. Mach. Learn. Res.* **2008**, *9*, 2579–2605. Available online: <https://www.jmlr.org/papers/v9/vandemaaten08a.html> (accessed on 1 January 2026).
39. Burton, S.D.; Brown, A.; Eiting, T.P.; Youngstrom, I.A.; Rust, T.C.; Schmuker, M.; Wachowiak, M. Mapping odorant sensitivities reveals a sparse but structured representation of olfactory chemical space by sensory input to the mouse olfactory bulb. *eLife* **2022**, *11*, e80470. [\[CrossRef\]](#)
40. Shani-Narkiss, H.; Beniaguev, D.; Segev, I.; Mizrahi, A. Stability and flexibility of odor representations in the mouse olfactory bulb. *Front. Neural Circuits* **2023**, *17*, 1157259. [\[CrossRef\]](#)
41. Bhattacharjee, A.S.; Konakamchi, S.; Turaev, D.; Vincis, R.; Nunes, D.; Dingankar, A.A.; Spors, H.; Carleton, A.; Kuner, T.; Abraham, N.M. Similarity and strength of glomerular odor representations define a neural metric of sniff-invariant discrimination time. *Cell Rep.* **2019**, *28*, 2966–2978.e5. [\[CrossRef\]](#)
42. Economo, M.N.; Hansen, K.R.; Wachowiak, M. Control of mitral/tufted cell output by selective inhibition among olfactory bulb glomeruli. *Neuron* **2016**, *91*, 397–411. [\[CrossRef\]](#)
43. Parabucki, A.; Bizer, A.; Morris, G.; Munoz, A.E.; Bala, A.D.S.; Smear, M.; Shusterman, R. Odor concentration change coding in the olfactory bulb. *eNeuro* **2019**, *6*, ENEURO.0396-18.2019. [\[CrossRef\]](#)
44. Bolding, K.A.; Franks, K.M. Complementary codes for odor identity and intensity in olfactory cortex. *eLife* **2017**, *6*, e22630. [\[CrossRef\]](#)
45. Rolls, E.T.; Baker, K.L.; Bhattacharjee, A.S.; Verhagen, J.V. Odor encoding by signals in the olfactory bulb. *J. Neurophysiol.* **2023**, *129*, 292–313. [\[CrossRef\]](#)
46. Shusterman, R.; Sirotin, Y.B.; Smear, M.C.; Ahmadian, Y.; Rinberg, D. Sniff invariant odor coding. *eNeuro* **2018**, *5*, ENEURO.0149-18.2018. [\[CrossRef\]](#)
47. Kato, M.; Okumura, T.; Tsubo, Y.; Honda, J.; Sugiyama, M.; Touhara, K.; Okamoto, M. Spatiotemporal dynamics of odor representations in the human brain revealed by EEG decoding. *Proc. Natl. Acad. Sci. USA* **2022**, *119*, e2114966119. [\[CrossRef\]](#)
48. Iravani, B.; Schaefer, M.; Wilson, D.A.; Arshamian, A.; Lundström, J.N. The human olfactory bulb processes odor valence representation and cues motor avoidance behavior. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, e2101209118. [\[CrossRef\]](#)
49. Zak, J.D.; Reddy, G.; Vergassola, M.; Murthy, V.N. Distinct information conveyed to the olfactory bulb by feedforward input from the nose and feedback from the cortex. *Nat. Commun.* **2024**, *15*, 3268. [\[CrossRef\]](#)

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.