TOWARDS FASTER MACHINE OLFACTION: ADVANCES IN HIGH-SPEED SENSING AND PROCESSING

by
Nik Dennler

A dissertation submitted to
University of Hertfordshire and Western Sydney University
in partial fulfilment of the requirements for the degree of Doctor of Philosophy (PhD)

Sydney, Australia
July 2024

Abstract

This thesis investigates the principles and methods for achieving high temporal resolution in odour sampling and processing for machine perception. Initially, a comprehensive review of fast olfaction mechanisms in insects and mammals is provided, highlighting the necessity of rapid sensing when exposed to the dynamics of turbulent odour plumes. The discussion extends to artificial olfaction technologies, emphasising current advancements in electronic nose (e-nose) systems and their applications. Initial evaluations of existing datasets and algorithms revealed significant limitations in a widely used gas sensor dataset, particularly due to a non-randomised measurement protocol and severe sensor drift. These issues rendered the dataset unusable for classification benchmarks. Multiple studies that are impacted by this were identified, where the example of a prominent neuromorphic few-shot odour-learning algorithm study was investigated further. In response, a set of best practices for future gas sensor data collection campaigns was established to ensure data reliability and reproducibility. Several data collection campaigns were conducted using a custom-built e-nose system, which was based on MOx gas sensors and fast peripheral electronic devices. A novel approach to data feature acquisition for odour classification was proposed, involving rapid temperature cycling of the gas sensors. This method enabled the recording of two datasets capturing diverse indoor and outdoor olfactory scenes, which were effectively distinguished using the acquired features. An extensive laboratory campaign followed, in which the e-nose system was evaluated against a benchmark previously used to explore the temporal odour discrimination capabilities of mammals. The results demonstrated that the e-nose could distinguish correlated odour pulse trains from anti-correlated ones at modulation frequencies up to 40 Hz and determine frequencies up to 60 Hz, surpassing mammalian capabilities. Additionally, the

system achieved odour classification at pulse widths as short as 10 milliseconds when employing 50 millisecond duty cycles for sensor temperature modulation, setting a new precedent in artificial olfaction. Further, for efficient processing of olfactory signals, neuromorphic computing principles were explored. The potential advantages of asynchronous sampling and data processing methods were discussed in the context of the physical characteristics of turbulent odour plumes. Various event generation and processing algorithms were critically reviewed, and discussed in the context of olfactory signals. An example study is provided, in which asynchronous event sampling is applied to heater-cycled MOx sensor e-nose data, and the effectiveness of different event encoding schemes was assessed. Finally, the thesis discusses the results by putting them in perspective with future research directions.

Thesis advisers:

Dr. Michael Schmuker
Professor
Biocomputation Group
University of Hertfordshire, Hatfield
United Kingdom

Dr. André van Schaik
Professor and Director
International Centre for Neuromorphic Systems
Western Sydney University, Penrith
Australia

Dr. Volker Steuber
Professor
Biocomputation Group
University of Hertfordshire, Hatfield
United Kingdom

Examiners:

Dr. Nada Yousif
Senior lecturer
School of Engineering and Computer Science
University of Hertfordshire, Hatfield
United Kingdom

Dr. Tim Pearce
Professor
School of Engineering
University of Leicester, Leicester
United Kingdom

This thesis is dedicated to my dear brother, my mother, my father,

and my beloved grandparents.

Acknowledgements

First, I would like to express my deep gratitude to Prof. Michael Schmuker and Prof. André van Schaik; for their expertise, generous guidance and mentorship, ongoing support, and trust in me. Prof. Michael Schmuker has introduced me to the field of olfaction, which has been both highly exciting and humbling. Prof. André van Schaik has further exposed me to the field of neuromorphic engineering and applications thereof, and allowed me to be part of a stellar research team. Special thanks goes also to Prof. Volker Steuber, who has supported me throughout my studies, despite the minimal overlap of our fields of research.

I would like to thank my colleague Shavika Rastogi, who has been on a similar—yet different—quest that relates to artificial olfaction. Throughout the whole duration of our studies, she has been my companion and provided support and advice whenever possible. Next, I would like to acknowledge my colleague Damien Drix, whose efforts in designing an electronic nose have been fundamental for the work presented here, and who has been teaching me the ropes in its operation. Further thanks goes to Dr. Aaron True, who has been a genuine source of inspiration, and has demonstrated that team work across continents can be just as exciting and fruitful.

Acknowledgements

I would like to thank my colleagues at University of Hertfordshire, in particular Samuel Sutton and Dr. Maria Psarrou as well as my colleagues at the International Centre for Neuromorphic Systems (ICNS) at Western Sydney University, in particular Dr. Nicolas Owen Ralph, Dr. Alex Marcireau, Dr. Yeshwanth Bethi, Pablo de Abreu Urbizagastegui, Dr. Ali Mehrabi and Dr. Evie Andrew. I would like to thank Prof. Paul Hurley, Dr. Travis Monk and Dr. Saeed Afshar for their open doors, ears and minds. Further, I would like to acknowledge NeuroNex Odor2Action network, as well as the community around the CapoCaccia and Telluride neuromorphic workshops. Special thanks goes to the ICNS admin team, in particular Kathy Glavas and Sandi Griffiths; who are dedicated to much of all the less-visible work that allows us to focus on research.

On a personal level; I am deeply grateful for my partner Adriana Alicia Lara Pages, who has shown me love, support and encouragement, and offered well-needed distraction throughout this journey. I thank my dear friends that I deeply connect with despite the distance; especially Nicolas Künzli, Annika Lysdal Pedersen, Céline Nauer, Hugo Dupont, Giulia Serino, Modestas Filipavicius, Timo Koch, Hubert Hodel, Duncan Tweed and Bojana Pavlovic. I thank current and former flatmates Alan McFetridge, David Morris, Galen Sultman and Charlotte Harrison; who celebrated my best, endured my worst, and made me feel at home. Further, I thank the Contact Improvisation community and my climbing buddies in London and in Sydney.

My deepest gratitude goes to my family; in particular to my brother, my mother, my father and my grandparents. I am eternally grateful for their unconditional support, which have allowed me to pursue my ambitions and dreams.

Table of Contents

List of Abbreviations

2H 2-heptanone

ADC . . . . . Analog-to-digital converter

AER . . . . . Address-event representation

AI . . . . . Artificial intelligence

AL . . . . . Antennal lobe

AO . . . . . Artificial Olfaction

ANN . . . . . Artificial neural network

CNT . . . . . Carbon-nanotube

CPC . . . . . Conducting polymer compound

DAC . . . . . Digital-to-analog converter

DFT . . . . . Discrete Fourier Transformation

DOAS . . . . Differential Optical Absorption Spectroscopy

DVS . . . . . Dynamic vision sensor

EAG . . . . . Electroantennogram

EB . . . . . Ethyl butyrate

EC . . . . . Electrochemical

EPL . . . . . External plexiform layer

EU . . . . . Eucalyptol

FBC . . . . . Frame-based camera

FFT . . . . . Fast Fourier Transformation

FPGA . . . . Field Programmable Gate Arrays

GC . . . . . Granule cell

GPCR . . . . G-protein-coupled receptor

IA . . . . . Isoamyl acetate

LiDAR . . . . Light Detection and Ranging

MC . . . . . Mitral cell

MCU . . . . . Microcontroller unit

MEMS . . . . Micro-electromechanical systems

ML . . . . . Machine learning

MOB . . . . . Mammalian olfactory bulb

MOx . . . . . Metal-Oxide

MRO . . . . . Mobile Robot Olfaction

MS . . . . . Mass spectrometer

MZI . . . . . Mach-Zehnder interferometer

OB . . . . . Olfactory bulb

OR . . . . . Olfactory receptor

ORN . . . . . Olfactory receptor neuron

PCA . . . . . Principal Component Analysis

PCB . . . . . Printed circuit board

PID . . . . . Photoionization detector

PLIF . . . . . Planar laser induced fluorescence

PN . . . . . Projection neuron

ppb . . . . . Parts-per-billion

ppm . . . . . Parts-per-million

QCM . . . . . Quartz crystal microbalances

RDF . . . . . Random Decision Forest

RMSE . . . . Root Mean Squared Error

SAW . . . . . Surface acoustic wave

SNN . . . . . Spiking neural network

STDP . . . . Spike-timing-dependent plasticity

SVM . . . . . Support Vector Machine

UAV . . . . . Unmanned aerial vehicles

UGV . . . . . Unmanned ground vehicles

UV . . . . . Ultraviolet

VLSI . . . . . Very large-scale integration

VOC . . . . . Volatile organic compound

Chapter 1

Introduction

"Anything worth doing, is worth doing right."

– Hunter S. Thompson

The olfactory world offers a complex landscape of a wide range of odours. They are rarely encountered in isolation, but vary in composition and concentration. The rich array of olfactory stimuli offers a nuanced depiction of the environment, which influences a wide spectrum of animal behaviours. Natural odourants are typically carried as dynamic plumes by airflow patterns that are governed by environmental factors. Turbulences in the airflow disrupt concentration gradients that were formed through diffusion, resulting in the plume being structured in discrete odour packages of varying concentrations. This creates an odour signal that is dynamic both spatially and temporally [FR82a; MM91; MEC92; SS00; MC04; CVV14; MSA21]. See Fig. 1.1a for an illustrative depiction of an odour plume.

In the following, I will explore the complex and turbulent physical environment that leads to intermittent and rapidly changing odour stimuli. I will then discuss various adaptations observed in the animal kingdom, highlighting how various insect and mammalian species have evolved sophisticated mechanisms to rapidly identify and react to olfactory cues. Further, I will examine how contemporary artificial systems achieve efficient and fast odour processing capabilities, and elaborate on applications. Finally, I will summarise this thesis' contributions to the field of machine olfaction in challenging the temporal limits of sensing and processing.

1.1 Fast odour signals in a turbulent physical environment

Detailed examinations of spatiotemporal odour plume structures have been conducted in aquatic and aerial/gaseous environments [MM91; MEC92; MC04; CVV14; RMV22]. In particular, it has been shown that different spatial scales are dominated by different physical phenomena. Molecular diffusion, for example, is a slow and small-scale process [MEC92], while turbulent diffusion exhibits vigorous characteristics and spans a broad spectrum of temporal and spatial dimensions [MEC92]. Odour plumes are shown to be of intermittent nature [VI99; CVV14], characterised by single odour packets (also "whiffs" or "bouts") or clusters of odour packets (also "clumps"), separated by periods of no signal (also "blanks") (see Fig. 1.1c).

The timescales of bout durations and inter-bout intervals, as well as bout amplitude (or odour concentration) and bout-per-clump counts, depend on many environmental factors, such as the wind's velocity field and the vector between the encounter point and the odour source [CVV14]. In particular, the odour concentration of bouts in turbulent plumes is distributed as a power law [MM91; CVV14; RAH08], see Fig. 1.1b. Odour concentration fluctuations for a stationary observer can exceed 100 Hz [Yee+95b], while individual odour encounters can last a few milliseconds or less [CVV14]. Consequently, by considering the statistics of those odour encounters, one can extract spatial information of the odour source and its surroundings [FR82b; MA91; Hop91; MEC92; Vic+01; Wei+02; MC04; SBH16]. In particular, measured bout concentration variations as well as intermittency are indicators of the plume dimensions [FR82b]. The degree of temporal correlation between two encountered odours at a single point can indicate their separating distance [Hop91]. Conversely, analysing the correlation between same-odour encounters at multiple points in space indicates the rel-

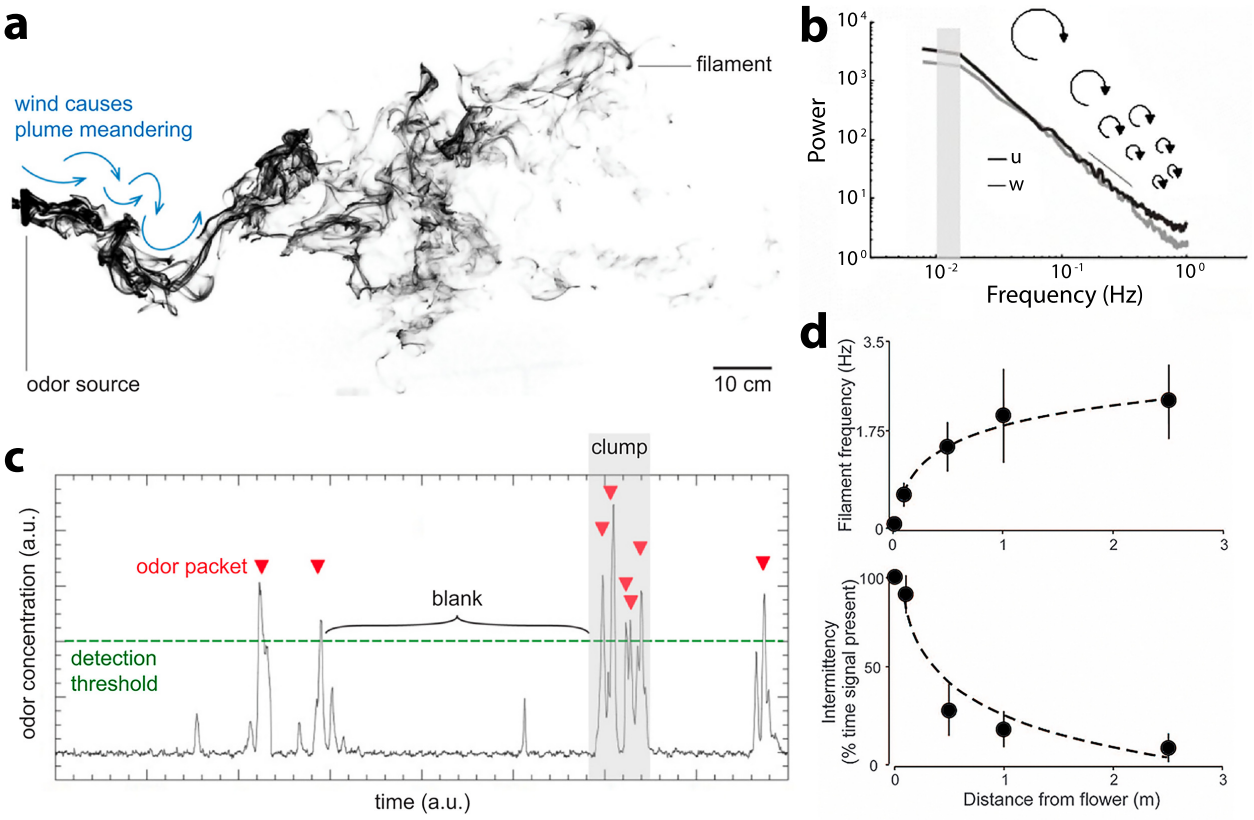

Figure 1.1: Odour plume dynamics are complex and require fast sensing. **a,** Meandering odour plume, visualised with $TiCl_4$ smoke. **b,** Energy spectra of wind velocities in two dimensions. Large eddies progressively cascade into smaller and smaller eddies until the turbulent kinetic energy is dissipated. **c,** Typical odourant concentration time series at a particular point in space, demonstrating different plume features (odour packets, blanks, clumps). **d,** Odour frequency (top) and intermittency (bottom) with increasing distance from the odour source (here, a flower). Panels a, b, c, d were adapted with permission from [SEE23], [RAH08], [SEE23] and [Rif+14] respectively.

ative position of the source [Wei+02]. Additionally, several plume features have been shown to reproducibly vary with distance and direction between the sensor and the odour source, such as the average bout count [SBH16], the concentration amplitude and first derivative of a bout [MC04], and the degree of intermittency [Rif+14] (see Fig. 1.1d). All those features can only be captured reliably if the plume is sensed and processed fast enough.

1.2 Fast olfaction in biology

Olfaction plays a vital role in various aspects of survival and behaviour, such as food localisation, predator-prey dynamics, mate selection and localisation, and other social interactions [CW08; Kad+22; Bla; Amo+08; HPB10; JDM14; Kha+21; Sul+15; BK04]. The ability to track the dynamics of fast odour signals is essential for success in these tasks. In the following, I elaborate on fast odour sensing in insects and in mammals. In particular, I will focus on mechanisms, temporal capabilities, and the ecological niche thereof.

Insect olfaction

The insect olfactory system (see Fig. 1.2a) starts at their antennae, which are covered with microscopic hair-like structures called sensilla, with pores that allow odourants to enter. Within each sensillum, there are one or multiple olfactory receptor neurons (ORNs), which are nerve cells specialised in detecting specific odour molecules [Kan+00]. Each olfactory receptor neuron (ORN) typically expresses one or multiple types of ORs on their cell membranes, allowing the insect to detect a broad range of odours through a combinatorial receptor code [VS07]. Once an odour molecule has bound to its corresponding receptor, a signal transduction cascade is initiated. This leads to changes in the intracellular levels of cyclic nucleotides and the opening of heterometric ligand-gated ion channels, ultimately generating a neuronal response that is transmitted to the antennal lobes for further processing [Suh+04]. There, the input from ORNs is organised into spatially separated glomeruli (see Fig. 1.2b), each receiving input from ORNs expressing the same OR type. Projection neurons (PNs) then relay the processed signals to higher brain centres, such as the mushroom bodies and the lateral horn, where the information is further integrated and associated with appropriate behavioural responses [Wil13b].

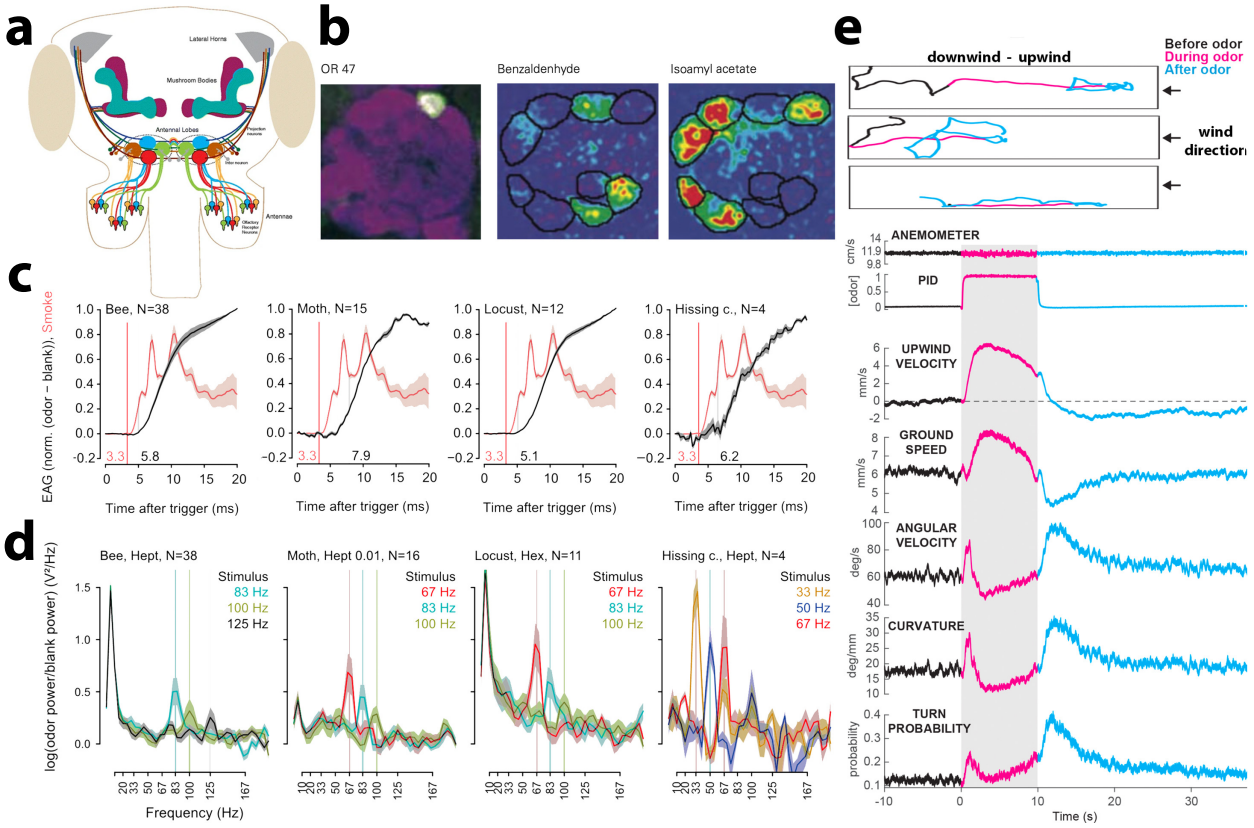

Figure 1.2: Fast odour sensing in insects. **a,** An overview of the insect olfactory system, showing stereotypic connections between olfactory receptor neurons, antennal lobes, mushroom bodies and lateral horns. **b,** Left: Olfactory receptor neurons that express the same olfactory receptor gene converge on the same glomerulus in the antennal lobe. Middle and right: Different odourants elicit different responses from a subset of the glomeruli. **c,** Olfactory transduction onset can occur in less than 2 ms. Shown here are $TiCl_4$ smoke signals (red) and electroantennogram (EAG) responses of different species to 2-heptanone (black). **d,** Antennal responses can track odour pulses at 125 Hz. Shown here are periodograms of the EAG responses for the three highest resolved pulse frequencies for different odours and species. **e,** Top: Example trajectories of three different walking *Drosophila melanogaster* before (black), during (magenta), and after (cyan) a 10 s odour pulse. Bottom: Calculated parameters of fly movement averaged across flies, showing behaviour that rapidly reacts to the presence and absence of odour stimuli. Panels a, b, c, d, e were adapted with permission from [CR16], [Kan+00], [Szy+14], [Szy+14], and [Álv+18] (CC-BY 4.0) respectively.

Insects exhibit remarkable abilities in rapidly detecting and processing odour stimuli, which enables them to track the dynamics of fast odour signals [Cri+22; SEE23]. Their olfactory receptor neurons' (ORNs) response latency is less than 2 ms [Szy+14; Ege+18] (Fig. 1.2c), and odour stimuli fluctuations can be resolved at frequencies of over 100 Hz

[Szy+14] (Fig. 1.2d), allowing for following fast concentration dynamics [BJS05; Gef+09; KLS11; Sch+08]. Evidence for odour identity inference at just tens of milliseconds exists in honeybees [Kro08], where Drosophila identify and react to particular odours within 85 ms [Bha+10], and mosquitoes can identify CO_2 packets of just 30 ms [DC11]. Distinguished responses — electrophysiologically and behaviourally — have been observed to be driven by odour onset asynchronies as short as 6 ms [SGS13], with evidence that complex odour dynamics are encoded by projection neurons in the insect brain [KLS15]. Further, adaptation mechanisms in the receptor neurons aid with resolving precise odour packet timing across background odour concentrations spanning at least four orders of magnitude [Gor+17].

For odour source localisation in turbulent environments, insects rely on a combination of odour and wind detection, employing strategies such as surging upwind when encountering odour and casting crosswind when losing the plume [BD14; Car21; SEE23]. The frequency and intermittency of odour packets play a crucial role in guiding motion, with agents biasing their movement upwind in response to these temporal features [KM74; MC94; BFC98; Álv+18; Dem+20]. Flies exhibit both temporal novelty detection and offset responses, adjusting their behavioural responses according to the temporal statistics of odour packets to extract spatial information for source localisation [Álv+18; Jay+23] (Fig. 1.2e). Additionally, insects can detect spatiotemporal odour concentration gradients using bilateral olfactory sensing, enabling them to navigate towards the centreline of the plume independent of wind direction [BH82; Tai+23; Kad+22]. Gradients in signal statistics can lead to the odour source [MA91], however integrating odour encounters over time requires retaining information in memory, posing challenges at greater distances from the source where odour packets are infrequent [VVS07; Rig+22b]. In such cases, "infotaxis" strategies (balancing random explo-

ration and exploitation of accumulated knowledge) can be more effective [VVS07; Rig+22b]. However, recent studies with fruit flies and mosquitoes have revealed a history dependence in flight decisions during odour tracking, suggesting that insects can hold information about odour plumes in memory and modulate navigation over longer timescales [Pan+18].

Insects can further exploit temporal cues to perceptually segregate mixed odours from different sources [Hop91; SEE23]. Various studies on mate choice [BFC98; NL02], host plant selection [And+11], and foraging [Szy+12; Seh+19; Sah+13] indicate that insects perceive two odourants as separate sources when their onsets are asynchronous and as one source when their onsets are synchronous. For example, when a male corn earworm moth detects its own species' female sex pheromone mixed with another species', it initiates a search flight; however, if both pheromones originate from the same source (indicating synchrony) it does not initiate [BFC98]. Similarly, honey bees and fruit flies prefer a mixture of an aversive and attractive odourant when the odourants arrive with a few-millisecond difference [Szy+12; Seh+19], suggesting spatial separation of sources. Animals may use other cues like recognising the target odourant during periods of its pure presence [Tai+23] or spatial sampling across antennae [BH82]. Fruit fly studies suggest they use temporal odour patterns to segregate odour sources [Seh+19], preferring asynchronous mixtures over synchronous ones. The neural mechanisms by which insects use those temporal cues for decoding spatial information about the odour sources remain largely unknown, posing questions and challenges to the field [SEE23].

Mammalian olfaction

For mammals, the process of olfactory sensing is similar to that in insects; however, it begins in the nasal cavity and in particular in the olfactory epithelium, where olfactory receptor neurons (ORNs) are located (Fig. 1.3a). Each ORN expresses only one type of olfactory receptor (OR) from a large gene family, enabling the detection of a wide range of odourants through combinatorial coding mechanisms [Mal+99]. Upon binding of an odourant molecule to its corresponding OR, a G-protein-coupled receptor (GPCR) signal pathway is activated, leading to an increase in intracellular cyclic adenosine monophosphate (cAMP) or inositol triphosphate (IP3), which then opens ion channels, resulting in neuronal depolarisation and the initiation of an action potential [BA91]. The olfactory signals are transmitted to and further processed in the olfactory bulb. ORNs project to mitral and tufted cells in glomeruli structures, where each glomerulus receives input from ORNs expressing the same OR type, thus maintaining the specificity of odourant detection. The mitral and tufted cells then relay the processed information to higher brain regions, including the olfactory cortex, where further processing and integration with other sensory inputs occur, leading to odour perception and the initiation of appropriate behavioural responses [She04]. Recent research has emphasised the importance of feedback mechanisms and neuromodulation in the olfactory system [LC16]. The findings suggest that the mammalian olfactory system is not only capable of detecting and discriminating a vast array of odourants, but also exhibits remarkable plasticity and adaptability, allowing for modulation of olfactory perceptions based on internal states and prior experiences.

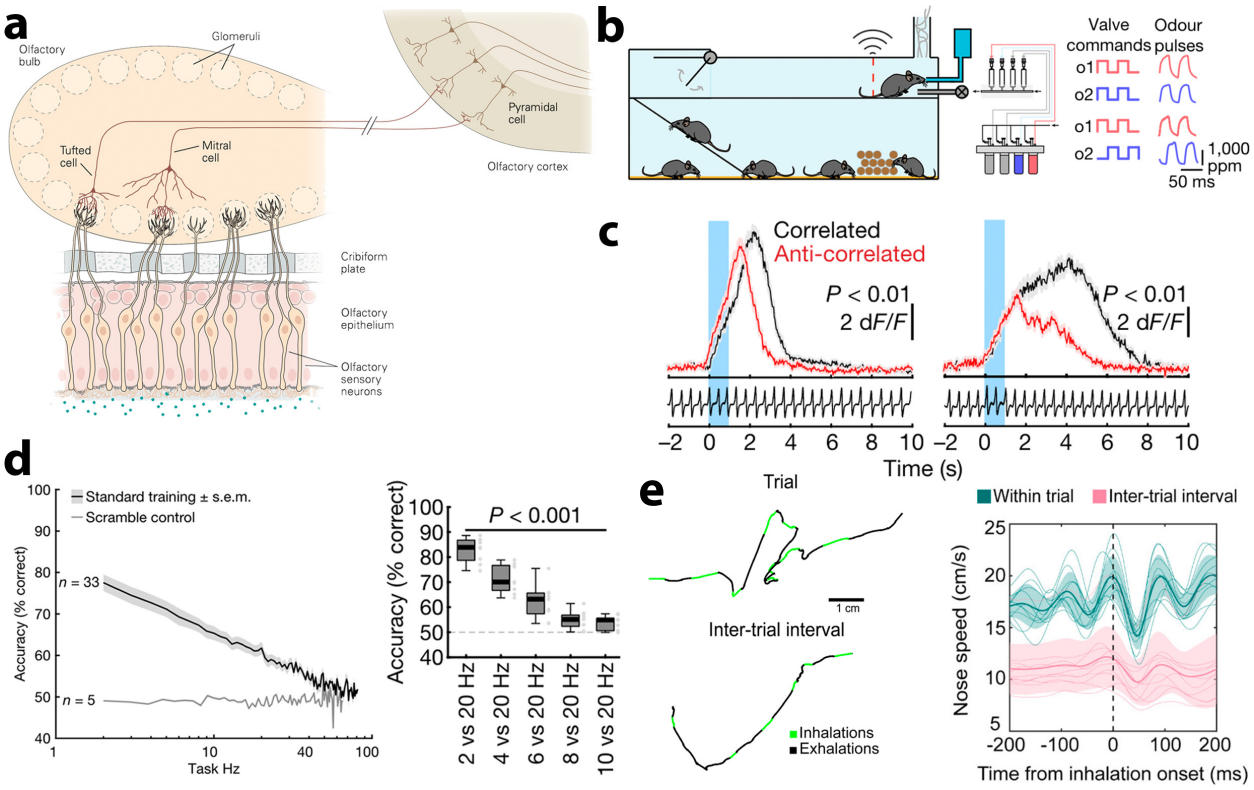

Figure 1.3: Fast odour sensing in mammals. **a,** Overview of the mammalian olfactory system, showing olfactory sensory neurons embedded in the olfactory epithelium, olfactory bulb with tufted cells and mitral cells, and the olfactory cortex. **b,** Automated olfactory stimulus delivery system, releasing odour pulses of short duration. **c,** In dual odour presentations, mice olfactory bulb output differs in their response kinetics between correlated (black) and anti-correlated odour pulse trains. **d,** In behavioural tasks, mice can accurately differentiate between correlated and anti-correlated odour pulse trains (left), as well as distinguish between two presented frequencies (right). **e,** Left: Example trajectories of mice in an odour navigation task. Right: Nose movement velocity is phase-locked and synchronised with inhalation and exhalation cycles. Panels a and e were adapted with permission from [Kan+00] and [Fin+21] (CC-BY 4.0) respectively, whereas b-d were adapted with permission from [Ack+21].

For a long time, demonstration of rapid olfaction in mammals was missing, and even thought absent. A possible underlying assumption could have been the fundamentally different—and substantially harder to study—odour sampling mechanism: while olfactory sensory neurons in insects are accessible and directly exposed to the odour stimulus, in mammals the odours usually need to be actively "sniffed" through the nasal cavity and pass a mucus layer to

reach the olfactory receptor neurons [Cri+22; MSA21], making it impossible to observe receptor-level stimulus dynamics without disturbing it. More fundamentally, multiple (but older) studies suggest that the odour pathway acts as a cascade of low-pass filters and that individual sniffs are the unit of information [DCD99; KUM06]. Therefore, the mainstream view has been that mammalian olfaction is slow and does not have access to rapid (sub-sniff) changes in odour concentration.

However, it has been shown early on that odour detection and discrimination can occur within a few hundred milliseconds [UM03; Abr+04]. Further evidence in awake rats and mice indicated that odours can evoke precisely sniff-locked activity in mitral/tufted cells of the olfactory bulb and odour-specific temporal spike patterns [CU10; Shu+11]. Optogenetic studies further suggested that information about evoked odour stimuli in the early olfactory system can be relayed with approximately 10 ms precision [Sme+11; Li+14; Reb+14]. In a recent landmark study, it has been shown that mice can decode and behaviourally differentiate odour stimuli at frequencies up to 40 Hz [Ack+21] (Fig. 1.3b-d), most likely by spatio-temporally averaging the neural activity across a large ensemble of olfactory receptor neurons [Abe04]. Experimentally, this was achieved through delivering temporally complex (but known) odour stimuli using high-speed odour delivery devices [RGS17; Ers+19a; Ack+21; Das+22] (Fig. 1.3b). Further, mitral cells and tufted cells of the mouse olfactory bulb can encode the dominant temporal frequencies present in odour stimuli up to at least 20 Hz [Das+22].

Mammals regularly overcome key behavioural challenges that require sampling olfactory stimuli efficiently, such as odour source separation [Ers+19a], odour-background segregation [Rok+14; Li+23], and odour localisation [Cat13; KSB12; Jac+15; Gir+16; Liu+20]. Recent studies reveal that mammals do indeed leverage the temporal dynamics of the odour stimuli for navigation tasks [Ack+21; Fin+21]. In particular, it has been shown that mice discriminate between stimuli derived from a single source or separated sources [Ack+21]. Further, navigation in noisy odour concentration gradients is possible with single sensors (without relying on stereo olfaction), where mice are observed to synchronise respiration and nose movement with tens of millisecond precision [Fin+21] (Fig. 1.3e).

1.3 Fast olfaction in machines

The physical landscape of the odour plume is complex, yet highly informative. The vast evidence that many animals efficiently decode and leverage this information for survival brings up the natural question of how machines perform in such tasks. Artificial odour and gas sensing have been an active field of research for over 40 years [Liu+12], both in academia and industry. Unsurprisingly, there are several parallels between biological and machine olfaction [Pea97a; Pea97b]. In the following, I will first elaborate on sensing mechanisms and then discuss different applications and their respective needs in terms of sensor response time and temporal precision.

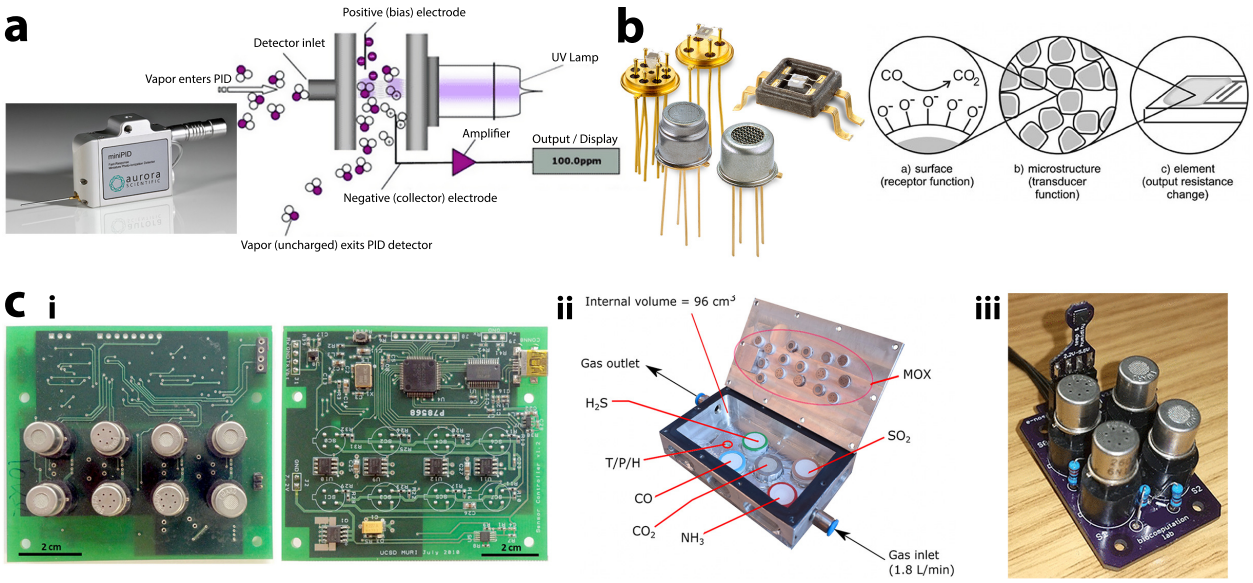

Figure 1.4: Different sensor modalities for artificial olfaction. **a,** Photoionization spectroscopy (PID), Aurora MiniPID device (left) and working principle illustration (right). **b,** Metal Oxide (MOx) gas sensors, devices (left) and working principles (right). **c,** Different instances of electronic noses, all based on MOx gas sensors. Panel a adapted with permission from [24a] and [24b], **b** from [24c] and [FKS06], and **c** i-iii from [Ver+14] (CC-BY NC), [Bur+21a] (CC-BY NC-ND), and [DS21] (CC-BY 4.0) respectively.

Sensor Modalities

Over the past decades, a broad variety of gas and odour sensor modalities have been developed, each with specific characteristics and ideal use cases. While these have been reviewed extensively [Pea+02; Liu+12; GGU23], I will mention only the ones relevant to this thesis. First, I will introduce Photoionization detectors (PIDs), an optical gas sensing technology that is often the choice for ground truth measurements. Then, I will elaborate more on Metal oxide (MOx) gas sensors, the preferred sensing modality for many applications. Finally, I will introduce the concept of the electronic nose.

Photoionisation detectors (PIDs)

Photoionisation detectors (PIDs) (Fig. 1.4a) are efficient gas detectors, which can measure volatile organic compounds (VOCs) and other gases in concentrations ranging from sub-parts per billion to 10000 parts per million (ppm). PIDs offer a very sensitive instrument for determining if and how much of a VOC is present, but in most cases, they are not very selective. Their attributes make them a preferred choice for ground truth gas concentration measurements; however, they are too bulky and expensive for most robotic applications.

For the analysis of a gas sample, the PID uses short-wavelength ultraviolet (UV) light to photo-ionise trace organic compounds in the sample, which is successful for any compound with an ionisation energy (E_I) lower than that of the UV lamp [Sta14]. The ionisation process results in in an electrophoretic motion of positively and negatively charged particles to the electrodes; the resulting net current is proportional to the concentration of ionisable compounds in the gas. Figure 1.4a illustrates the working principle of the sensor.

While the ionisation processes are very fast (femtoseconds to milliseconds), the response time of a PID is determined through the purge-and-flush rate of the gas in the detection chamber and is typically much slower (milliseconds to seconds). A typical PID houses a UV lamp with a single dominant photon energy $E_\nu = h\nu$. In this way, the sensor acts as a binary filter: every compound with $E_I < E_\nu$ is ionised and triggers a response, while everything with $E_I > E_\nu$ is ignored [Sta14]. There are different design choices for PIDs, which offer various trade-offs between bias voltage, sensitivity, and low-concentration response linearity.

Further, combining two PIDs with different ionisation energy [Ack+21], or using a device with two detecting units [SH04], has been shown to be effective. With such, one can distinguish between two compounds in a mixture. In both cases, it is crucial that E_ν is below the ionisation energies of any ambient air compound, which, fortunately, are quite high [Sta14].

Metal oxide gas sensors

Metal oxide semiconductor gas sensors (Fig. 1.4b) have established themselves as a common choice for many gas sensing applications, due to their tunable sensitivity, their space efficiency, and their low price. They consist of materials combining a metal and an oxide layer, such as SnO_2 and ZnO [Dey18], and most often a hotplate to bring the sensor to its working temperature regime of several hundred degrees.

The sensing principle can be understood as the combination of a receptor function (its ability to interact with a target gas) and a transducer function (its efficiency to convert the chemical interaction into an electrical signal). The receptor function is defined through the chemical properties of the oxide layer, while material morphology and crystallographic structure determine the transducer function [Wan+10]. In clean air, the surface layer adsorbs the oxygen from the air and thus traps negative charges. This bends both the valence band E_v and the conduction band E_c upwards, thus creating an electron-depleted region. When the target gas reacts with the oxide layer, the electrons are removed, reversing the process of band bending, which consequently causes the conductivity to increase. An illustration of the sensing principle is shown in Fig. 1.4b.

Historically, one challenge of using MOx sensors has been their relatively long impulse response duration, which can be on the order of tens to hundreds of seconds [PLT12a; DS21]. A large part of the impulse response is caused by the slow recovery phase, which can be orders of magnitude larger than the sensor response latency. There are different ways to partially overcome this drawback. Hardware solutions suggest, for example, purging of the sensor site [Gon+11a], active sniffing [Web+18], reducing size and micro-fabricating the sensors [Gar+10; Liu+18; GGU23], and de-capping the sensors [BVM19]. Another hardware approach is to modulate the temperature of the sensor hotplate. Traditionally, MOx gas sensors use a fixed temperature for their heating element, thus the reaction on the sensor site is isothermal. However, it has been shown that the rapid variation of the hotplate can decrease the response time of the sensor to a target gas dramatically [Ver+14]. The authors used oscillations with sub-second duty cycles, and achieved response times of 9 s compared to the 300 s resulting from constant temperature. Further studies examined this phenomenon by choosing different temperature profiles, such as temperature steps [Bau+18] and pulses [Xin+19]. Another promising approach is to modulate the temperature in-the-loop using the sensor output in a self-adapting way, which is described in [HT19] and more recently in [Di +21].

Software or post-measurement signal processing solutions can aid in further reducing the sensor response time. Many approaches rely on constant-temperature sensing mode and process the recordings using filters or filter banks. Cascaded filtering [Di +95; Mue+09; SBH16] can enhance fast transitions of the signal, giving rise to "bouts" (as defined earlier) [SBH16]. Similar but refined approaches have been suggested [BM19; BVM19], reducing response time and recovery time. Further, it has been shown that the slow response of MOx sensors can be mitigated by signal deconvolution, given the availability of an adequate

sensor model [MBM19]. By using multiple sensor instances with slight variations in response dynamics, similar results can be achieved without such a model, a process called "blind deconvolution" [MBM19]. Another proposed approach was to acquire the data at high resolution, then use Kalman filtering and absolute-deadband sampling [DS21]. This allowed for extracting fast onset events that may be used for informative stereo sensing. Using the described approaches — stand-alone or combined — has led to a decrease of reported response times down to seconds or even below [Xin+19; BVM19; MBM19; DDS22]. However, given the challenges in both delivering odours at high fidelity and in acquiring ground-truth measurements, it is difficult to define meaningful metrics or benchmarks that would allow adequate comparisons between the approaches.

Further challenges to using MOx sensors are the severe susceptibility of the sensor site to changes in temperature and humidity [Wan+10] and sensor drift. The latter describes the gradual and unpredictable variation of signal response over time when exposed to identical analytes under the same conditions [Ziy+10; Ver+13]. Drift is mainly caused by chemical and physical interactions on the sensor site, such as sensor ageing (re-organisation of the sensor surface over time) and sensor poisoning (irreversible or slowly reversible binding of previously measured gases or other contamination). Drift can partially be overcome by carefully crafting the experimental procedure when performing recordings, e.g., by randomising the order of analytes presented to avoid any correlation between gas identity and sensor drift, or by regularly measuring a reference gas and re-calibrating the system.

Electronic nose

An electronic nose, or e-nose, describes a device consisting of an array of multiple gas sensors, the necessary signal readout electronics, and a computing device [PD82]. Ideally, the selectivity of the gas sensors is broad and partially overlapping while maintaining high mutual information across the set, such that a large odour space can be measured. This strategy is particularly feasible for MOx sensors, as they offer a space- and cost-sensitive footprint while providing the ability to sense broadly. This makes them the preferred choice for many e-nose designs. However, other sensor modalities can be used too, such as conducting polymers [FL95], quartz crystal microbalances (QCM) [Sar+13], surface acoustic wave (SAW) sensors [Mat+19], micromechanical cantilever elements [Lan+99], biofunctionalised Mach-Zehnder interferometers [Lap+22], and more. Beyond the type of sensing element, there are different choices to make when designing an e-nose from scratch, such as the dimensionality of the sensor array (the number of different sensors), the performance of the used electric components (bandwidth, resolution, sampling rate, etc.), the computing device (microcontroller, FPGA, neuromorphic processor, etc.), and the way of processing the signals (sampling schemes, on-board algorithms, etc.). Different instances of electronic nose designs are shown in Fig. 1.4c.

Applications

Slow and static

Some common applications of artificial olfaction lie in the areas of industrial production and automotive technology; e.g., the detection of hazardous gases in mines [Kul+05; Sha+19] or polluting gases from vehicles [FT06]; medical applications, e.g., sensing volatile metabolic compounds for non-invasive screenings and diagnostics [Cha+97; DAm+10; Cha+18b; Wil18;

Zul+20]; food science, e.g., the discrimination of different edible compounds [Pea+93; Gar+94; ISY16; Van+22] or the identification of food spoilage [Jun+23]; domestic monitoring, e.g., the detection of fires [SPP06; FSM18] or formaldehyde gas traces [Gao+11]; as well as environmental, ecological, and agricultural studies [Fin+10; Wil13a; Ter+24], e.g., the monitoring of greenhouse gases [Dom+24], plant disease diagnostics [Cel+17], and agricultural odour measurements [MHP97].

Fast and dynamic

While the above-listed examples are often satisfied by static and relatively slow-responding gas sensing systems, more recent applications involve the sensing platform being mobile, which requires faster sensors. Both ground-based and aerial robotic olfactory navigation and mapping have been studied extensively, where many challenges have been identified and addressed.

Early on, indoor tests with odour sensors mounted on unmanned ground vehicles (UGVs) revealed that airflow disturbance by the vehicle adds turbulences to the environment; however, constant displacement velocity can help with stabilising the airflow, leading to new gas source localisation strategies [Lil+01]. Further studies on plume tracking and navigation strategies using UGVs have been performed [LD04; PGC07; HTK08; Loc+08a; Loc+08b; Ram+11; Lu13; Mam+13; Tak+14; ZLC15; HLC19], yielding valuable insights, e.g., that simple surge anemotaxis strategies result in efficient but less bio-plausible plume tracking [HTK08], that there are no obvious advantageous initial vehicle orientations and positions if there is no prior knowledge about the environment or source location [Lu13], or that plume tracking can be improved by swarms of identical vehicles compared to single entities [Mam+13]. Applications for gas sensing UGVs have been explored, such as gas leakage de-

tection and localisation [Wan+01; Ben+12], gas pipeline examinations [KBP09], methane distribution mapping on landfill sites [Her+12], gas discrimination and mapping in emergency response scenarios [Fan+19], hazardous gas detection in mines [KC21], and industrial plant inspection and supervision [Fis+24].

The use of unmanned aerial vehicles (UAVs) with olfactory sensing capabilities, particularly small drones, has gained traction in recent years. They come in different sizes, ranging from several kilograms down to a few grams [Dui+21]. Several navigation and mapping approaches have been tested using UAVs. In particular, adaptive odour sampling strategies using MOx gas sensors were shown to be advantageous when compared with predefined trajectories for the task of gas distribution mapping [Neu+12]. Several bio-inspired navigation algorithms were deployed onto micro-drones and compared in outdoor source localisation experiments [Neu+13b]. Nano-drone swarm strategies have been explored for efficient CO_2 mapping [NHB19]. Deep reinforcement learning is suggested for UAV path finding [Wu+19]; however, gas sensors have not yet been implemented in the design. Bio-hybrid systems have been suggested that combine moth antennae electroantennogram interfaces with MOx sensors, enabling small UAVs to perform odour source localisation [And+20]. Three-dimensional odour source localisation algorithms are suggested and deployed onto micro-drones, while carefully evaluating the effects of propeller-induced disturbances, sensor location, and environmental scenarios [EM20; EJM23]. For overcoming the difficulty of odour source localisation in large search spaces, an algorithm based on simulated annealing is suggested [YJM20]. It has been shown that the basic infotaxis algorithm can be improved by combining it with Gaussian Mixture Models [Par+21], demonstrating impressive mean search times. Physical advection-diffusion models can be combined with binary gas detection measurements (in contrast to measuring precise gas concentration), allowing for efficient gas source identification [Wie+22].

Similarly, considering physical models can improve UAV swarm coordination in dynamic environments, such as when exposed to rapid wind direction fluctuations [Hin+23]. Other studies explore probabilistic measures for gas source localisation, such as probabilistic gas-hit maps [OMG23], or the combination of Voronoi tessellation, particle swarm optimisation, and Bayesian inference [Pra+23]. While UAV-based gas sensing poses different challenges in terms of propeller-induced airflow disturbances, payload restrictions, and ideal sensor placement [Fra+22], they have been used successfully for many applications. Examples include monitoring carbon capture and storage sites [Neu+13a], near-field characterisation of industrial methane emissions [Nat+15], volcanic plume recordings [Are+19], source localisation of hazardous airborne chemicals [HLC19], swarm-based indoor air quality monitoring [NHB19], UAV monitoring in mining areas [Ren+19], disaster scene management [Wu+19], the monitoring of wastewater treatment plants [Bur+21b; Bur+21a], and gas source localisation in multi-building scenarios [YJM20] and cluttered environments [Dui+21].

Much of the reviewed literature indicates that — if dealing with sufficiently realistic environments — merely following concentration gradients is not effective or efficient enough in locating an odour source [Her+12; Neu+13b; ZZ23]. Instead, it is necessary to consider the much more complex plume dynamics, which can be estimated using situation-specific physical models [Hin+23], be sampled by using swarms of entities scattered across multiple locations in space [EJM23], and/or be measured by using fast enough sensors [GFS24].

1.4 Research objectives

In summary, adequately responding to the challenges posed by a dynamic and rapidly changing odour landscape is crucial for success in many robotic tasks and applications. For such, the requirement for high temporal resolution odour sampling and processing is evident, which is reflected in the remarkable timescales at which animals have evolved to process and respond to odours. During the scope of this thesis, different research objectives have been formulated, aiming to reduce latency and increase the speed of an artificial agent for sensing, processing, and responding to olfactory stimuli.

1. **Identify and Mitigate Dataset Bias in Gas Sensing**

   Investigate potential biases in widely used MOx gas sensor datasets, with a focus on dataset design and validation, to improve the reliability and accuracy of gas classification algorithms.

2. **Enhance Model Validation in Neuromorphic Odour Learning**

   Replicate and critically evaluate neuromorphic odour-learning algorithms and validation protocols. Emphasise the need for rigorous validation to ensure the applicability of these models to real-world odour identification tasks.

3. **Optimise Sensor Modulation for Urban Olfactory Scene Recognition**

   Explore the potential of rapid temperature modulation and multi-modal sensor data to improve recognition accuracy in urban olfactory environments.

4. **Develop a High-Speed, Miniaturised Electronic Nose**

   Design and validate a compact electronic nose capable of processing millisecond-scale odour pulses and encoding stimuli at frequencies comparable to biological olfaction, for applications in environmental monitoring and beyond.

5. **Investigate Neuromorphic Principles for Olfactory Processing**

   Analyse the potential of neuromorphic approaches to address the challenges of spatial and temporal sparsity in odour plumes, focusing on asynchronous acquisition and processing techniques.

6. **Evaluate Encoding Schemes for High-Resolution Odour Data**

   Compare different spike-based encoding schemes derived from asynchronous sampling techniques, identifying processing methods that leverage complex temporal patterns in sensor responses.

1.5 Contributions of this work

Addressing these research objectives has led to multiple scientific contributions and publications, which are elaborated on in the following:

1. I identified a previously unnoticed limitation in a widely used MOx gas sensor dataset, allowing a trained classifier to infer the analyte gas of a trial before gas deployment. I revealed that gases were recorded in temporally clustered batches, where the drifting sensor response baseline correlates with batches, inadvertently facilitating gas identification. I found numerous studies that used this dataset without addressing this issue, potentially

leading to overestimated accuracies in reported gas classification algorithms. I underscore the importance of careful dataset examination to avoid misleading results and discuss best practices for future data collection protocols. The findings are described in Section 2.1, and have led to the following journal publication:

Nik Dennler, Shavika Rastogi, Jordi Fonollosa, André van Schaik, and Michael Schmuker. "Drift in a popular metal oxide sensor dataset reveals limitations for gas classification benchmarks." In: *Sensors and Actuators B: Chemical* 361 (2022), p. 131668. DOI: 10.1016/j.snb.2022.131668

2. I replicated a study on a neuromorphic odour-learning algorithm, which uses the mentioned dataset for evaluation. I provide evidence that dataset limitations significantly impact the evaluation. Further, I exposed the model's insufficient generalisation capabilities, and proposed a much simpler method that achieves comparable or superior performance on the task. This highlights the need for further validation of the model's applicability to real-world odour identification tasks. Chapter 2.2 describes this in detail and is currently under review at *Nature Machine Intelligence*:

Nik Dennler, André van Schaik, and Michael Schmuker. "Limitations in odour recognition and generalization in a neuromorphic olfactory circuit." In: *Nature Machine Intelligence* 6 (2024), pp. 1451–1453. DOI: 10.1038/s42256-024-00952-1

3. I evaluated a MOx electronic nose dataset consisting of urban olfactory scene recordings, and showed the efficacy of modulating the sensor operating temperature at cycle periods much shorter than the literature suggestions. Further, I investigated using different sensor modalities for the task and concluded that a classifier with access to gas sensor responses may yield higher recognition accuracies compared to using humidity, temperature, and pressure. The findings are described in Section 3.1, and have led to a conference talk and the following conference proceedings:

Damien Drix[†], Nik Dennler[†], and Michael Schmuker ([†] denotes equal contribution). "Rapid Recognition of Olfactory Scenes with a Portable MOx Sensor System using Hotplate Modulation." en. In: *2022 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*. Aveiro, Portugal: IEEE, 2022, pp. 1–4. DOI: 10.1109/ISOEN54820.2022.9789654

4. I introduced a miniaturised electronic nose with high-bandwidth sensor readouts and advanced algorithms, which used the principle of short heater cycles introduced above. I demonstrated its capabilities in classifying brief, millisecond-scale odour pulses and encoding stimuli at frequencies up to 60 Hz, which is well beyond what has been achieved previously using MOx sensors. I showed that the system matches the rapid detection and recognition capabilities of animal olfaction, offering potential applications in environmental monitoring and other fields. A detailed elaboration on this study is provided in Section 3.2, which is currently under review at *Science Advances*:

Nik Dennler, Damien Drix, Tom PA Warner, Shavika Rastogi, Cecilia Della Casa, Tobias Ackels, Andreas T Schaefer, André van Schaik, and Michael Schmuker. "High-speed odor sensing using miniaturized electronic nose." In: *Science Advances* 10.45 (2024), eadp1764. DOI: 10.1126/sciadv.adp1764

5. I discussed how Neuromorphic principles may suit the demands of processing the complex dynamics inherent to olfaction. Investigating the challenges posed by the spatial and temporal sparsity inherent to odour plumes, and the resulting intermittency in encounters, I suggest that the asynchronous acquisition and processing of change-event samples may be advantageous when compared to regularly sampled data. I review and critically discuss the literature on Neuromorphic olfaction, and suggest which mechanism may be particularly suited for the inherent properties of the turbulent odour space. The discussion is found in Section 4.1, and will serve as the basis of an upcoming publication.

6. Finally, I compared different encoding schemes that are based on asynchronous send-on-delta sampling, using MOx e-nose sensor responses that are phase-locked to short temperature modulation cycles. The tested encoding schemes—namely rate code, latency code, and rank code—yielded significantly worse performance results compared to event reconstruction or the raw signal. I discussed that this may indicate dependencies on temporally more complex patterns in the sampled curves, suggesting processing schemes that leverage precise spike-timings. The study is described in Section 4.2, and has led to a conference talk and the following conference proceedings:

Nik Dennler†, Damien Drix†, Shavika Rastogi, André van Schaik, and Michael Schmuker († denotes equal contribution). "Rapid Inference of Geographical Location with an Event-based Electronic Nose." In: *The 9th Annual Neuro-Inspired Computational Elements (NICE), 2022, UTSA, USA Issue: 1.* Vol. 1. Association for Computing Machinery, 2022

Chapter 2

Limitations in Data and Algorithms

"Le doute n'est pas une état bien agréable,

mais l'assurance est un état ridicule."

– Voltaire

In this chapter, I explore a journey familiar to many doctoral candidates in the computational sciences. It begins with the development and evaluation of novel algorithms, often carried out with a blend of optimism and rigorous scientific inquiry, using publicly available datasets as benchmarks. Early results frequently suggest promising avenues for practical application and innovation; however, the optimism may be tempered by the realities encountered beyond the controlled environments of these datasets. When the algorithms are finally applied to real-world scenarios, the complexities and unpredictabilities of practical contexts can present challenges that are starkly different from those foreseen during the algorithm's developmental phase.

In the first part, I introduce an electronic nose dataset [Ver+13], which has facilitated a wide range of studies in the field. It is esteemed for its comprehensive data space, meticulously detailed documentation, and its accessibility. Then, I reveal its limitations, which—unfortunately—constrain the dataset's use in gas classification benchmarks. Finally, I discuss a set of guidelines and best practices for data collection and processing protocols.

The second part discusses how the discovered limitations affect a study of particular interest [IC20] for fast machine olfaction. The referred work introduced and evaluated a bio-inspired algorithm for artificial olfaction and implemented it on a neuromorphic platform. It has since received broad exposure and has brought traction to the field of neuromorphic olfaction. Besides its dependency on the dataset in question, I reveal multiple drawbacks and limitations of the performance evaluation, leading to the question of the potential applicability of the proposed algorithm to real-world scenarios.

2.1 Limitations in MOx sensor dataset for gas classification benchmarks

During the early stages of my studies, I was exploring different gas identification algorithms, with my particular interest being in finding methods that can classify a gas in a very short time. For this objective, it is paramount to have a well-parameterised dataset at hand on which the algorithms can be evaluated. I believed to have found such a dataset in the work by Vergara et al. [Ver+13], which consists of thousands of MOx sensor electronic nose measurements, recording different analyte gases under varying environmental conditions.

The surprising discovery occurred when it was noted that the first tested classification method (a linear Support Vector Machine) could learn and infer the analyte class from short data windows that were sampled *before* the sensors were physically exposed to the gas. A detailed analysis followed, tracking down the reason for this unexpected behaviour. The answer was found in the metadata, particularly the recording timestamps embedded in the filenames, revealing that the space of free parameters (gas identity, wind speed, etc.) had not been sampled at random—as indicated in the publication—but in gas-specific batches.

This led to a coupling between the recording time and the class analyte. As described earlier, MOx sensors are most often affected by sensor drift, resulting fluctuations in the sensor baseline response as well as in its responsiveness that are not directly caused by the analyte, but by either environmental factors (e.g. temperature, humidity, pressure) or sensor condition (e.g. age and poisoning). This sensor drift, together with the non-random recording protocol, renders the dataset useless for gas classification evaluations.

The remainder of this section is adapted from the peer-reviewed publication, which is accessible under the CC-BY license:

Nik Dennler, Shavika Rastogi, Jordi Fonollosa, André van Schaik, and Michael Schmuker. "Drift in a popular metal oxide sensor dataset reveals limitations for gas classification benchmarks." In: *Sensors and Actuators B: Chemical* 361 (2022), p. 131668. DOI: 10.1016/j.snb.2022.131668

The co-authors—in the following abbreviated as "we"—contributed as follows to this work: S.R. aided in an exhaustive literature scan across the studies that have cited and/or used the dataset in question. J.F. provided us with insights on how the dataset was collected and added valuable points to the discussion. A.v.S. co-supervised the project. M.S. added valuable points to the discussion and co-supervised the project. All co-authors assisted in editing the final manuscript. My contributions were the following: Performing the initial analysis and finding the limiting factors in the dataset. Performing a thorough analysis of the dataset. Scanning the literature. Drafting and editing the manuscript. Editing and revising the manuscript during peer-review.

Abstract

Metal oxide (MOx) gas sensors are a popular choice for many applications, due to their tunable sensitivity, space efficiency and low cost. Publicly available sensor datasets are particularly valuable for the research community as they accelerate the development and evaluation of novel algorithms for gas sensor data analysis. A dataset published in 2013 by Vergara and colleagues contains recordings from MOx gas sensor arrays in a wind tunnel. It has since become a standard benchmark in the field. Here we report a latent property of this dataset that limits its suitability for gas classification studies. Measurement timestamps show that gases were recorded in separate, temporally clustered batches. Sensor baseline response before gas exposure were strongly correlated with the recording batch, to the extent that baseline response was largely sufficient to infer the gas used in a given trial. Zero-offset baseline compensation did not resolve the issue, since residual short-term drift still contained enough information for gas/trial identification using a machine learning classifier. A subset of the data recorded within a short period of time was minimally affected by drift and suitable for gas classification benchmarking after offset-compensation, but with much reduced classification performance compared to the full dataset. We found 18 publications where this dataset was used without precautions against the circumstances we describe, thus potentially overestimating the accuracy of gas classification algorithms. These observations highlight potential pitfalls in using previously recorded gas sensor data, which may have distorted widely reported results.

Introduction

Over the last 50 years, artificial olfaction has evolved from an almost niche field of study into a thriving interdisciplinary research area. Many use cases have been addressed, for example the detection of hazardous gases or pollutants [SJR86], spoilage localization [Mai+06], mobile olfactory robotics [LLD06], health monitoring [AJT20] and medical screening [Loi+13]; and artificial olfaction is expected to address many more use cases in the future [Cov+21]. A key challenge in artificial olfaction is to identify a range of odorants at high specificity. One way to achieve this is to use an array of multiple gas sensors, each with a rather large selectivity and low specificity, and extract the identity of the presented odor using pattern recognition. Metal oxide (MOx) gas sensors are widely used for such sensor arrays. Their sensing layer can be tuned to different analyte classes and they require little electronic periphery, which simplifies sensor design, reduces cost and saves space. One big drawback of MOx sensors is their susceptibility to sensor drift—the gradual and unpredictable variation of signal response over time when exposed to identical analytes under the same conditions [Ziy+10]. Drift is mostly due to chemical and physical interactions on the sensor site, such as sensor ageing (reorganization of the sensor surface over time) and sensor poisoning (irreversible or slowly reversible binding of previously measured gases or other contamination). Environmental effects such as changes in humidity, temperature or pressure also affect the sensor response. The impact of sensor drift can be reduced by careful experimental design that avoids any correlation between gas identity and drift, for example by randomizing the order of analytes presented. Where this is not possible or not desired, it is essential to be aware of the presence of drift and design analysis algorithms accordingly.

Setting up an electronic olfaction system still requires custom design of electronics and data analysis procedures. These designs and algorithms must be based on reliable data. There are many parameters that can affect MOx sensor recordings; e.g. environmental conditions like temperature and humidity, technical constraints like wind tunnel design/construction, flow control of analytes, turbulent dispersal, gas availability and associated safety requirements, among others. Previously recorded datasets from reputable sources are therefore popular in the Artificial Olfaction / Mobile Robot Olfaction (AO/MRO) community, since they reduce the need for recording data in the initial design stages. A number of datasets are publicly available, covering a range of tasks and use cases [RDR10; Ver+13; Ver+12; Fon+14; Fon+15; Ziy+15; Fon+16; BJM18; Gam+19].

One of the most popular datasets contains MOx sensor data sampled in a wind tunnel, for different gases and different experimental parameters, over a time of 16 months [Ver+13] (downloadable at [Ver+]). This publication has been cited more than 100 times[1]. The dataset has been used as a benchmark for gas classification algorithms in at least 18 publications [BBK18; VL16; IC20; CL19; Zho+19b; Mon+16; Fan+18; Gam+21; Zho+19a; KH20; MRS18; AGS19; Ver18; MG15; GSD16; CCP20; Wan+21a; Li+21]. It has also been used for gas source location estimation [SBH16; Bur19; BM19; BM20b] and other applications [Lee+17; Mon+17; SHS18; MYS20; CLG19; GB21].

Here, we reveal a fundamental limitation of the dataset published in [Ver+13]. First, we observed that gases were not presented in random order, but in distinct batches, sometimes recorded weeks or months apart. In consequence, the sensor recordings were contaminated by slow baseline drift effects that are characteristic for the time of recording. We show that

[1] According to Google Scholar as of December 2021

Col. no.	Sensor model	Col. no.	Sensor model
1	TGS 2611	5	TGS 2600
2	TGS 2612	6	TGS 2600
3	TGS 2610	7	TGS 2620
4	TGS 2602	8	TGS 2620

Table 2.1: Metal Oxide (MOx) sensors included in each 8-sensor array. All sensors were manufactured by Figaro USA, Inc[Fig].

since both gas identity and sensor baseline correlate with time, it is possible to identify trials using a specific gas only by looking at the baseline response, before any gas has been released. In addition, we show that even after correcting for slow drift by subtracting the average of the first few sensor readings of each experimental trial, residual short-term drift effects are characteristic enough to identify trials where specific gases have been used, using the baseline alone. Moreover, when further minimizing the impact of drift by selecting the least-affected subset of recordings and compensating for drift as much as possible, the gas classification performance is far inferior to the numbers we obtained when using the full dataset. Therefore we conclude that this dataset is only of limited use for gas classification benchmarking, and that previously reported classification results based on this dataset are likely overestimating the true accuracy of gas recognition. Finally, we give a perspective on how the measurement protocol could be improved to mitigate this problem, and elaborate on what tasks the dataset can be appropriately used for, i.e., tasks that are not affected by the drift contamination.

Dataset

The dataset in question [Ver+13] consists of 18000 time-series measurements recorded over a period of 16 months from a 72 MOx gas sensor array-based chemical detection platform exposed to 10 different analyte gases (Acetone, Acetaldehyde, Ammonia, Butanol, Ethylene, Methane, Methanol, Carbon monoxide, Benzene, and Toluene). The sensor platform consisted of nine modules, each equipped with eight MOx sensors (see Table 2.1 for sensor types). It was placed in a $2.5\,\mathrm{m} \times 1.2\,\mathrm{m} \times 0.4\,\mathrm{m}$ flat-bed wind tunnel, at six different distances from the gas inlet, perpendicular to the wind direction (see Figure 2.1a for a schematic). Each sensor module was integrated with a sensor controller, which enabled data collection at 12-bit resolution and a sampling rate of 100 Hz. Gas flow was adjusted by computer-supervised mass flow controllers. A multiple-step motor-driven exhaust fan controlled the wind speed.

Different experimental conditions were tested, namely three different wind speeds set by the fan ($0.1\,\mathrm{m\,s^{-1}}$, $0.21\,\mathrm{m\,s^{-1}}$, $0.34\,\mathrm{m\,s^{-1}}$) and five different sensor operating voltages (4.0 V, 4.5 V, 5.0 V, 5.5 V, 6.0 V). Before each measurement, a combination of the experimental parameters *gas*, *location*, *wind speed*, *operating voltage* was selected, until each combination was repeated 20 times. Each measurement lasted for 260 s, where gas was released between $t = 20\,\mathrm{s}$ and $t = 200\,\mathrm{s}$. Before and after each experiment, the wind tunnel was ventilated at the maximum speed ($0.34\,\mathrm{m\,s^{-1}}$) for two minutes to assert the reestablishment of sensor response baseline.

The data is deposited as raw sensor data with one file per trial and parameter combination. The time of recording of individual measurements was encoded as part of the name of the file containing the time-series, alongside the parameters used in that recording and the trial sequence.

For our analysis, we interpolated and re-sampled the data for dealing with missing data points, and further converted the sensor voltage readings V_{sensor} given in the dataset to sensor resistance values R_{sensor}, according to Eq. 2.1,

$$R_{sensor} = 10\,\mathrm{k\Omega} \times \frac{3.11\,\mathrm{V} - V_{sensor}}{V_{sensor}}\,. \tag{2.1}$$

The readings of sensor 1 for all boards were discarded due to excessive sensor noise. Figure 2.1b shows the responses of a sensor board to one gas in a typical trial. Unless stated otherwise, we used the wind tunnel location P4 B5 (wind-downstream from the gas source, see Figure 2.1a for wind tunnel schematics), as we expected a high gas exposure at that location.

Results

Non-random order of gas measurements

We extracted the times of recording from the filenames to analyse the temporal order of measurements. Figure 2.1c shows when each the 18000 measurements have been made, arranged by gas identity and sensor position. It is evident that gases have not been measured in random order, but in separate batches that cluster in time. Only rarely do measurements of different gases overlap in time (as for Ammonia, CO and Toluene); more often, measure-

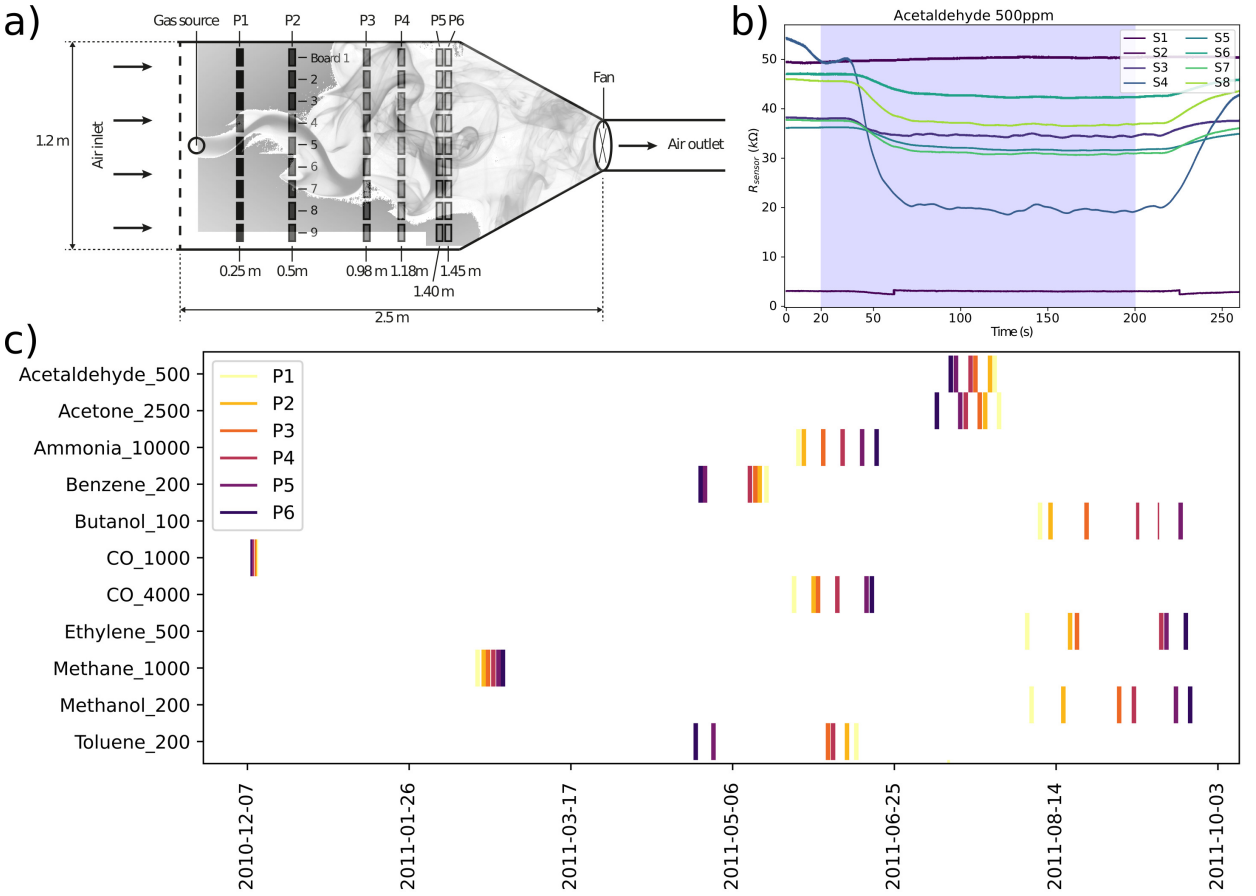

Figure 2.1: Experimental procedure of Vergara et al. [Ver+13] **a)** Sensor boards with 8 MOx sensors each are placed at different locations in a wind tunnel. Air gets sucked in by a fan, and a gas source can be opened or closed. Adapted from [Ver+13]. **b)** Measured sensor resistance for all sensors on one sensor board (location P4, module 5, Acetaldehyde, $0.21\,\mathrm{m\,s^{-1}}$ airflow velocity, $6\,\mathrm{V}$ operating voltage, trial 1). The shaded portion denotes the period during which the analyte was injected into the wind tunnel. **c)** Event-plot of timestamps for all gas trials contained in the dataset. Each vertical line represents 300 trials, which were performed too close to each other for them to be visually distinguishable in this representation. The row name indicates the measured gas and its concentration in parts-per-million (*ppm*). The blank spaces indicate periods where there have been no data collected. CO at 1000 ppm was removed from further analysis since significantly fewer trials were performed than for the other analytes.

ment batches are several weeks apart (e.g., Toluene and Methane). In no case have gases been alternated on a per-trial basis. In addition, we observed that also other experimental parameters like distance-to-source, wind speed, sensor temperature were selected in a sequential fashion rather than in random order (not shown).

The batched arrangement of gas identity and parameter settings is not evident from the description of the dataset provided by the authors, neither in the original paper, nor in the documentation contained in the UCI repository [Ver+]. Describing the experimental protocol, it was stated that (quote) *"This measurement procedure was reproduced exactly for each gas category exposure, landmark location in the wind tunnel, operating temperature, and airflow velocity in a random order and up until all pairs were covered."*[Ver+13]. This could be read as to imply that all experimental parameters that define an experiment were selected randomly before each trial, including which gas to release—which would mitigate, to a large extent, the detrimental effect of baseline drift on gas identification benchmarks—which is not the case, as we show here.

Drift in baseline over time

We investigated the sensor baseline across trials, where here we defined baseline as the sensor readings measured before gas is released into the wind tunnel. Figure 2.2a shows the trial-wise average of sensor baseline values at times $t < t_{release} = 20\,\mathrm{s}$, for a fixed sensor board location, operating temperature and airflow velocity, versus the date of recording. We observed that baseline varies significantly over time. Long-term drift can be observed as significant discontinuities between recording sessions. Since gas presentations were batched, the baseline pattern often correlated with gas identity. In addition, substantial baseline drift could be observed within some recording sessions.

Spatial distribution of baseline variations

By design, the gas plume does not distribute homogeneously across the wind tunnel, but disperses in a turbulent manner.

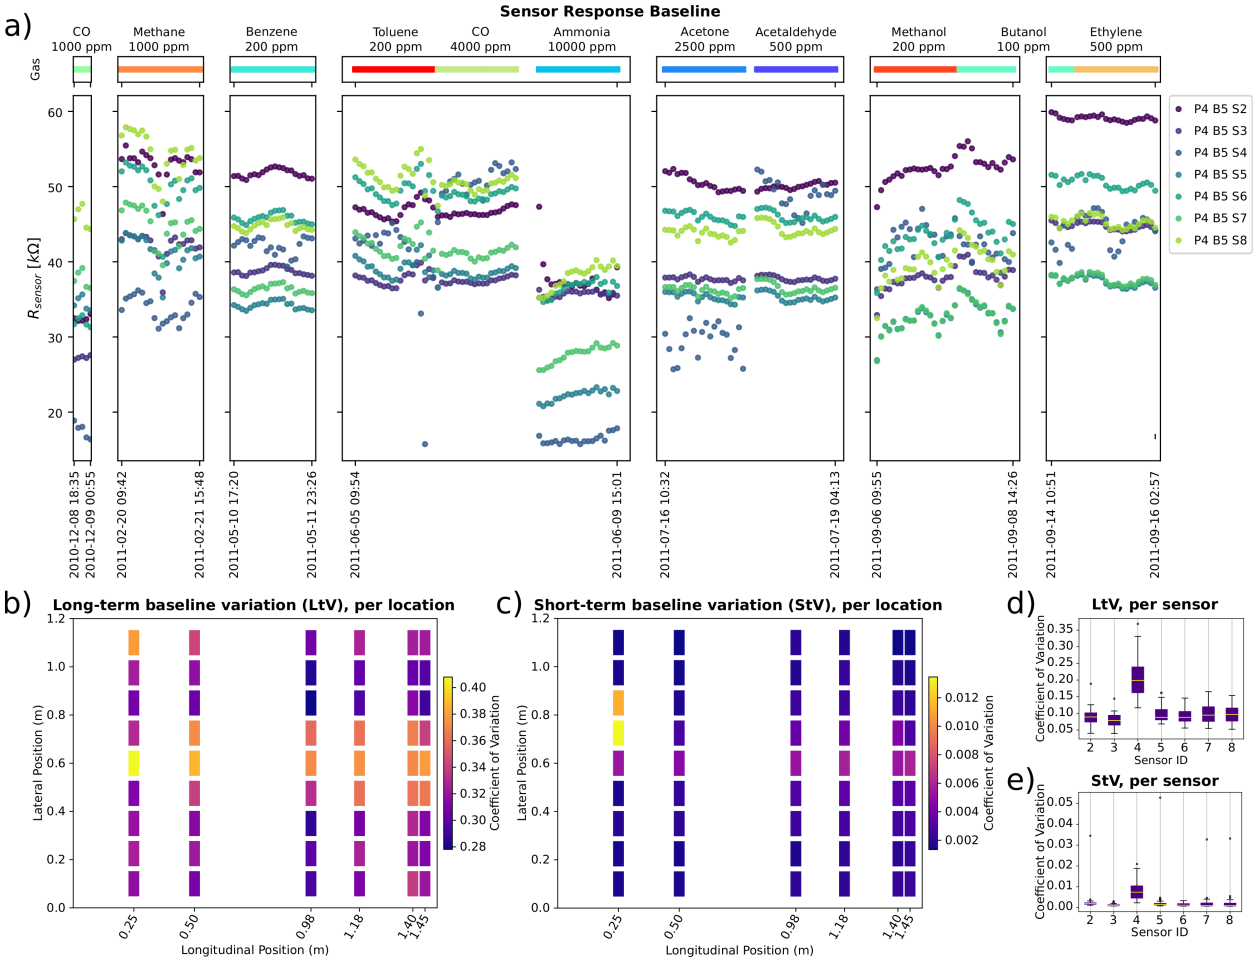

Figure 2.2: Drift analysis of Vergara et al. dataset [Ver+13]. **a)** Baseline for each sensor and experimental trial. Dots represent the mean sensor resistance during the time before gas release (20 s). Top row indicates the gas and its concentration (in *ppm*) used in the corresponding sessions. **b)-e)** Local baseline variation analysis using the coefficient of variation, for spatial wind tunnel location (**b)**&**c)**) and sensor board (**d)**&**e)**). **b)**&**d)** display the long-term baseline variation across the whole experimental duration (16 months), where **c)**&**e)** display the averaged within-trial, short-term baseline variation. Data shown here was obtained with wind flow speed 0.21 m s^{-1} and hotplate voltage 6 V. All ten gases and sensors 2 - 8 were considered. For **a)**, only location 4 and board 5 were considered (see Figure 2.1a for wind tunnel schematics).

Consequently, the total gas exposure at different sensor sites could vary, which may alter each sensor's response differently. Here we investigated how variations in the baseline response were distributed across the wind tunnel and the sensor board, as a proxy for sensor drift effects.

To quantify these variations, we calculated the coefficient of variation (c_v) of the baseline, for each sensor and each board location. The coefficient of variation is also known as normalized root-mean-square deviation, or relative standard deviation, and should allow for a comparison of the variation that is agnostic to the absolute sensor values.

Generally, c_v is given by the fraction between the standard deviation σ and the mean μ (Eq. 2.2),

$$c_v = \frac{\sigma}{\mu}. \tag{2.2}$$

We discriminated between *long-term* baseline variations over the whole duration of the experiment, and *short-term* drift within single trials. To quantify long-term variations, σ and μ were computed from the distribution of trial-wise averages of sensor baseline values, thus c_v described the variation of the baseline across the duration of the whole data collection period. For short-term variations, c_v was taken as the average of the σ-to-μ ratios of the single-trial baseline values.

We observed a distinct spatial pattern in the distribution of long-term baseline variations across the wind tunnel (Figure 2.2b). The long-term drift effects were strongest in sensor boards close to the center line of the wind tunnel, where gas concentration was expected to be highest. This observation suggests that long-term drift could be caused by exposure to the sample gas. Long-term drift affected all sensors, although sensor 4 was affected most strongly (Figure 2.2c. This sensor is a Figaro `TGS 2602`, which is targeted towards "Air pollutants (VOCs, ammonia, H2S)" according to Figaro's website.

The values for within-trial short-term coefficients of variation were naturally lower in magnitude but exhibited a similar pattern as observed for long-term variations, both spatially and per-sensor (Figures 2.2d and 2.2e). This indicates that also within-trial drift was highest for those sensors that were exposed to the highest gas concentrations.

Gas Clustering and Classification

Since both the baseline drift and the identity of the gas used in a trial correlate with time, we tested how much information about the gas could be obtained from the baseline signal alone.

Figure 2.3a shows a Principal Components Analysis (PCA) plot of the raw baseline values for each gas, at a range of times after the start of trials. Each plot presents a snapshot of a 100 ms time window, within which the time-series data of the sensor responses was averaged and used for the PCA. The PCA was computed using all trials in all windows, thus each snapshot is a projection of the data into this shared PC space. Distinct gas-specific clusters could be observed already at $t = 0\,\mathrm{s}$, before any gas was released into the tunnel at $t_{release} = 20\,\mathrm{s}$. The clusters change slightly between 30 and 40 seconds, which we assume is when the gas had reached the sensors.

Next, we attempted to compensate for long-term drift effects by subtracting the average of the first 100 ms window, i.e. for $t \in [0.0\,\mathrm{s}, 0.1\,\mathrm{s})$. The data then only contains the difference of the sensor response relative to the start of a trial. This is a standard procedure when dealing with MOx-sensor data. For Figure 2.3b we computed a PCA on the data compensated for long-term drift and used the same windows as before to visualize the evolution of sensor responses. By design, at $t = 0\,\mathrm{s}$ there are no visible clusters. Interestingly, although the

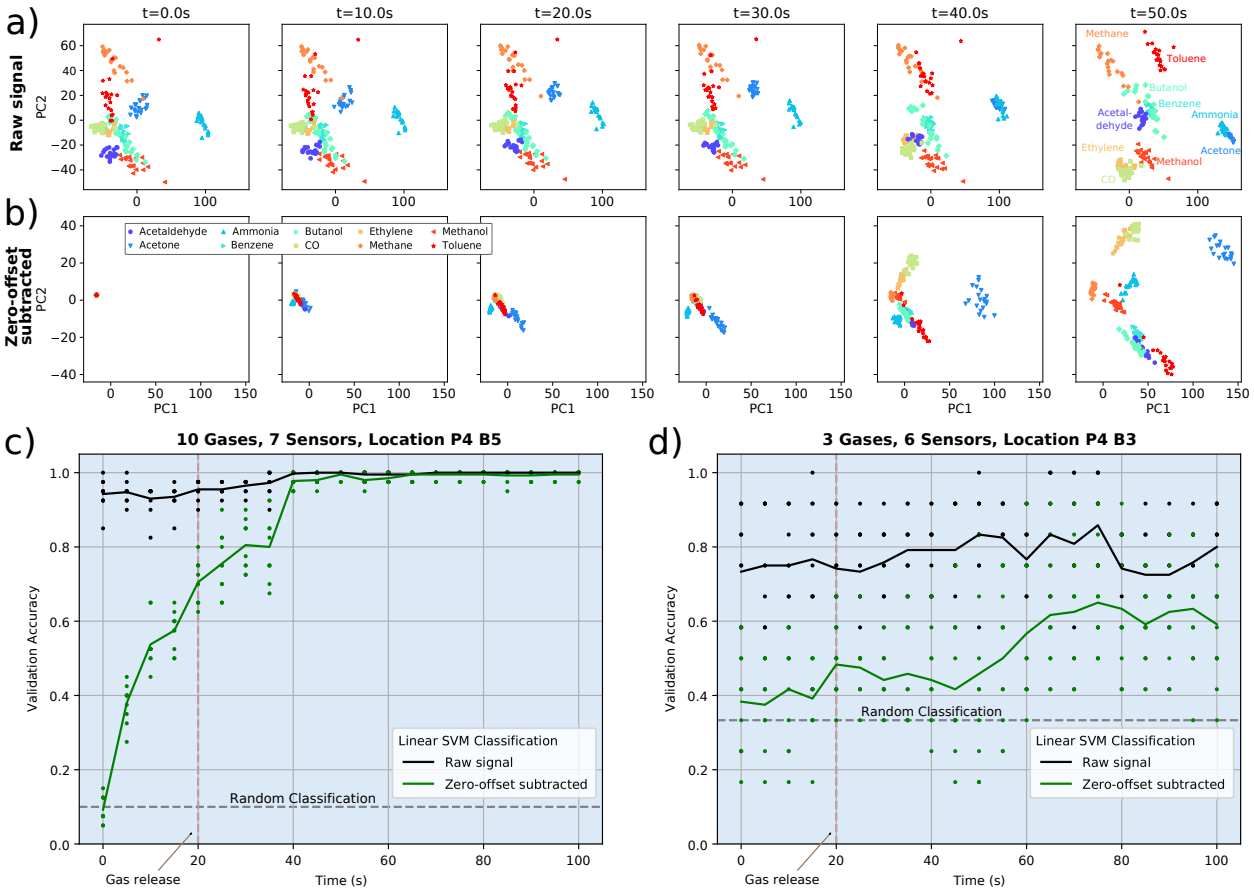

Figure 2.3: PCA analysis and SVM gas classification of Vergara et al. dataset [Ver+13]. **a)&b)** Principal Component Analysis (PCA) of samples within a 100 ms time window, at different starting times. Each color and shape corresponds to a different gas. For **a)**, the raw resistance signal was considered, where for **b)**, the zero-offset was removed by subtracting the mean resistance in the first 100 ms for each sensor. **c)&d)** Classification results using a Linear Support Vector Machine classifier. The trials for each gas were randomly split in training and validation datasets with a ratio of $80-20$. Black corresponds to the raw resistance signal, whereas green corresponds to the zero-offset subtracted signal. For all experiments shown here, the wind flow speed was fixed at $0.21\,\mathrm{m\,s^{-1}}$, and the hotplate voltage was set to 6 V. For **a)-c)**, all ten gases and sensors 2 - 8 have been considered, at location 4 and board 5 (see Figure 2.1a for wind tunnel schematics). For **d)**, the gases Methanol, Ethylene and Butanol have been considered, measured with sensors 2-3 and 5-8, at location 4 and board 3.

zero-baseline has been subtracted, we still observe the formation of clusters before the release of the gas. We interpret this observation as the manifestation of short-term drift within trials (*cf.* Figures 2.2d and 2.2e). It indicates that short-term drift also changed over time, in a way that correlated with gas identity.

These observations were confirmed using a time-windowed supervised classification approach with a soft-margin Support Vector Machine (SVM) classifier. We used a linear kernel with regularization parameter $C = 1.0$. The classifier was trained and tested separately for each time window, using the same features as for the PCA (time-series data of the sensor responses). We used a 4-to-1 random training to test split, i.e. training on 80% of the trials in each window and testing using the remaining 20%, repeated 10 times with different 4-to-1 random splits. Figure 2.3c shows the classifier performance. As expected, the classifier yielded near-perfect gas recognition performance on the raw data (i.e., without compensation for long-term drift), with an average accuracy of 94.3% already on the first time window of a trial, for $t \in [0.0\,\mathrm{s}, 0.1\,\mathrm{s})$. Test accuracy increased slightly for later time windows. From $t = 40\,\mathrm{s}$ on it converged at 100%. We assume that this is when the sensor board was maximally exposed to the gas.

Compensating for baseline offset did not rectify these classification artefacts. While test accuracy was random for the time window at $t = 0\,\mathrm{s}$, it was clearly above random already at $t = 5\,\mathrm{s}$ (Figure 2.3c, green line). It increased further to around 80% at $t = 35\,\mathrm{s}$, before making a step to near 100% at $t = 40\,\mathrm{s}$.

Taken together, we observed that the time window before gas exposure contained enough information to identify the gas used in a particular trial, even before the gas has been released into the wind tunnel. Baseline compensation for long-term drift reduced the extent of the problem, but there was still sufficient information contained in the short-term drift dynamics that allowed identification of the gas used in a given trial far above chance level. Noteworthy

here is that even gases that are measured in close temporal proximity (such as CO at 4000 ppm and Ammonia) separated well in PC space. We suspect that not only time-related sensor ageing, but also the slow recovery phase after gas exposure and permanent sensor poisoning could play a role in causing the baseline drift effects we observed here.

Restricted data subset

Based on our findings in Section 2.1.1 - 2.1.4, we selected a subset of the data that would be least affected by the drift effects we observed. We selected this subset by three constraints. First, only Methanol, Ethylene and Butanol were considered, since they have been measured within close temporal proximity (see Figures 2.1c and 2.2a). Second, we removed Sensor 4 from the analysis, as it appears to be particularly affected by drift (see Figures 2.2c & 2.2e). Third, we used data from sensor board 3 rather than sensor board 5, since our analysis suggested that it was, on average, less affected by drift (see Figures 2.2b & 2.2d).

We repeated the SVM classification task in this, according to our analysis, less compromised subset. The results are displayed in Figure 2.3d. The classification accuracy for the raw signal is initially still well above chance level at around 75%, without changing significantly after gas release. This indicates that long-term drift effects are pronounced enough to enable trial identification even in the restricted dataset. Moreover, classification accuracy increased only very slightly after gas onset. This indicates, paradoxically, that actual gas exposure made little difference for "gas" recognition in the restricted dataset.

The picture changed after compensating for long-term drift by subtracting the baseline offset at $t = 0\,\mathrm{s}$. Classification accuracy was only slightly above the chance level of 33.3% until gas release. After gas release, accuracy slowly increased to slightly above 60%. Therefore, we conclude that the restricted dataset is suitable as a gas classification benchmark when compensating for baseline offset. It should be noted though that a gas recognition accuracy of 60% is much lower than what we and others have reported for the original dataset. On the other hand, the sensor board we selected was located slightly lateral to the downwind axis from the source, therefore likely not as strongly exposed to the gas plume, which potentially affects classification performance negatively (but also apparently reduces sensor drift). A fair comparison between the results from subset and full-set is therefore non-trivial.

Discussion and Conclusion

In our analysis, we have shown that the different gases have been measured in time-separated batches and not in random order, which makes the data susceptible to sensor drift effects. We have shown that the sensor response baseline correlates with the time and order of measurement, consistent with long-term drift behaviour. We have also shown that the sensor response baseline alone is enough for 'accurate' gas classification, even after compensating for long-term drift by removing the offset at $t = 0\,\mathrm{s}$. This means that the dataset cannot be used for gas classification benchmarks without further precautions.

In an attempt to alleviate this limitation, we identified a subset of the dataset, which, under certain conditions, could be used for gas classification benchmarking. The subset contains three gases that have been measured in close temporal proximity, at a location that appears to be less affected by drift, while disregarding one most affected sensor. After applying long-term drift compensation to this subset, we observed what would be expected from a clean

experiment: Gas identification accuracy was near chance level at the beginning of a trial and rose only after the gas has reached the sensor. However, gas classification performance under those conditions was much lower than when using the full dataset, in spire of the reduced complexity of the task due to the smaller number of gases.

We therefore conclude that there is substantial information about the gas identity in the baseline. It must be assumed that this information will also interfere with actual gas sensor response. Therefore, the classification accuracy of a benchmark will overestimate the accuracy that could be obtained without the drift effect. These findings suggest that many, if not all, of the previous studies using this dataset overestimated the performance of their gas recognition algorithms. In fact, previous studies have acknowledged the exceptionally high classification accuracy obtained on this dataset, compared to others [Mon+16; Gam+21].

We identified 18 publications that are potentially affected [BBK18; VL16; IC20; CL19; Zho+19b; Mon+16; Fan+18; Gam+21; Zho+19a; KH20; MRS18; AGS19; Ver18; MG15; GSD16; CCP20; Wan+21a; Li+21]. At least one study reported a classification accuracy of 100% [AGS19]. None of those 18 studies described subtracting the baseline or other attempts that could address the detrimental effects of drift. Only four studies described data normalisation efforts [VL16; Ver18; BBK18; MRS18] that would remove the absolute scaling of the data, which should be a standard procedure when dealing with MOx sensor data. Since the different studies use different subsets of the dataset, and only a small fraction of the teams provided their analysis code [Gam+21; IC20], a thorough numerical comparison of their algorithmic performance is beyond the scope of this work.

Due to the popularity of this dataset as a benchmarks for electronic olfaction, the overestimation of accuracy that could be obtained may have distorted the state-of-the-art in gas recognition. The uncritical use of this dataset may have impeded progress in the field to a considerable extent, for example by casting unjustified doubt on other datasets where classification scores were lower (but possibly more relevant). We hope that the analysis presented here may enable a more realistic assessment of gas classification algorithms, and further encourage the collection and sharing of novel gas sensor datasets.

It should be noted that this dataset, despite its limitations, is an excellent example of how datasets should be shared. It contains the raw measurement data and all timestamps of the recordings. This is unfortunately not common practice in the field—often, only derived features are shared. We expressly acknowledge the effort Vergara et al. have made to share the data as accurately as possible. Only through their diligence and attention to detail was it possible to identify the underlying limitations.

The dataset still has unique features which make it a tremendous resource for machine olfaction research. It is one of the very few available datasets which have been recorded with a very high temporal resolution in a wind tunnel. Therefore, it includes temporal dynamics of odor concentration which are due to turbulent dispersal. This feature of the dataset has given rise to a study demonstrating that information about source proximity can reliably be extracted from turbulent plumes using metal oxide sensors [SBH16], which has been replicated independently [BM20b] and confirmed using newly recorded data [BM19]. Such studies are not affected by the adverse effects discussed in the present study, since they do not attempt to identify odorants, but focus only on the temporal dynamics of odorant concentration induced by turbulence, which is largely independent of odorant identity.

Our study highlights that it is still difficult to obtain clean and reliable data for gas recognition benchmarks. Besides the challenges in designing and manufacturing a gas sensor setup, planning a recording campaign robust against drift could hold its own pitfalls that may not be evident from the outset, and even go undetected for years after publication, even for highly cited datasets. Our findings highlight once more the importance of thoroughly checking the validity of third-party datasets before using it as a basis to develop algorithms for gas sensing.

A few recommendations emerge from our analysis towards best practices for designing MOx gas sensor datasets and sharing them. First and foremost, it is imperative to use a reference gas at short time intervals that will allow the identification and quantification of deviations in sensor response. Second, individual gases or mixtures should ideally be presented in a pseudo-randomized order, as should any varied parameter (e.g. wind speed, hotplate voltage). If a randomized presentation order is not feasible, one should record multiple batches for the same set of parameters at separate points in time. Training and testing data splits should then be selected from batches that were time-separated (as in [Asa+17]), which would allow for a more realistic performance evaluation. Finally, external parameters that could affect sensor behavior should be measured and reported, e.g. ambient temperature and humidity, and the exact time of the recording. Some MOx gas sensor datasets that implemented such principles are openly available [Fon+15; Ziy+15; BJM18].

Reliable data is the foundation for progress in the development of algorithms for gas sensing. The large number of citations of the original publication of the dataset analyzed here indicates that such data is much sought after and of high value for the community. It underlines the requirement for future efforts to record and publicly share gas sensing data for the progress of the field as a whole.

2.2 Limitations in neuromorphic algorithm for robust odour recognition

The groundbreaking study published by Imam & Cleland [IC20] introduces a bio-inspired and event-based algorithm for machine olfaction and proposes an implementation on the neuromorphic *Loihi* chip by Intel [Dav+18]. It distinguishes itself from other studies by promising one-shot odour learning and robust recall within milliseconds, which makes it most relevant to the topic of fast machine olfaction. In fact, it was this particular study that caught my personal interest in artificial olfaction in the first place, and therefore led me to start my doctorate researching the subject. After revealing the limitations of the dataset described earlier, it was a logical consequence to examine if and how strongly the proposed algorithm is affected by such limitations. After all, if demonstrated to be unaffected, the method would not only surpass similar approaches in accuracy, speed, and power consumption but also solve the long-standing problem of sensor drift.

However, the subsequent analysis revealed that the algorithm is indeed quite strongly affected by sensor drift. A deep dive into the methodology and existing codebase further revealed limitations in the evaluation protocol. In particular, it appears that the denoising-and-classification algorithm has been trained and tested on the same data instances. When tested on previously unseen data, the classification performance dropped drastically.

Summary of the original paper

The authors [IC20] introduce a novel algorithm that is inspired by the mammalian main olfactory bulb (MOB), particularly by the external plexiform layer (EPL). The method targets the problem of few-shot online learning and identification of odour samples under noise, uses gas sensor response traces as input data, and is implemented on the neuromorphic platform *Loihi* [Dav+18]. In the following, I will elaborate on the network architecture and plasticity rules, on the evaluation protocols, as well as on the reported results.

Network architecture and plasticity rules

The suggested network architecture is an elegant abstraction of the EPL, displayed in Fig. 2.4a. Some core principles of MOB computation are incorporated, which is i) the specific lateral-inhibitory topology of the EPL that includes columnar mitral cell (MC) and granule cells (GCs), ii) the discretised spike timing-based computation that resembles gamma-band oscillations in the EPL, iii) the global GC excitation and local MC inhibition, iv) the time-scale of the inhibition-induced MC spike delays that approximate these of gamma-band oscillations, v) continuous learning abilities, and vi) neuromodulation embedded in a dynamic optimisation framework.

The presented instance of the model is composed of 72 columns, each comprising a single two-compartment MC principal neuron and up to 50 inhibitory GC interneurons. MCs connect sparsely and globally/intercolumnar to GCs via excitatory synapses, while GCs connect densely and locally/intracolumnar to MCs via inhibitory synapses.

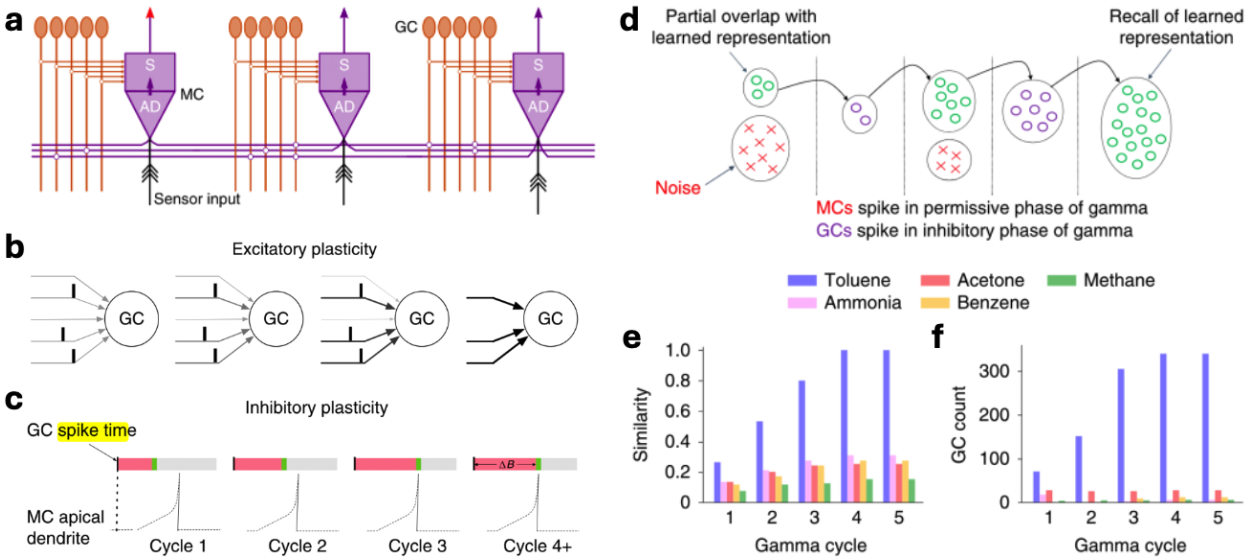

Figure 2.4: Overview of the methods and results demonstrated in Imam et al. [IC20]. **a)** Neuromorphic circuit architecture, consisting of Mitral Cells (MCs) and Granule Cells (GCs). **b)** During training, coincident MC spikes activate GCs, strengthening their synaptic weights while weakening and eliminating other inputs. **c)** During training, the duration of GC-mediated inhibition of co-columnar MCs (red bar) extends until its release (green) aligns with spike initiation in the MC apical dendrite. The learned inhibitory weights induce a blocking period ΔB, suppressing spike propagation in the MC soma. **d)** Illustration of iterative denoising for an occluded test sample. Partially correct MC representations activate some correct GCs, whose inhibition iteratively refines MC activity towards the learned representation. **e)** The Jaccard similarity to toluene, evoked by the occluded-toluene stimulus, increased over five gamma cycles, leading to its classification as toluene. For clarity, only five odourants are shown. **f** The number of toluene-tuned GCs activated by the occluded-toluene stimulus increased progressively over five gamma cycles as MC spiking patterns converged toward the learned toluene representation, with minimal recruitment of GCs tuned to the other nine odourants. Figure and caption adapted, with permission, from Imam et al. [IC20]

.

The model expresses oscillations to constrain MC spike timings, reflecting the cyclic inhibition seen in biological olfactory bulbs, e.g. the gamma-band (30 Hz - 80 Hz) oscillations found in the EPL layer. This oscillatory constraint enables sensory integration within permissive and inhibitory epochs, allowing the network to iteratively process sensory inputs over successive cycles.

During odour learning, spike-timing-dependent plasticity shapes GC-MC synapses, allowing GCs to become selective for specific MC spike patterns that represent higher-order stimulus features. In turn, GCs inhibit their respective MCs, adjusting MC spike times based on stimulus-specific patterns. Through this mechanism, the network learns to encode odour signatures as spatiotemporal patterns, which are refined across gamma cycles (see Fig. 2.4b&c).

During recall, plasticity is disabled and occluded inputs fed into the network. These activated a fraction of the GCs, which in turn modified the corresponding postsynaptic MC spike times. The networks topology allows for an iterative (across gamma cycles) minimisation of the distance between the inputs' network representation and the learned odour by progressively increasing the fraction of activated GCs (see Fig. 2.4d). With this, an occluded (or noisy) input can be restored if previously learned.

Evaluation protocol

The network was tuned to the particular gas sensor dataset [Ver+13] described earlier (see Section 2.1). The 72-columns match the number of gas sensors across all eight arrays distributed in the wind tunnel.

From the 20 repetitions per analyte gas and experimental parameters (here $v_{wind} = 0.21\,\mathrm{m\,s^{-1}}$, $V_{heater} = 500\,\mathrm{V}$, $d = 1.18\,\mathrm{m}$, location L4, constant odour concentration), one such repetition was selected and sampled at a single time point at $t = 90\,\mathrm{s}$. Each sensor response was discretized into 16 levels of activation, normalised by the full range across all stimuli. The discrete 72-element sensor vector was further sparsened by setting the smallest 50% of the values to zero. Collectively, the training set consists of 10 samples of a 72 times 4 bit data feature.

For most of the demonstrated evaluations, the testing set was created by altering samples from the training set. In particular, 60% of the values in each samples were replaced with impulse noise, i.e. uniformly distributed 4 bit values. This procedure was repeated 100 times for each instance of the training set, resulting in a total test set of 1000 samples.

After training, the networks' plasticity was disabled and the testing data was used to activate the MCs. The resulting representation of the network was observed across multiple gamma cycles. Ideally, the similarity between the networks' response to a noisy stimulus and the corresponding noise-free stimulus should increase progressively as the number of activated GC increases. This is evaluated by tracking the Jaccard similarity index [LW71] across gamma cycles (see Fig. 2.4e&f).

Generally, the Jaccard similarity index, J, between two sets A and B is defined as:

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|} \tag{2.3}$$

where $|A \cap B|$ is the size of the intersection of A and B, and $|A \cup B|$ is the size of their union. Here, it is defined as the number of spikes in the intersection of two odour representations, divided by the number of spikes in their union. Where applicable, test samples were classified as a known odourants if the Jaccard similarity exceeded a threshold of 0.75 in the fifth gamma cycle.

Results

Given the described evaluation protocols, the authors report on reliable odour identification despite strong destructive interference. Remarkably, this is shown in a one-shot scenario, with only one training sample per class. These one-shot classification results were favourably benchmarked against various machine learning methods; among these Principal Component Analysis (PCA) and a seven-layer dee[autoencoder. Further, the authors reported on successfully overcoming catastrophic forgetting, which is achieved by continuously adding new GCs and thus emulating neurogenesis mechanisms known in mammals [Mor+09]. Finally, the proposed network was implemented and evaluated on Intels' neuromorphic platform Loihi [Dav+21]. The authors demonstrate that processing a single sniff requires around 2 ms, which is five times faster than the sensors' sampling rate and thus would allow for continuous and online operation.

The remainder of this section is adapted from the peer-reviewed publication, with permission from *Springer*:

Nik Dennler, André van Schaik, and Michael Schmuker. "Limitations in odour recognition and generalization in a neuromorphic olfactory circuit." In: *Nature Machine Intelligence* 6 (2024), pp. 1451–1453. DOI: 10.1038/s42256-024-00952-1

The co-authors—in the following abbreviated as "we"—contributed as follows to this work: A.v.S. co-supervised the project. M.S. co-supervised the project, and co-edited the draft of the manuscript. All co-authors assisted in editing of the final manuscript. My contributions were the following: Replicating the original analysis and modifying it to shed light on potential limitations. Identifying the limitations in the algorithm. Performing a throughout analysis, including proposing a new algorithm that challenges the suggested evaluation metrics. Drafting and editing the manuscript. Editing and re-vising the manuscript during peer-review.

Abstract

Neuromorphic computing is one of the few current approaches that have the potential to significantly reduce power consumption in Machine Learning and Artificial Intelligence, and has drawn vast inspiration from considerations of biological systems and circuits. In their work, Imam & Cleland presented a neuromorphic odour-learning algorithm that is inspired by mammalian olfactory bulb circuitry, which they assessed by considering its performance in "rapid online learning and identification" of gaseous odorants and odorless gases (short "gases") using a set of gas sensor recordings of different odour presentations and corrupting them by impulse noise. We replicated parts of the study and discovered limitations thereof, which are 1) that the dataset used suffers from sensor drift and a non-randomised measure-

ment protocol that render it of limited use for odour identification benchmarks, and 2) that the model is restricted in its ability to generalise over repeated presentations of the same gas. Therefore, a validation of the model that goes beyond restoring a previously learned data sample remains to be shown, in particular its coherence with the attributed capabilities of robustness and broad generalisation beyond experience, as well as its suitability to realistic odour identification tasks.

Introduction

Imam & Cleland's [IC20] algorithm takes inspiration from the neural pathways of the external plexiform layer of the mammalian olfactory bulb. Gas representations are built by an iterative approach of applying spike-time dependent plasticity rules to sequential gamma-frequency spike packages, on the basis of a dataset consisting of recordings from 72 Metal Oxide (MOx) gas sensors mounted in a wind tunnel [Ver+13] (Fig. 2.5a). They validate the model's capability to learn and robustly identify gases by computing and thresholding the Jaccard similarity coefficient between gas sensor recordings and representations arising from artificially occluded sensor recordings. Further aspects such as neuromodulation, contextual priming and neurogenesis are explored. The implementation and operation of the algorithm on Intel's Loihi neuromorphic platform [Dav+18] presents a major milestone in neuromorphic computing due to the high complexity and biological realism of the underlying network model. The authors claim to describe a gas identification framework that is superior to other models in terms of common classification metrics, that generalises "broadly" beyond experience and that can be deployed into environments containing unknown contaminants

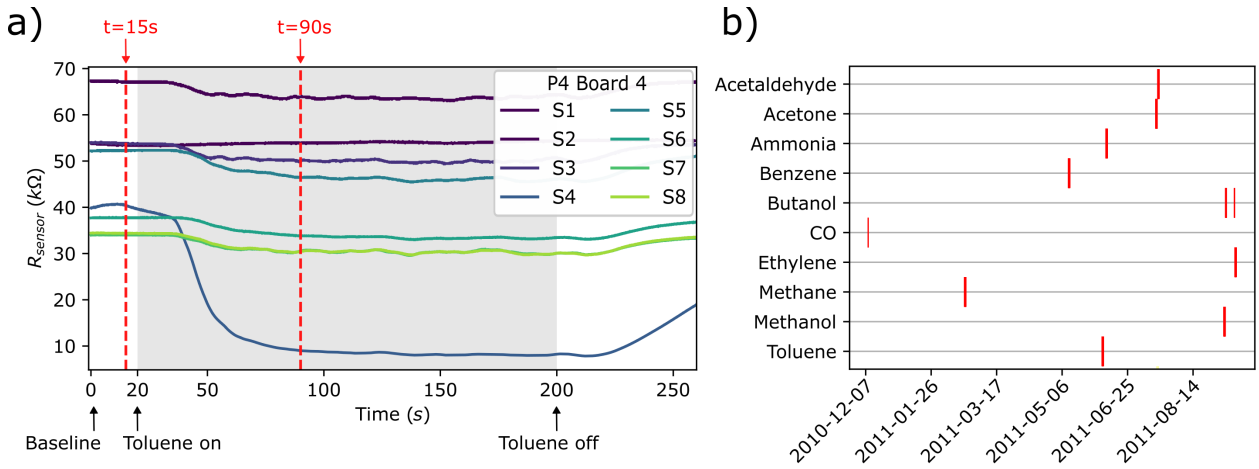

Figure 2.5: Overview of the data collection protocol. **a)** Example sensor resistance measurement for all sensors on one sensor board (location P4, module 4, Toluene, 0.21 m s^{-1} airflow velocity, 5 V operating voltage, trial 2). The shaded area denotes the period during which the gas is injected into the wind tunnel, the red dotted lines indicate where the data is sampled in [IC20] ($t = 90$ s) and in this work ($t = 15$ s and $t = 90$ s). **b)** Timestamps of the gas sensor recordings from which, in [IC20], one trial per gas was sampled. Each vertical line represents multiple trials (up to 20), which were performed too close to each other for them to be visually distinguishable in this representation. Adapted from [Den+22b].

.

and other sources of interference [IC20]. In addition, the study has been referred to as a demonstration on how a neuromorphic network can learn and discriminate odours [Dav+21; Chr+22; Dav21; Inta; Intb; Lef20]. Below we demonstrate limitations of the study that call these statements into question.

Results

Drift contamination of data set

The first limitation of the study relates to restrictions in the dataset used to validate the olfactory bulb network. The issue with the dataset [Ver+13] (Fig. 2.5a) has been described in detail in the previous section. In short: Recordings were acquired in gas-specific batches over the course of nine months (Fig. 2.5b), and not in a randomised order. The non-

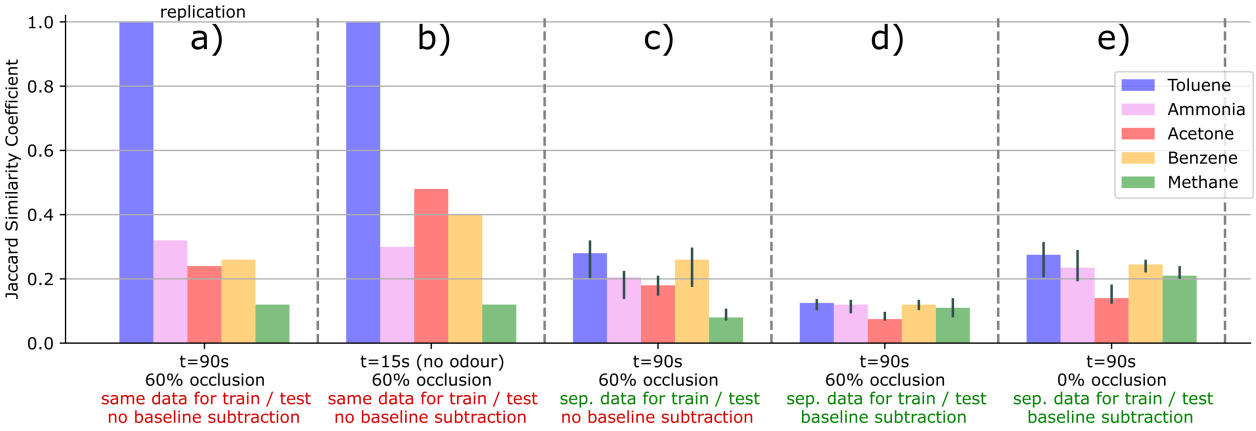

Figure 2.6: a) Jaccard similarity coefficient of the networks response to occluded Toluene and the learned odour representations, after five successive gamma cycles. Replicated from [IC20]. b) Spurious recognition in absence of gas in the wind tunnel. c,d,e) Recognition failure on different repetitions, c) without and d) with baseline subtraction, and e) without sample occlusion. Median and interquartile range across 10 Toluene representations are displayed, and only five out of 10 gases are depicted for clarity.

randomness, together with the dominating presence of sensor drift contaminations, allows for successful gas classification *before* the gas is presented, which renders this dataset largely unsuitable for classification tasks [Den+22b]. Drift contaminations could be partly mitigated by subtracting the baseline, i.e., the sensor response right before gas exposure, which is a widely used procedure when using MOx sensor data [HLG99]. No baseline subtraction was performed in the discussed study, suggesting that the reported findings about odour learning and recognition may be skewed, and potentially invalid, due to the limitations of the dataset.

We repeated the simulations for a range of conditions, using the publicly available code by the authors. As in the original work, the model was trained on 10 gases, and tested on 10 occlusions for each gas, i.e., 100 samples total. If not otherwise stated, 60% of the data was occluded by impulse noise when testing. We successfully replicated the authors' Jaccard similarity coefficient plot ([IC20], Fig. 4b), using the same raw data points for composing training and test sets, and sampling the recordings at $t = 90s$ (Fig. 2.5a, 2.6a). The result

appears to demonstrate robust recognition of the Toluene gas instance. Paradoxically, the same level of "recognition" of Toluene can be obtained in the absence of gas, using samples obtained at $t = 15s$, *before* the release of odour into the wind tunnel at $t = 20s$ (Fig. 2.5a, 2.6b). Therefore, the high Jaccard score for Toluene should be considered an artefact of sensor drift, and is unsuitable to substantiate a capability of the model to recognise odours [Den+22b].

Restrictions in capability to generalise beyond training data

In addition to the dataset's limitations, we found restrictions in the model's capability to generalise over different recordings of the same stimulus. In machine learning, generalisation refers to a model's ability to accurately make predictions on new, unseen data that is drawn from the same distribution as the training set. It is the capacity of a trained model to adapt and perform well on previously unseen examples, rather than simply memorizing the training data [Gar+13]. Generalisation is an important property of any pattern recognition system. The authors convincingly show that the model can restore input patterns corrupted by impulse noise. However, in most instances the authors tested recognition on the same sample that was used for training, occluding 60% of the sample with noise. 40% of each training sample were present unchanged in the corresponding testing sample. By using overlapping portions of data for both training and testing, the statistical independence required for a robust assessment of generalisation is not maintained. A real odorant recognition and signal restoration system would rarely encounter the exact same stimulus twice, once in a clean and once in a corrupted version. Therefore, assessing the model's capability to recognise and restore patterns from separate recordings is essential to judge its relevance.

For most gas and parameter combinations, the dataset contains 20 repetitions. We repeated the experiment using separate repetitions for training and testing and found that gas identity could not be recognised in occluded samples (Fig. 2.6c). Recognition scores were further reduced when subtracting the baseline from training and testing data (Fig. 2.6d). In this configuration, aimed at mitigating sensor drift, recognition across repetitions failed even for samples without any noise occlusion (Fig. 2.6e).

Conclusion

We conclude that the capability of the proposed model to identify learned odourants appears to be limited to corrupted versions of the training data. It failed to generalise to data outside the training set: Repetitions of learned gases were not recognised if these repetitions were not part of the training data. Imam & Cleland's model is an elegant example for an implementation of a biologically plausible model on neuromorphic hardware that can restore learned signals corrupted by noise. However, due to the restricted generalisation capability of the model, and to the limitations of the data used, it cannot be claimed that it solves the problem of odour learning and identification under a realistic scenario. We hope that raising awareness about these limitations paves the way towards improved neuromorphic models for robust gas recognition that can solve real-world odour recognition tasks.

2.3 Conclusion

In the first part of this chapter, I found that a widely used MOx electronic nose dataset has been recorded in temporally clustered and gas-specific batches, rather than by using an analyte-agnostic and randomised measurement protocol. As the drifting gas sensor response baseline correlated with those batches, it was possible to infer the analyte gas seconds *before* the sensors were physically exposed to the gas, thus artificially inflating any classification metrics. An exhaustive literature review of studies that used this dataset for benchmarking novel gas classification algorithms led us to conclude that the reported accuracies may not be reflective of the actual performance of the tested algorithms, likely skewing the state-of-the-art in gas classification. I compiled a set of best practices and and constructive suggestions to consider when performing data campaigns with MOx gas sensors.

In the second part, I examined one prominent algorithm study in neuromorphic olfaction that used the discussed dataset, and proposed to solve complex olfactory tasks at very short timescales. I demonstrated that the evaluation of the proposed algorithm is, in fact, affected by the limitations in the data and does not deliver positive results when tested on less affected data. Furthermore, I discovered issues in the evaluation itself: Testing was performed on a rendered subset of the training set, a practice that is known to artificially inflate performance metrics. I showed that the algorithm fails to generalise to data outside the training set, and discuss the implications thereof.

Chapter 3

Fast Odour Sensing with Electronic Nose

"Fasten your seatbelts; it's going to be a bumpy night."

– Bette Davis

While the limiting factors in both existing datasets and algorithms have temporarily stalled the progress in developing algorithms that rapidly recognise odours and gases, they inspired and motivated us to acquire our own datasets.

In this chapter, I will elaborate on two electronic nose data collection campaigns, which allowed for detailed analysis, brought many unexpected insights, and are anticipated to open new pathways in artificial olfaction. Both were based on electronic nose devices designed and constructed by Damien Drix (D.D.). In the first study, we explored how sub-second sensor heater cycles can be leveraged for recognising urban olfactory scenes. In the second—much more extensive—study, I collaborated with sensory neuroscience researchers from the Crick Institute in London, and collected electronic nose data using an olfactometer that was most recently used to reveal the temporal capabilities of the mammalian olfactory system [Ack+21]. The analysis and evaluation shed light on the fast nature of the electronic nose, which in certain cases even exceeds the temporal capabilities of the mammalian system.

3.1 Rapid odour classification using MOx sensor hotplate modulation

During his research, D.D. designed an electronic nose prototype based on MOx sensors, which can set heater temperature and read out sensor response values at one kilohertz. With this, he collected multiple datasets of different urban olfactory scenes. During the scope of this thesis, I have analysed two of them. The first was used to compare different asynchronous sampling strategies, which I will discuss in Section 4.2. The second was used to explore the viability of using sensor response windows, phase-locked with short (150 ms) heater cycles, for the discrimination of olfactory scenes, which is elaborated on in the following.

The remainder of this section is adapted from the peer-reviewed conference proceedings (© 2022 IEEE):

Damien Drix†, Nik Dennler†, and Michael Schmuker († denotes equal contribution). "Rapid Recognition of Olfactory Scenes with a Portable MOx Sensor System using Hotplate Modulation." en. In: *2022 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*. Aveiro, Portugal: IEEE, 2022, pp. 1–4. DOI: 10.1109/ISOEN54820.2022.9789654

The co-authors—in the following abbreviated as "we"—contributed as follows to this work: D.D. designed the electronic nose, wrote the acquisition software, collected the data, and drafted the the manuscript sections on e-nose design and data collection. M.S. supervised the project. All co-authors conceptualised the study and assisted in editing the final manuscript. My contributions were the following: Suggesting the use of short heater cycles for the analysis of the olfactory scenes, performing the formal analysis, including proposing novel feature extraction methods, algorithms, and evaluation methods, and drafting the manuscript.

Abstract

A café, the metro, a supermarket, a book store — many locations of everyday life have a specific smell. Recognising such olfactory scenes could inform personal activity tracking, environmental monitoring, and assist robotic navigation. Yet it is unclear if current Metal-oxide (MOx) sensor technology is sensitive and specific enough to achieve this. Factors like sensor drift, and sensitivity to ambient humidity and temperature further complicate the recognition of olfactory scenes. Hotplate temperature modulation has been suggested as a method to counter these drawbacks. We present an electronic nose based on MEMS-MOx sensors that support rapid hotplate temperature modulation with a 150 ms period. We recorded different natural olfactory scenes in an urban context. A linear SVM was able to recognise four olfactory scenes in single hotplate cycles with near-perfect performance when trained and tested on the same day, and 73% accuracy when tested in the same locations on the next day. Gas sensor responses yielded higher recognition accuracy than humidity, temperature, and pressure, which were also partly-location specific. Our results indicate that hotplate modulation enables recognition of natural odour scenes across extended timespans. These findings encourage the use of MOx-sensors as rapid sensing devices in natural, uncontrolled environments.

Introduction

The use of electronic noses (e-noses) has become popular in many areas, such as industrial and environmental monitoring [BCR00], hazard control [Sol+22], mobile olfactory robotics [LLD06], and medical screening [Loi+13]. Despite shortcomings like drift and cross-sensitivity to humidity, MOx sensors are an attractive choice for electronic noses due to their

low cost and availability. The current generation of sensors are made using MEMS techniques which enable faster modification of the sensing site temperature, and these fast modulation cycles have been shown to decrease the integration time and improve the specificity of the responses to different analytes [Ver+14].

The recognition of olfactory scenes is an interesting but challenging problem, as it requires a portable e-nose [Che+21] that can operate in uncontrolled natural environments susceptible to changes in temperature, humidity and pressure.

Electronic nose design

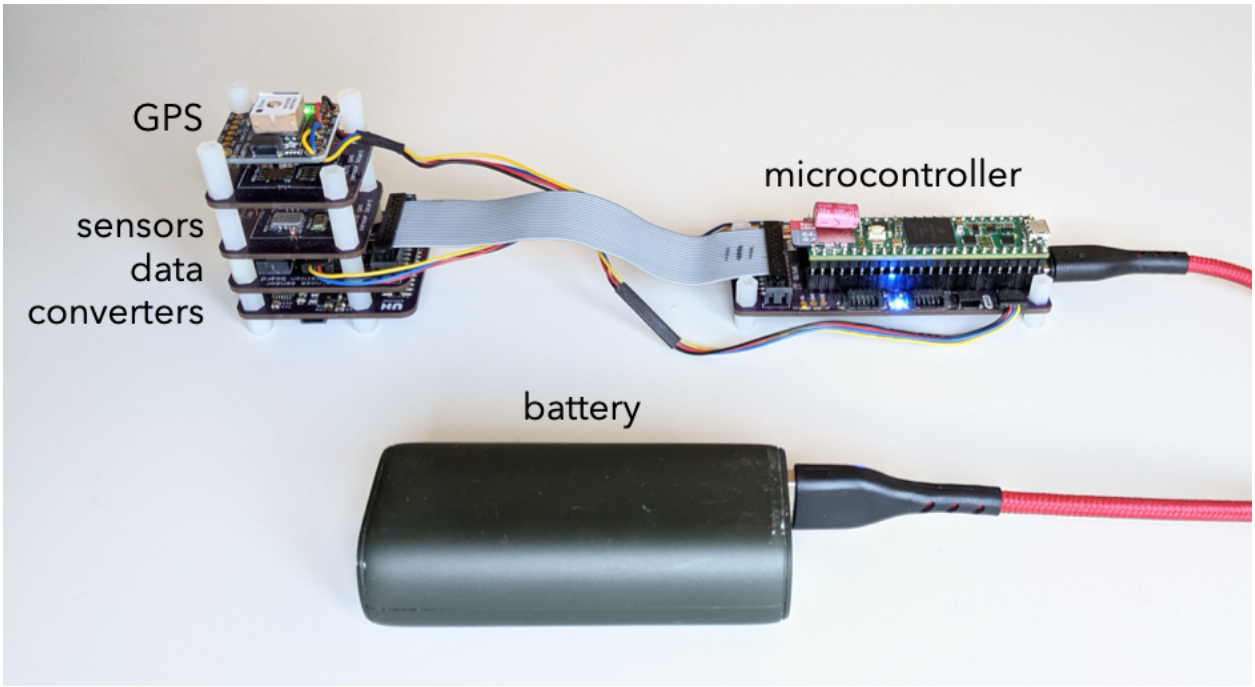

Figure 3.1: Overview of the electronic nose system, showing the sensor board, microcontroller board and power source.

We designed an electronic nose with three goals in mind: (1) to investigate heater modulation techniques, (2) to take advantage of MEMS gas sensors and their faster response times, and (3) to enable field recordings untethered to a computer. Our design uses off-the-shelf components and consists of two main parts: a microcontroller board based on a Teensy 4.1 microcontroller (PJRC.com) for data processing and storage, and a sensor board hosting the sensors, associated analogue circuitry, and data converters, in a portable unit (Fig. 3.1).

We use four metal-oxide sensors grouped in three MEMS packages: MiCS 4514 and MiCS 5914 (SGX Sensortech) and CCS801 (AMS/ScioScience). These can sense various kinds of reducing and oxidising gases including various volatile odor compounds (VOCs), hydrocarbons, carbon monoxide, hydrogen, nitrogen oxides and ammonia.

Heater modulation requires a way to measure the hot plate temperature, and a way to regulate the power delivered to the resistive heating element. In metal-oxide gas sensors with resistive hotplates, the heater resistance increases quasi-linearly with temperature. Absolute hotplate temperature can be estimated by continuously measuring the heater resistance, combined with calibration information obtained from the manufacturer's datasheet. We use a DAC (DAC7554, Texas Instruments) and an amplifier (TS924, STMicroelectronics) to set the voltage, and a sense resistor of known fixed value to measure the resulting current (Fig. 3.2). This enables control of the heater resistance, temperature, and power.

The gas sensor response fluctuates rapidly with the hot plate temperature, independently from the speed at which chemical reactions occur on the sensor surface (Fig. 3.3b). We therefore sample the hotplates and sensors synchronously, using two simultaneously-sampling, 8-channel ADCs (ADS131M08, Texas Instruments). These form a closed control loop with the DAC, running at 1 kHz. The high sampling rate is is required to support accurate control of the heating elements, which have thermal time constants on the order of 20 ms.

A GPS module provided position and time. One environmental sensor measured the temperature and humidity in close thermal proximity to the sensors (MS8607 from TE Connectivity, also providing barometric pressure). Another sensor measured the temperature and humidity of the surrounding air, uninfluenced by the heaters (SHT31 from Sensirion). The microcontroller recorded the data on-the-fly to an SD card at a data rate of 150 kB/s, enabling multiple-hour recordings.

Data collection

We acquired a dataset of natural olfactory scenes recorded in different urban indoor locations: 'Café', 'Metro', 'Bookstore' and 'Supermarket'. Each location was visited once per day, on two consecutive days. We recorded the four gas channels, ambient temperature, relative humidity, and pressure. Heater power was modulated with a period of 150 ms as described in Fig. 3.3.

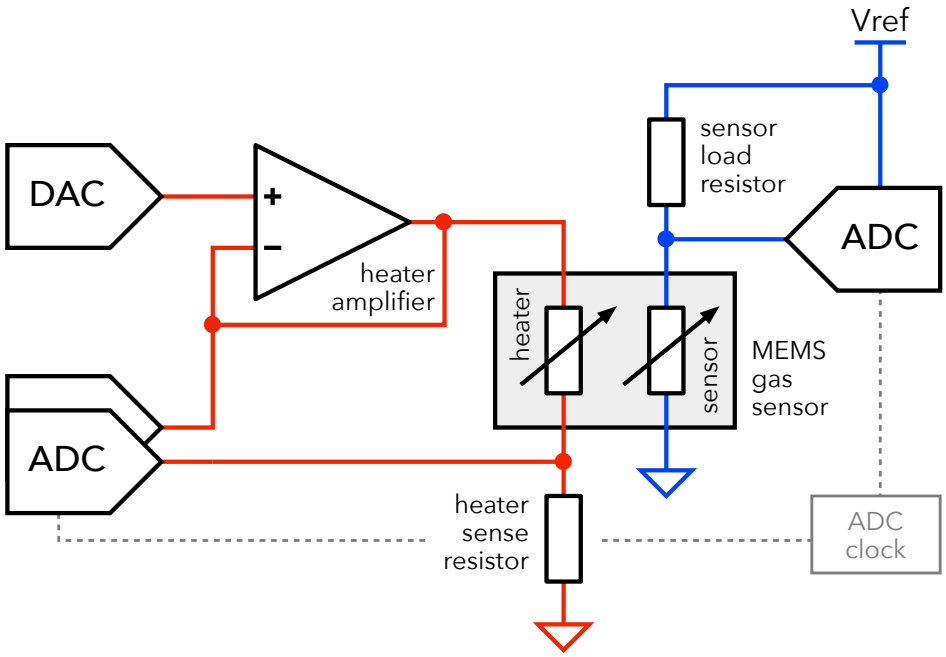

Figure 3.2: Simplified schematic of the sensor board showing one out of four sensor channels. Each channel contains its own heater circuit (in red) and gas sensor circuit (in blue). The heater amplifier tracks the DAC voltage (gain=1) and supplies the required heater current (up to 35 mA). Crosstalk was minimized by spatially separating the electrical components, by avoiding parallel PCB traces, and by implementing two separate ground connections.

The data was divided into training and testing sets with a four-fold cross-validation procedure. For each day and location, we selected a contiguous block containing 25% of the data and randomly picked 400 heater cycles from this block for testing. We then drew 1200 training cycles from the remaining 75% of the data. The whole procedure was repeated four times with non-overlapping test blocks, to allow the models to be validated on the entire dataset. This yielded a total of 12800 cycles divided in four sets: training (day 1), testing (day 1), training (day 2) and testing (day 2).

Inference of olfactory scenes from baseline-normalised sensor cycles

We then investigated whether one can recover the label of the olfactory scene from the time course of a single 150 ms heater cycle, either from the gas sensor conductance or from the environmental sensors.

First, the gas signal was normalised to the same minimum and maximum values over each heater cycle (Fig. 3.4a). This mitigates the baseline drift often seen in gas sensor datasets [Ver+], which is problematic for classification tasks [Den+22b]. While this normalisation step yields curves that, at first sight, look very similar across locations, they differ substantially and reproducibly if one compares them closely (see insets in Fig. 3.4a). Observing the 2D projections in the space of principal components (Fig. 3.4b), separable and class-aligning clusters emerge not only for for single day recordings, but also for the entire data spanning two days.

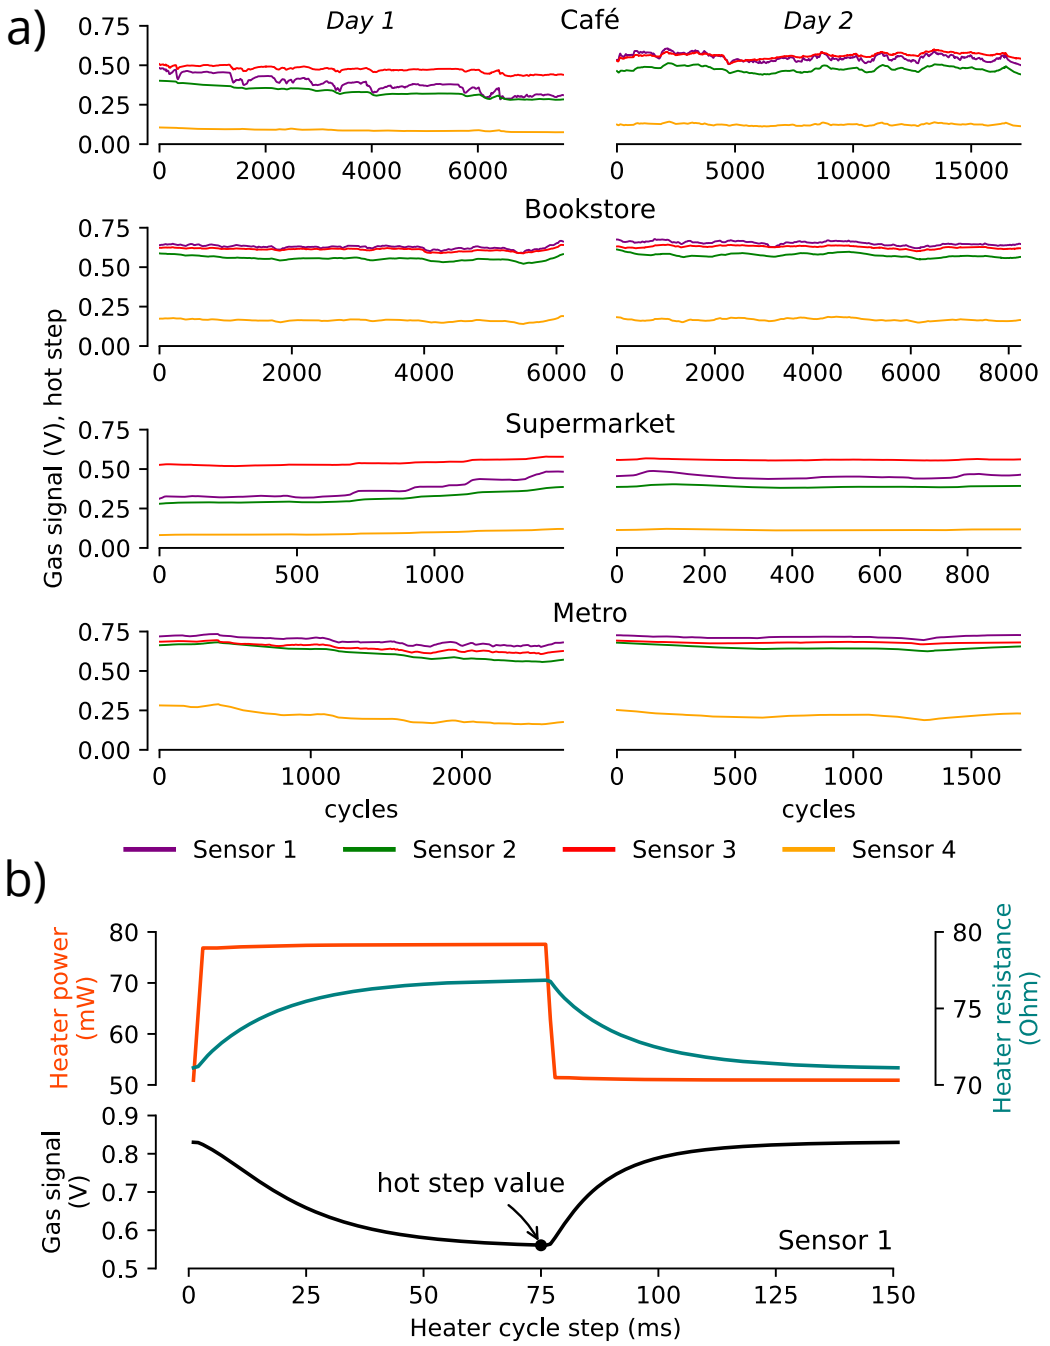

Figure 3.3: Overview of the gas sensor data. a) Gas sensor response at the hottest part of each heater cycle, across four locations visited on two different days. b) Schematic of the heater modulation cycle. Each cycle steps between two fixed power levels (different for each channel) with a 150 ms period. Hot plate temperature is inferred from heater resistance. Volatiles and environmental factors affects the shape of the temperature-induced response of the gas sensors.

If one considers the environmental sensor responses (Fig. 3.4c), it is evident that at each location there is a drift in temperature and humidity, causing a strong overlap of the responses across the classes, making a separation challenging. While pressure is more stable for each class at a given day, there appears to be an offset when comparing different days.

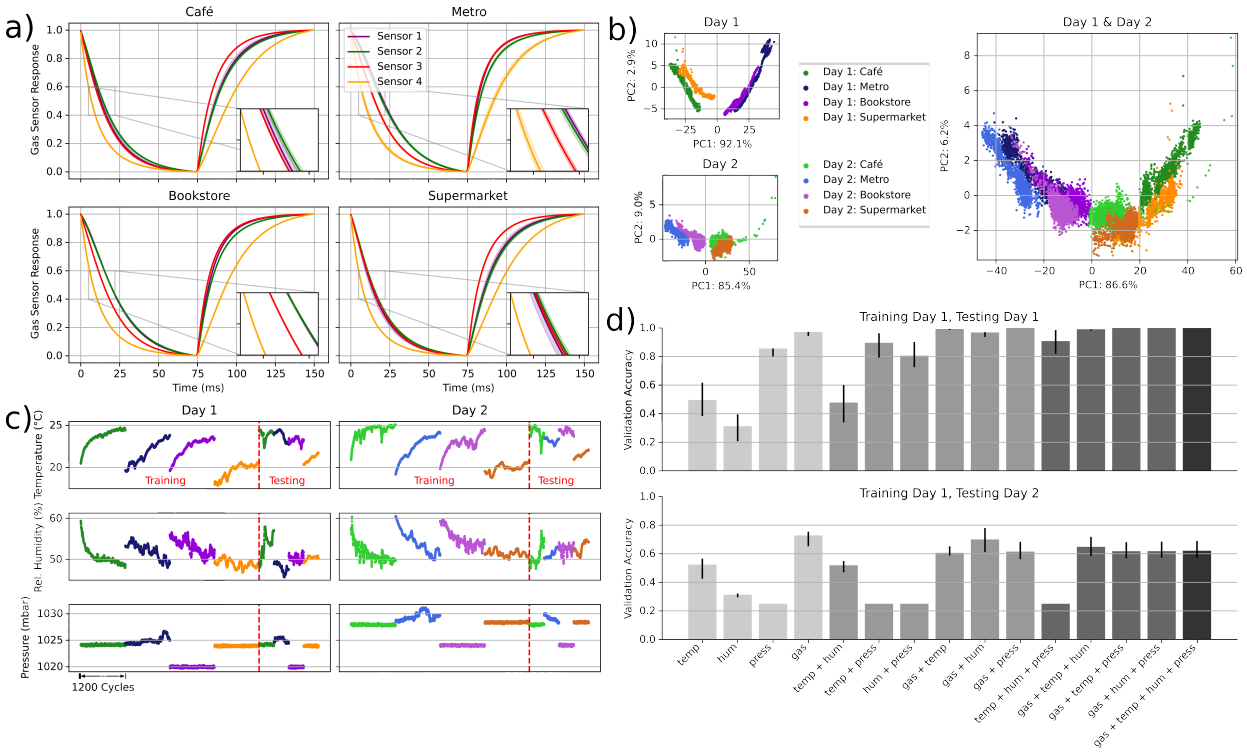

Figure 3.4: Analysis results. a) Normalised gas sensor cycles in four locations. b) PCA projection of sensor responses. Left: separate days, right: both days in joint PC space. c) Environmental variables of one cross-validation train-test split overlap for the different classes (one value per cycle). d) Validation accuracies for an SVM classifier, for different sensor modalities and their combinations.

Those qualitative observations were confirmed by evaluating a linear classifier trained on a subset of the data. For the different sensors (temperature, humidity, pressure, gas), all possible combinations were formed and the respective day-1 training set was used to train a soft-margin, linear-kernel Support-Vector-Machine (SVM) [CV95]. Each model was then validated via the classification accuracy achieved on the day-1 testing set and on the day-2

testing set (Fig. 3.4d). When training and testing on data from the same day (day 1), the model trained on the gas sensor responses achieves the highest accuracy scores by a large margin (97.2% validation accuracy vs. 49.6%, 31.2% and 85.6% for temperature, humidity and pressure respectively). When training on day 1, and testing on day 2, again the gas sensor results surpass the ones achieved by having the environmental sensors only (72.8% validation accuracy vs. 52.3%, 31.3% and 25.0%). In both cases, there seems to be no obvious advantage in combining the gas sensors with any of the environmental sensors.

Conclusion

In this work, we presented a novel design for a portable e-nose, which takes advantage of the properties of state-of-the-art MEMS gas sensors, allowing high-frequency heater control and sensor readout. We demonstrate that the phase-locked sensor response relative to heater cycles constitutes a promising feature for classifying natural olfactory scenes from sub-second samples. Further, we showed that the information present in the gas sensor signal does not appear to be fully explained by cross-sensitivities to the ambient air temperature, humidity and pressure: each location's unique olfactory signature seems to play a role as well. Future work should focus on evaluating the reproducibility of olfactory scene recognition across a wider range of conditions, and over longer intervals of time.

3.2 High-speed odour sensing using miniaturised electronic nose

The impressive efficacy of using sub-second heater cycles motivated a comprehensive follow-up study, characterised by enhanced control over experimental parameters—requiring levels of precision attainable solely within a laboratory setting. To achieve this required to establish a collaboration with researchers from the Crick Institute in London, which most recently published an influential study on evaluating the temporal resolution of the mammalian olfactory system [Ack+21]. An essential part of this was a customised odour delivery device with particularly high temporal signal fidelity.

An optimised and more compact successor of the previously demonstrated electronic nose was designed by D.D. While it consists of identical or similar sensing elements and periphery, a different microcontroller allowed for better control of the recording protocols. In the experiments, sensor heater modulations of either short cycles (50 ms) or a constant high temperature were deployed, then a set of odours at different temporal patterns was analysed. The outcomes of the study were equally exciting and surprising: We demonstrated that—for the first time—it was possible to match the temporal resolution of mammalian olfaction in robotic systems.

The remainder of this section is adapted from the following peer-reviewed publication:

Nik Dennler, Damien Drix, Tom PA Warner, Shavika Rastogi, Cecilia Della Casa, Tobias Ackels, Andreas T Schaefer, André van Schaik, and Michael Schmuker. "High-speed odor sensing using miniaturized electronic nose." In: *Science Advances* 10.45 (2024), eadp1764. DOI: 10.1126/sciadv.adp1764

The co-authors—in the following abbreviated as "we"—contributed as follows to this work: D.D. designed the electronic nose, wrote the acquisition software, co-conceptualised the experiments, aided in collecting and analysing the data, and drafter parts of the methodology. T.W. assembled and prepared the odour delivery device, provided codes for deploying odour stimuli, and drafted parts of the methodology. S.R. and C.D.C. aided in collecting the data. T.A. drafted parts of the methodology. A.S., A.v.S., and M.S. acquired funding and supervised the project. All co-authors assisted in editing of the final manuscript. My contributions were the following: Co-conceptualising the experiments, leading the collection of the data, performing formal analysis of the data, creating visualisation, drafting the manuscript, as well as editing and revising during submission and peer-review.

Abstract

Animals have evolved to rapidly detect and recognise brief and intermittent encounters with odour packages, exhibiting recognition capabilities within milliseconds. Artificial olfaction has faced challenges in achieving comparable results — existing solutions are either slow; or bulky, expensive, and power-intensive — limiting applicability in real-world scenarios for mobile olfactory robotics. Here we introduce a miniaturised high-speed electronic nose; characterised by high-bandwidth sensor readouts, tightly controlled sensing parameters and powerful algorithms. The system is evaluated on a high-fidelity odour delivery benchmark. We showcase successful classification of tens-of-millisecond odour pulses, and demonstrate temporal pattern encoding of stimuli switching with up to 60 Hz. Those timescales are unprecedented in miniaturised low-power settings, and demonstrably exceed the performance observed in mice. For the first time, it is possible to match the temporal resolution of animal olfaction in robotic systems. This will allow for addressing challenges in environmental and industrial monitoring, security, neuroscience, and beyond.

Introduction

The sense of olfaction is found all across the animal kingdom, and is crucial for survival and guiding behaviors such as navigation, food detection, predator avoidance, and mate selection [BK04; Ste+05; HDC06; CW08; KSB12; Sul+15; Kha+21]. Success in these tasks often hinges on the ability to swiftly and accurately detect and recognise scents [RAH08; BD14; Van+15; Dem+20], particularly when dealing with odour plumes characterised by brief and intermittent encounters[MC94; Vic+01] generated by turbulent dispersion processes [MM91; JMJ02; CMC18]. Concentration fluctuations in odour plumes can exceed 100 Hz [Yee+95b], while individual odour encounters can last single milliseconds or less [CVV14] (see Fig. 3.8a). Many environmental cues are embedded in the fine structure of the odour plume [Hop91; MC94; SBH16], which various organisms have evolved to use for their advantage. For instance, *Drosophila* olfactory receptor neurons can transduce odours in less than two milliseconds and resolve odour stimuli fluctuations at frequencies exceeding 100 Hz [Szy+14]. Similarly, honeybee projection neurons decode odour identity in tens of milliseconds after stimulus onset [Kro08], while mosquitoes can identify CO_2 packets of just 30 ms [DC11]. A recent landmark study in mice has revealed their ability to discriminate rapid odour fluctuations, enabling them to distinguish temporally correlated from anti-correlated odours at up to 40 Hz, which facilitates source separation in complex environments [Ack+21].

Research on mobile olfactory robotics [LLD06] has flourished over the last decade; driven by promising applications and solutions across various domains [Fra+22], and bootstrapping on the well established field of artificial olfaction[Cov+21]. The latter has demonstrated its effectiveness in domains where static and slow measurements are sufficient, such as the detection of hazardous gases or pollutants [SJR86], spoilage alert systems [Mai+06], health

monitoring [Far+19], and food sciences [Jun+23]. However, many recent applications call for unmanned ground or aerial vehicles (UGV / UAV) to perform odour source localisation and navigation tasks [KR08; Bur+19; JMI21; Fra+22], which rely heavily on sensing the environment fast and efficiently, considering plume dynamics [Kad+22].

Typically, mobile olfactory robots incorporate electronic noses; devices that are characterised by arrays of multiple gas sensors and associated peripheral electronics [PD82]. They offer distinct advantages over conventional analytical methods such as Photoionization Detectors (PID) and mass spectrometers (MS), notably in terms of portability, power efficiency, cost-effectiveness, and sensitivity to a wide range of odours and volatile compounds. The most widely employed sensing components are Metal-oxide (MOx) gas sensors [RBW08], which offer the significant advantage of a sensing layer that can be tuned through 1., modifications to its chemical structure, and 2., variations in operating temperature achieved by local heating, allowing for effectively detecting a diverse range of analyte classes. Their minimal requirements for electronic peripheral components streamline sensor design, lower costs, and conserve valuable space. Further reductions in latency, form factor and power consumption were enabled through latest MEMS-based MOx sensors [Liu+18; GGU23], facilitating seamless integration into electronic circuits [RHB18].

However, the relatively slow response and recovery times of MOx sensor electronic noses pose challenges for widespread adoption, and are prohibitive for many potential robotic applications [Wan+22a]. For this reason, various studies have investigated sensor response times and tried to improve them. Recent advancements in both hardware [Ver+14; Mar+14b] and software [MGB12; Di +14; SBH16; MBM19; BM19; BVM19; Xin+19; DS21] have significantly reduced response and recovery times from the orders of hours or minutes [PLT12a] down to

tens-of-seconds or seconds [BVM19; DS21]. Nevertheless, those timescales remain orders-of-magnitudes slower than what's observed for olfactory sensing in animals, potentially stalling progress on critical challenges in tracking of greenhouse gas emissions [Dom+24], ecological and environmental monitoring [BM20a; Ter+24], aerial-based wild fire detection [Wan+23] disaster management [Fan+19], and more.

In this work, we are pushing the limits of artificial olfaction with a high-speed, miniaturised electronic nose that can resolve odour pulses in the millisecond range. We propose an integrated e-nose design of MEMS-based MOx sensors and fast sampling periphery, as well as a set of powerful algorithms for control, sensing, and signal processing. We demonstrate the systems ability to operate at unprecedented temporal timescales when classifying short odour pulses, as well as when discriminating temporal characteristics of rapidly switching odour pairs. The challenge of deploying rapid and complex odour stimuli in a controlled and precise fashion [MSA21] is overcome by using a high temporal-precision olfactometer setup, which most recently has been used for showcasing the temporal odour recognition capabilities in mice [Ack+21; Das+22].

In the first part, we elaborate on the material and methods used. In particular, we discuss the electronic nose design including sensor heater modulation heuristics and responses to ambient air, the odour delivery setup including a description of the olfactometer and its calibration, the experimental protocol, and the analysis methods.

Afterwards we discuss the results, where we initially elaborate on the proposed design of the electronic nose and the feedback control methods, with which we achieve thermal response times that allow for ultra-fast heater cycles — orders-of-magnitudes faster than what is suggested in the literature. Later, we show that the electronic nose can successfully classify

the odour of short pulses, with durations down to tens of milliseconds. This is achieved by rapidly switching the sensor heater temperature, then extracting phase-locked data features to train machine learning classifiers. Further, we demonstrate the system's ability encode and infer temporal features in a task involving rapidly switching odour pairs, up to modulation frequencies of 60 Hz, which we show to match and even exceed the demonstrated capabilities of mice on equivalent tasks [Ack+21]. This is achieved by controlling the heater temperature to be constant, permitting for sensor response feature extraction from the frequency domain.

Finally, we discuss the results and its implications, and identify some example use cases that may benefit highly from using fast sensing modalities.

Methods

Electronic nose

Circuit board design: The electronic nose uses readily available components and is illustrated in Figure 3.5a. It features a *Raspberry Pi Pico* microcontroller and incorporates eight MEMS-fabricated MOx gas sensors of four different types, grouped into four packages.

The sensor packages comprise two *SGX Sensortech* MiCs-6814 sensors (sensors 1 and 5) and two *ScioSense* CCS801 sensors (sensors 2,3,4 and 6,7,8) , capable of detecting a wide range of reducing and oxidising gases, including volatile odour compounds (VOCs), hydrocarbons, carbon monoxide, hydrogen, nitrogen oxides, and ammonia. Figure 3.8c displays an optical microscopy image of the sensor structure. For sensing, the DC resistance across the sensing element is measured. The integrated micro-hotplates allow for operation at temperatures of up to 500 °C.

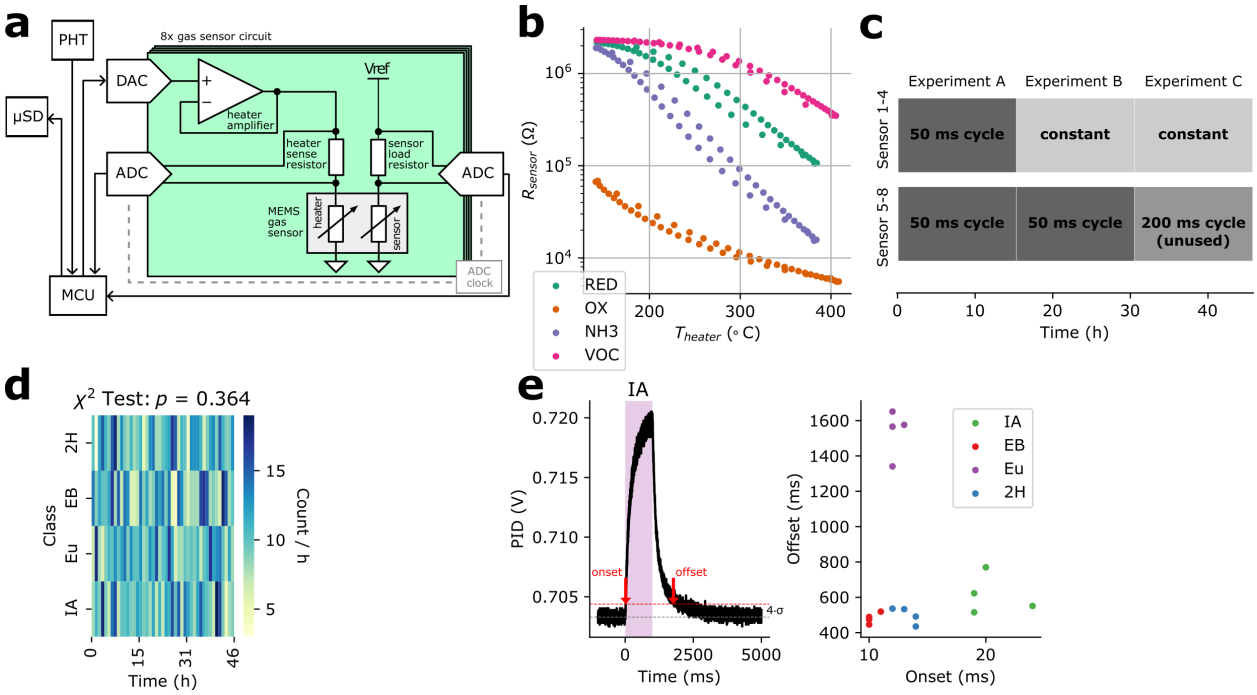

Figure 3.5: Supplementary figure for experimental setup. **a**, Electronic nose design, displaying how the microcontroller unit (MCU) sets and reads out the sensor heaters in a closed loop, while reading out the analyte dependent sensor resistances. Further, the MCU connects to an environmental sensor (PHT) and a micro SD card. **b,** R-T curve of a 50 ms temperature cycle between 150 °C and 400 °C without external stimulus, displaying how the sensor response closely follows the hotplate temperature. **c,** Different sensor hotplate settings over time. For each experiments, all the stimuli were presented in randomised order. **d,** Heatmap depicting the distribution of odour presentations over a set of 1 hour time intervals. A χ^2 test was performed to assess the randomness of class distribution over time intervals, with the computed p-value indicated as 'p'. **e,** PID response to a 1 s isoamyl acetate pulse. Grey-dotted and red dotted lines denote mean of pre-stimulus baseline and 4 standard deviations threshold respectively. Where the response crosses the threshold upwards (downwards), the odour onset (offset) is registered. **f,** Extracted odour onsets (w.r.t. $t = 0$ ms) and offsets (w.r.t. $t = 1000$ ms) for 1000 ms pulses of different odours. For all experiments, the odourants are abbreviated as follows. IA: isoamyl acetate; EB: ethyl butyrate; Eu: cineol; 2H: 2-heptanone; blank: odourless control.

For controlling the micro-hotplates, eight operational amplifiers (2x *STMicroelectronics* TS924) are employed, and the sensor heater voltage is configured using two Digital-to-Analogue Converters (DACs), specifically the *Texas Instrument* DAC60004, which offer four channels each at 12 bits and 1 kHz. Additionally, two Analogue-to-Digital Converters (ADCs), the *Texas*

Instrument ADS131M08, are used to read sensor and heater resistances. These are differential, simultaneous-sampling ADCs which read out the eight gas channels and the eight temperature channels in lockstep at 24 bits and 1 kHz (the two ADCs share the same clock). To monitor environmental conditions, a digital pressure-humidity-temperature sensor, the *TE Connectivity* MS8607, is included, which samples data at 24/16/24 bits and 50 Hz. Real-time data logging is facilitated through the inclusion of a microSD card. The device's power needs, ranging from 1.2 W to 1.5 W, allow for multiple days of continuous operation on a pocket-sized battery pack, making it suitable for extended field recordings or robotic environments.

Sensor heater modulation: To implement controlled heater modulation, continuous measurement of the hotplate temperature and regulation of power delivered to the resistive heating element are essential. Each heater voltage V_{heat} was adjusted using a DAC and an associated amplifier, while the resulting current I_{heat} was monitored using an ADC in conjunction with a fixed-value sense resistor R_{sense}. From these two quantities one can compute the heater resistance $R_{\text{heat}} = V_{\text{heat}}/I_{\text{heat}}$ and dissipated power $P_{\text{heat}} = V_{\text{heat}} I_{\text{heat}}$. Because the device did not directly measure V_{heat}, the resistance calculated by substituting the known control and sense voltages $V_{\text{dac}} - V_{\text{sense}} \approx V_{\text{heat}}$ was subject to errors, of which transient errors caused by lag and settling time in the DAC and amplifier were deemed the most significant. These affected the sample acquired immediately after a change in control voltage; therefore we used a Kalman filter [WB+95] to estimate R_{heat}, setting the measurement uncertainty proportionally to the rate of change of the control voltage V_{dac}.

The kind of resistive heating element present in our design exhibits a quasi-linear relationship between hotplate resistance and temperature [Ras+10], so we used a linear model to map the heater resistance to a calculated hotplate temperature. The parameters of that linear model were set before recording the data by measuring the heater response to a series of power steps, and matching it with calibration data from the manufacturer's datasheet, namely the nominal hotplate temperature delta above ambient air temperature at nominal heating power. This calibration procedure was repeated on an approximately weekly basis to account for the possible ageing of the sensors.

Achieving short temperature steps of a duration not much greater than the thermal time constant of the hotplate presented a number of challenges. Because the shortest steps consisted of only 25 samples (25 ms at 1 kHz), we employed a combination of open-loop and closed-loop control. In that scheme, we operated the heater at a constant voltage for the duration of each step. That voltage was selected based on a learned mapping between a desired temperature change, and the control voltage required to achieve it. We adjusted the mapping after every step to compensate for the effects of airflow and ambient temperature fluctuations, using a proportional controller with a relatively slow adaptation rate ($0.1\,\mathrm{V\,°C^{-1}\,s^{-1}}$). With this, fast and repeatable temperature modulation patterns can be obtained without introducing artefacts due to e.g. control loop oscillations.

For experiments with a constant heater temperature, achieving fast temperature changes was not an issue, but care was taken to avoid the artefacts caused by the DAC's 12-bit quantisation of applied heater voltage. These quantisation steps of about 0.7 mV led to small but measurable transients in the recorded sensor signal. Since these transients were in

the same frequency band as the signals of interest, we decided to also keep the heater voltage constant during each stimulus. We adjusted it to eliminate the temperature error after each stimulus, which was sufficient since the thermal environment of the sensors changed only slowly.

Sensor responses to ambient air: The resistance of MOx sensors depends not only on the presence of gases, but also on the operating temperature. In heater cycle mode and odourless air (see Fig. 3.9b, left), the sensors exhibited nearly-exponential relationships between the recorded sensor resistances and the hotplate temperatures in the range of 200 °C and 400 °C, with deviations at lower temperatures (see Fig. 3.5c). A small deviation can be observed between the trajectories corresponding to heating and cooling, however this may be attributed to uncertainties in estimating the temperature of the sensor, which is close to but not necessarily equal to that of the hotplate. The resistances returned to their initial values after a completed cycle, without significant hysteresis.

Odour delivery setup

Reagents: All odourants were obtained in their pure liquid form from Sigma-Aldrich, and contained in 15-ml glass vials (27160-U, Sigma-Aldrich). The odorants ethyl butyrate, isoamyl acetate and cineol were diluted 1:5 with mineral oil, while 2-heptanone was diluted 1:20 with mineral oil. 2-heptanone was diluted to a lower concentration as a 1:5 dilution saturated some of the gas sensors.

Olfactometer: Odours were presented using a custom made olfactometer capable of constructing temporally complex stimuli with temporal bandwidths of up to 500 Hz (Fig. 3.8d). This temporal olfactory delivery device (TODD) has been outlined previously[Ack+21; Das+22]. The device consisted of 8 independent channels which contained either odour (diluted with mineral oil) or pure mineral oil. These 8 channels were grouped into two sets of 4. Each set of four consisted of an odour manifold, which contained odours or pure mineral oil in glass vials which were fed by a common air flow. Each channel in this odour manifold was fed into its own high speed valve on a separate valve manifold. Each high speed valve could be opened and closed at frequencies of up to 500 Hz. On each valve manifold, one of the channels containing pure mineral oil was set to remain open indefinitely, acting as a 'carrier' valve (grey valves in Fig. 3.8d). When a trial was triggered, this carrier valve flow was reduced in accordance with the amount of additional airflow generated by the other valves on the manifold, therefore maintaining a continuous rate of air flow through the system. In some cases, the carrier valve was not simply reduced, but was used to generate temporally complex airflow to compensate for the temporal patterns generated using the other valves in the system. Signals to the valves were convolved with a high frequency 500 Hz continuous signal, referred to as 'shattering'. This shattering was included as it has been previously found to improve the temporal fidelity of the resultant odour signal. The airflow to the TODD was maintained at a rate of 1 L/min using a custom closed loop-feedback flow controller.

Calibration: To ensure a continuous total airflow whilst maintaining a high signal fidelity, prior to the electronic nose recording session, the output of the olfactometer was measured with both a PID (200B miniPID, Aurora Scientific) and flowmeter (AWM5101VN, Honeywell). The PID was positioned a short distance away from the output of the olfactometer (>

2cm) and calibration trials were presented. These were selected in a way to making sure to cover all the different valve combinations of the experiment. The PID response to presented odours was measured and the fidelity estimated. If the fidelity was found to be too low, the rate of flow into each channel was tuned to increase the fidelity. Next, the PID was replaced with the flowmeter, and the same selected calibration trials were presented. If the rate of flow varied during the trial presentation, the compensatory flow or carrier flow was modulated to return the net flow back to pre-trial levels. The flow through the odour valve was kept constant so as to not alter the odour signal fidelity. Airflow was modulated by altering the duty-cycle of the valve shattering. Once there was no visual change in the rate of flow between the trial and pre-trial levels the flowmeter was removed and the olfactometer was deemed to be calibrated.

Fidelity Calculations: For quantifying the olfactometers' temporal fidelity after calibration, we deployed single-odour pulse trains of different frequencies and obtained simultaneous PID and flow meter recordings. Here, the odourant Ethyl Butyrate was used, as its ionisation energy is well-suited for the used PID. In particular, at $t = 0\,\mathrm{s}$ the carrier flow valves opened while odour valves remained closed, for a duration of $10\,\mathrm{s}$. At $t = 10\,\mathrm{s}$ the odour valve and odourless compensation valve deployed anti-correlated pulse trains of various frequencies, for $2\,\mathrm{s}$ (see Fig. 3.5a (left) for the $10\,\mathrm{Hz}$ example). The fidelity for each square pulse was calculated as the value of peak to trough, normalised to the peak to the baseline value. The fidelities shown in Fig. 3.5a (right) were computed as the mean and standard-deviation across all square pulses fidelities of a particular modulation frequency.

Experimental protocol

Electronic nose placement: The electronic nose was attached to a movable arm and fixed in place downstream of the olfactometer outlet, with a distance of approximately 3 cm from outlet nozzle to the gas sensors. To ensure that the gas flow reached all the sensors on the board, we fine-tuned the alignment of the electronic nose with respect to the nozzle by trial-and-error until a strong response was obtained on all channels.

Heater modulation and odour delivery protocol: Three experiments with different sensor heater conditions were performed (see 3.5d for the conditions over time):

1. Sensor 1-8: 50 ms cycles between 150 °C and 400 °C.

2. Sensor 1-4: constant temperature of 400 °C, Sensor 5-8: 50 ms cycles between 150 °C and 400 °C.

3. Sensor 1-4: constant temperature of 400 °C, Sensor 5-8: 200 ms cycles between 150 °C and 400 °C (not used).

For each heater condition, odour stimuli of different pulse widths and concentrations (controlled by adjusting the shattering duty cycle of odour and mineral oil valves) were presented. After each odour stimulus, there was a 30 s recovery phase before the next stimulus onset. The set of different stimuli included:

- 4 odours + two control ('blank') vials, 50 repetitions 1 s odour pulses, at 100% concentration.

- 4 odours, 20 x 1 s odour pulses, at concentrations of 20%, 40%, 60%, and 80%.

- 4 odours, 5 repetitions of shorter odour pulses in the range [10, 20, 50, 100, 200, 500] ms, at 100% concentration.

- 2x6 odour pairs, 5 repetitions of 1 s anti-correlated pulse trains, at frequencies in the range [1, 2, 5, 10, 20, 40, 60] Hz, at 100% concentration.

- 6 odour pairs, 5 repetitions of 1 s correlated pulse trains, at frequencies in the range [1, 2, 5, 10, 20, 40, 60] Hz, at 100% concentration.

Within each experimental run, all the stimuli were presented in a fully randomised order. Fig. 3.5e shows the distribution of odours over time, binned in 1 h time intervals. A statistical χ^2 test was performed, confirming that the null hypothesis can't be rejected ($p = 0.364$), i.e. that the trials are in fact randomised. Importantly, the odour delivery protocol has not been synchronised with the sensor heater modulation phase.

PID recordings and odour onset/offset determination: A shortened version of the odour delivery protocol was deployed and recorded with the PID. Fig. 3.5e (left) displays a PID response to a 1 s isoamyl acetate pulse. For all the odours, the mean and standard deviation of the pre-stimuls baseline were computed, and a threshold of 4 times the standard deviation (4σ) defined. This was used to estimate an upper bound for the time from theoretical stimulus onset to odour exposure at the sensing site, and a lower bound from theoretical stimulus offset to the purging of the sensor site. Fig. 3.5e (right) displays all the extracted onset and offset values, indicating that the odour may reach the sensor as rapidly as in 10 ms, while the purging may take several hundreds of milliseconds. While PIDs are extremely fast, they too have a finite and odour dependent response time, thus the actual times may be shorter than this.

Pulse classification analysis

Feature extraction and validation: For evaluating what data features may be most suitable for the rapid classification of short odour pulses, we used experiment B, where sensor 1-4 were operated at a constant temperature of 400 °C, and sensor 5-8 employed heater cycles of 50 ms in the range of 150 °C and 400 °C.

Different sensor data features were used and evaluated: Data windows of 50 ms starting at a given time t relative to the simulus onset at $t_{onset} = 0\,\mathrm{s}$ were used to extract 1. raw data from constant heater sensors, 2. pre-stimulus ($t_{pre} = -5\,\mathrm{s}$) baseline subtracted data from constant heater sensors, 3. raw data from cycled heater sensors, and 4. pre-stimulus ($t_{pre} = -5\,\mathrm{s}$) baseline normalised data from cycled heater sensors. For the normalisation in the latter, the procedure is illustrated in Fig. 3.6c and described in the following. The extraction of the feature $G(\mathcal{D}_s, t)$ can be summarised as applying a chain of a sensor-wise logarithmic transformation and a maximum scaling, to both a 50 ms baseline data snippet before stimulus onset (here, $t_{pre} = -5\,\mathrm{s}$) and to a snippet at time t, and then computing their vector difference:

$$G(\mathcal{D}_s, t) = \frac{\log(H(\mathcal{D}_s, t))}{\max(\log(H(\mathcal{D}_s, t)))} - \frac{\log(H(\mathcal{D}_s, t_{pre}))}{\max(\log(H(\mathcal{D}_s, t_{pre})))} \tag{3.1}$$

Here, $H_{cycle}(\mathcal{D}_s, t)$ describes a kernel extracting data from the sensor recordings $\mathcal{D}_s$, using a window that begins at time t and ends after one full heater cycle (e.g. 50 measurements). The sensor index is denoted as s. The pre-stimulus normalisation is a necessary and commonly used step that eliminates the first-order sensor baseline (i.e. here the data 5 seconds before the stimulus), which is known to often be contaminated the effects of sensor drift and fluctuations in environmental conditions [Den+22b].

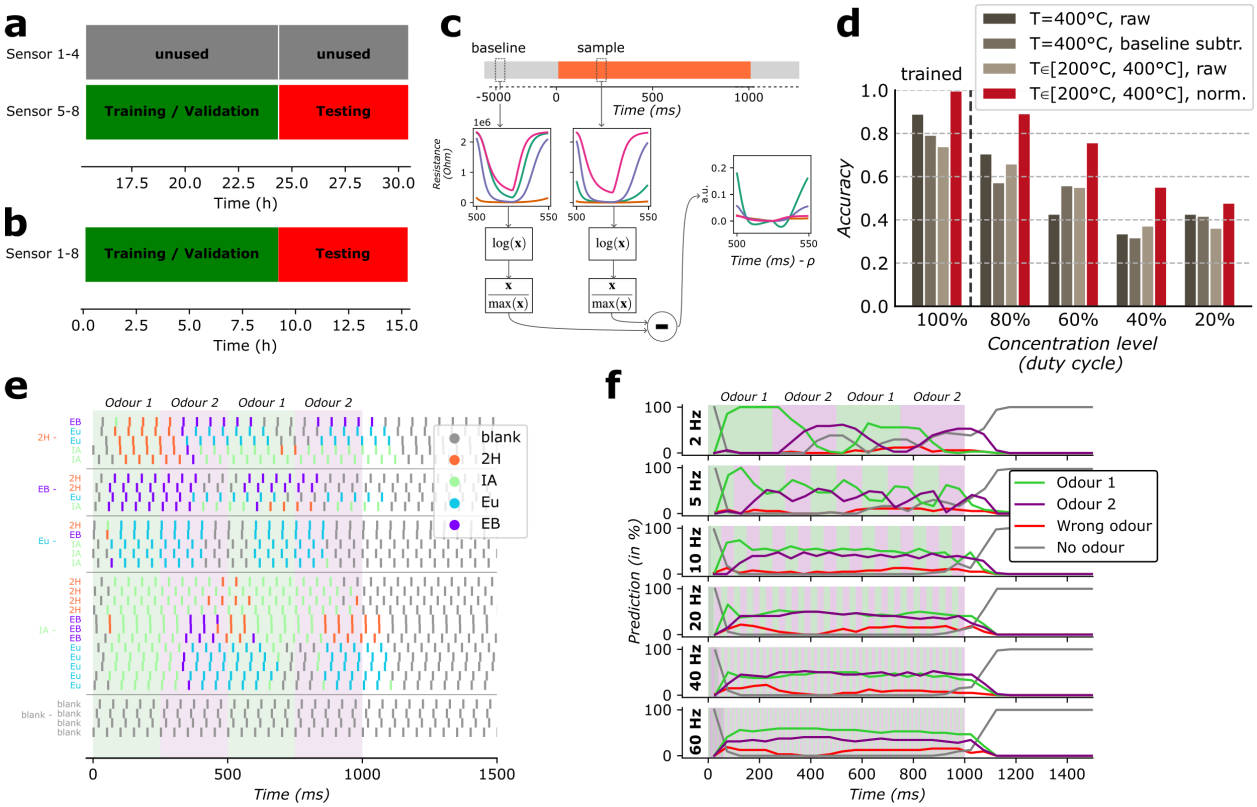

Figure 3.6: Supplementary figure for fast odour classification. **a,** Data splitting for robustness analysis of the rapid heater modulation data features (see Fig. 3.9e). **b,** Data splitting for evaluating the dynamic classification of millisecond odour pulses (see Fig. 3.10). **c,** Normalisation procedure for the heater modulation data feature. Time shifted by cycle phase ρ w.r.t. odour onset, for visual guidance only. **d,** Accuracy scores for a k-nearest neighbours (k-NN) classifier trained on 50 ms data features from 1000 ms odour pulses at full concentration, and tested on 50 ms features from 1000 ms odour pulses at different concentration levels (tuned by adjusting the duty cycle of the micro-valves). Features are compared for constant heater sensor readings (raw and baseline-normalised) and cycled heater sensor readings (raw and normalised, as described in c)) **e,** Odour stimulus classification for anti-correlated odour patterns over time. An RBF-kernel SVM classifier was trained on 50 ms features from 1000 ms odour pulses, and tested on odour anti-correlated odour patterns of various frequencies. Shown here is is a 2 Hz pattern. **f,** Classification correctness over time (evaluated via the true odour presence), for anti-correlated odour patterns of different switching frequencies.

The data was split into one set for training & validation, and one set for testing — with a ratio of 60% to 40% (see Fig. 3.6a). The former was used to train and validate a k-NN classifier using different features via cross-validation. The latter was used to evaluate the performance with the different features used.

Classifier training was performed on data features from sensor responses between 500 ms and 1000 ms after stimulus onset, where the stimuli were 1000 ms odour pulses of the gases 2H, EB, IA, Eu and blank, at 100% concentration. Testing was performed on equivalent data features, however now the stimuli concentration was sampled in the range [20, 40, 60, 80, 100]%, and the blank class was omitted. Fig. 3.6d displays the achieved performance using the different data features. The normalised cycled-heater data feature $G(\mathcal{D}_s, t)$ outperforms the other tested features, both in accuracy at 100% concentration, as well as for the reduced concentrations. For clarity, Fig. 3.9e shows a subset of Fig. 3.6d.

Dynamic pulse classification: For the dynamic classification of short odour pulses, experiment A was used with all 8 sensors modulated on a 50 ms period betweeen 150 °C and 400 °C. Again, a 60% vs. 40% split for training & validation vs. testing was performed, where the former was used to determine a suitable classifier and its hyper-parameters via cross-validation, while the latter served to evaluate the performance of the classifier (see Fig. 3.6b). Data features were extracted as described in Eq. 3.1. For training, the underlying data for the features are the sensor responses for the subsequent 2000 ms after the onset of 1000 ms odour stimuli, for concentrations in the range [20, 40, 60, 80, 100]%. The data features were labelled according to their time t as follows:

Procedure 1: Training data labelling procedure

Input: Parameters t_{onset}, t_{offset}, τ, d, and $\mathcal{O}_{stimulus}$

Output: Label y for the feature

if $t_{onset} \leq t < t_{offset} - \tau + d$ **then**
    Feature is fully within measurable odour pulse.
    if $\mathcal{O}_{stimulus}$ is not 'blank' **then**
        $y = \mathcal{O}_{stimulus}$
    else
        $y =$ 'blank'
else if $t_{offset} + d \leq t$ **then**
    Feature is fully after measurable odour pulse.
    $y =$ 'blank'
else
    Feature timing is ambiguous; exclude data feature from training set.
    $y =$ 'rejected'

Here, $t_{onset} = 0\,\text{ms}$ and $t_{offset} = 1000\,\text{ms}$ are the stimulus onset and offset respectively, $\tau = 50\,\text{ms}$ is the feature duration, $d = 10\,\text{ms}$ is the upper bound stimulus delay as computed earlier, $\mathcal{O}_{stimulus}$ is the stimulus odour of the corresponding trial, and y is the prescribed label of the data feature in question. This procedure is illustrated in Fig. 3.10a.

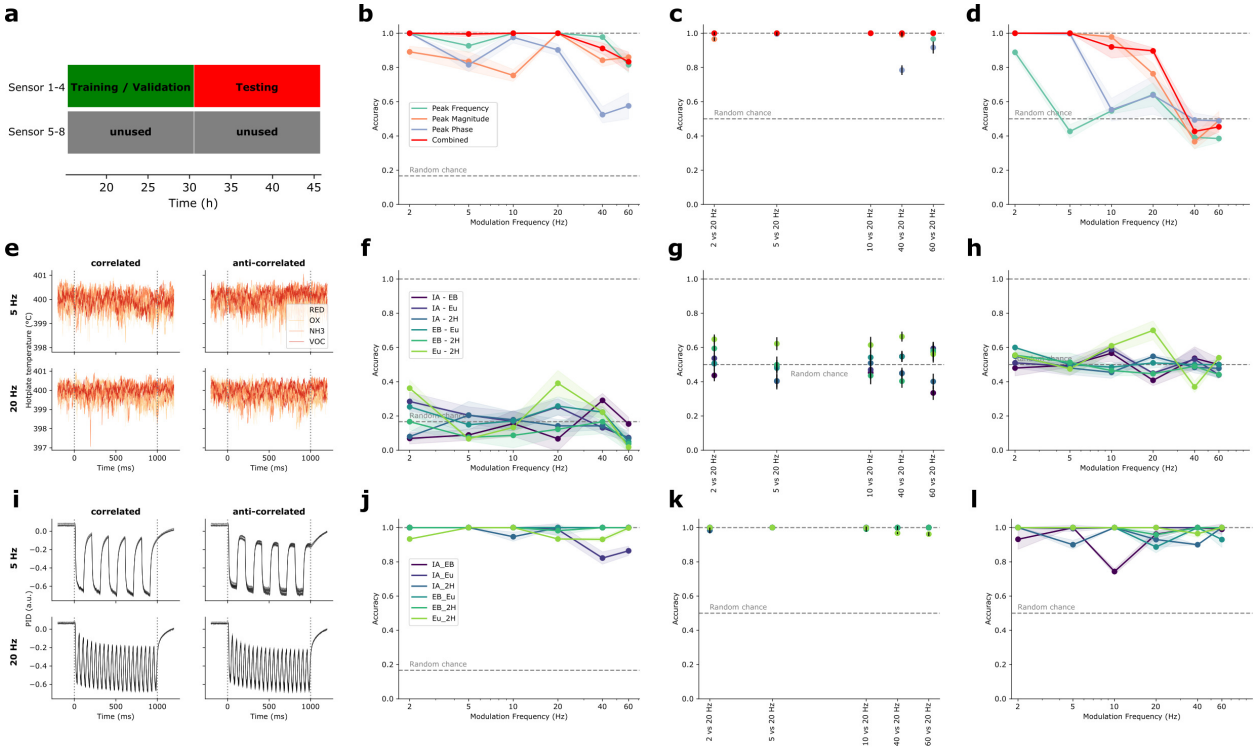

Figure 3.7: Supplementary figure for temporal pattern discrimination. a, Data splitting for evaluating the temporal pattern discrimination performance (see Fig. 3.11). **b-d,** Validation accuracy plots for different extracted DFT-spectrogram peak features using the MOx gas sensor resistances. **b,** Modulation frequency classification, **c,** pairwise modulation frequency classification, and **d,** correlated vs anti-correlated modulation discrimination. **e,** MOx heater temperature values for different odour modulations. Here shown are data for the odour pair IA (isoamyl acetate)- EB (ethyl butyrate), 5 trials each for 5 Hz correlated, 5 Hz anti-correlated, 20 Hz correlated and 20 Hz anti-correlated respectively. **f-h,** Test accuracy plots for different odour pair modulations, using the MOx heater temperature values. **i,** Photoionisation Detector (PID) responses for different odour modulations (odour stimuli as in **e**). **j-l,**, Test accuracy plots for different odour pair modulations, using the Photoionisation Detector (PID) responses. For all the classification tasks, an ensemble of Random Forest Classifiers was used. The mean and error estimations arise from repeating training and testing with different random seeds.

For the normalised data features, several classification algorithms were trained and validated via five-fold stratified cross-validation. The best performing algorithm with corresponding hyper-parameters was selected, which here was a Support Vector Machine (SVM) classifier with radial basis function kernel [VTS04] ($C = 1e3$, $\gamma = 1e-4$, balanced class weight). Ultimately, an ensemble classifier was composed from the five SVMs trained on each split.

For testing how well the trained classifier performs on shorter odour pulses, the features were extracted from sensor response data for the subsequent 2000 ms after the onset of odour stimuli of different durations, at 100% concentration. The stimulus durations fall within the set {10, 20, 50, 100, 200, 500, 1000}ms. For each data feature, the classifier predicted the odour, which is illustrated as a raster plot in Fig. 3.10b. The predicted odours y were compared against the actual stimulus odour $\mathcal{O}_{stimulus}$ and divided in predicting 'correct odour', 'wrong odour' and 'no odour', resulting in Fig. 3.10c. To extract the accuracy for each pulse duration, a confusion matrix was composed by — for each trial — comparing the most predicted non-blank class against $\mathcal{O}_{stimulus}$, across multiple trials. The on- and offset times correspond to the elapsed time from odour onset to first non-blank prediction, and from odour offset to first 'blank' prediction, respectively.

An analogous procedure was followed for testing the trained classifier on anti-correlated patterns of odour pairs, resulting in predictions over time, as shown in Fig. 3.6e & 3.6f.

Temporal structure analysis

For the temporal structure analysis, i.e. the determination of the frequency and the phase-shift of the two-odour pulse trains, the constant heater sensor data (i.e. sensors 1-4) of experiment B and C were used. In particular, experiment B was used for training and validation (i.e. finding and evaluating a suitable data feature and classification algorithm), where experiment C was used for testing, see Fig. 3.7a for an illustration.

Feature extraction: For each data trial, we extracted sensor data $\mathcal{D}_s$ from $t = t_{onset}$ to $t = t_{offset} + b$, where $b = 100\,\text{ms}$ to account for the stimulus delay and potential sensor lag. The data was then log-transformed and differentiated before applying a discrete Fourier transformation $\mathcal{F}(.)$, using the fast Fourier Transformation algorithm [CT65]:

$$L(\mathcal{D}_s) = \mathcal{F}\left(\frac{d}{dt}\log(\mathcal{D}_s)\right) \tag{3.2}$$

All triplets [frequency, magnitude, phase] were extracted, and sorted according to the magnitude. For each of the four sensors, the triplet with the highest magnitude was selected, collectively composing a 12-dimensional data feature.

Temporal structure classification: The data features and potential classifiers were evaluated on a 10-fold cross-validation using the training & validation data (i.e. experiment B) and the three different tasks described earlier. We decided on utilising a Random Decision Forest (RDF) classifier [Ho95] (balanced class weight, $N_{tree} = 100$), and on using the same 12-dimensional data feature for all tasks. See Fig. 3.7b-d for an evaluation of the data fea-

tures. For each cross-validation training split we trained a RDF classifier, then combined them to form an ensemble classifier for each task, which was finally evaluated on the testing data (i.e. experiment C). For the validation using hotplate temperature and PID data, the analogous pipeline was used, except that we omitted the log-transformation.

Comparing electronic nose performance with mouse performance

Performance analysis of the electronic nose to discriminate odour correlation structure was carried out in the style of a previously published experimental dataset (See Ackels et al. 2021 [Ack+21]). This allowed for a direct comparison of the electronic nose performance with that of mice during an operant conditioning task. A complete description of the experimental conditions and data analysis can be found in the original paper. In brief, two cohorts of up to 25 mice were housed in a common home cage system [Ers+19b] that is used as an automated operant conditioning setup. Mice were trained to discriminate perfectly correlated from perfectly anti-correlated odour stimuli switching at frequencies ranging from 2 Hz to 81 Hz. Task frequency was randomized from trial to trial. Odours were presented with a multi-channel high-speed odour delivery device similar to the one used in this thesis. During a go/no-go task animal performance was rated based on their lick responses to $S+$ (rewarded) and $S-$ (unrewarded) stimuli. For roughly half of all mice, the correlated pattern was $S+$ and the anti-correlated pattern was $S-$. In the other half of the group this reward valence was reversed. All stimuli were 2 s long. A water reward could be gained by licking so that licking was detected for at least 10% of the stimulus time during an $S+$ presentation (a 'Hit'). Licking for the same amount of time during $S-$ presentation resulted in a timeout interval of 7 s. In all other response cases, the inter-trial interval was 3 s and no water reward was delivered. All behavioural performance within a specified trial bin was calculated as a

weighted average of $S+$ versus $S-$ performance:

$$\text{Performance} = \frac{(Hit/S+) + (CR/S-)}{2} \tag{3.3}$$

in which $S+$ is the total number of rewarded trials, $S-$ is the total number of unrewarded trials, Hit is the total number of rewarded trials in which a lick response was detected, and CR (correct rejection) is the total number of unrewarded trials in which no lick response was detected.

Results

High-speed electronic nose and odour delivery system

We constructed a portable high-speed and high bandwidth electronic nose, which leverages the advantage of MEMS-based gas sensors and their rapid response times. We emphasised form factor and power consumption considerations that allow for sophisticated field measurements under space and power constraints, such as mobile robotic platforms[JMI21]. Our design (Fig. 3.8c & 3.5a) consisted of the following elements: a microcontroller for data processing and storage, eight analogue metal-oxide MEMS gas sensors (Fig. 3.8d), associated analogue circuitry and data converters, and a combined pressure, humidity and temperature sensor.

Ideal MOx sensor operation requires the sensing site to be heated to several hundred degrees. The sensor response is highly dependent on the temperature and its variation over time. Previous studies have shown that a modulation of the sensor's operating temperature often leads to better and faster gas discrimination performances [Ver+14], however the suggested sensor heater cycle durations were on the orders of seconds to minutes [Ver+14; Xin+19; NT20; Di

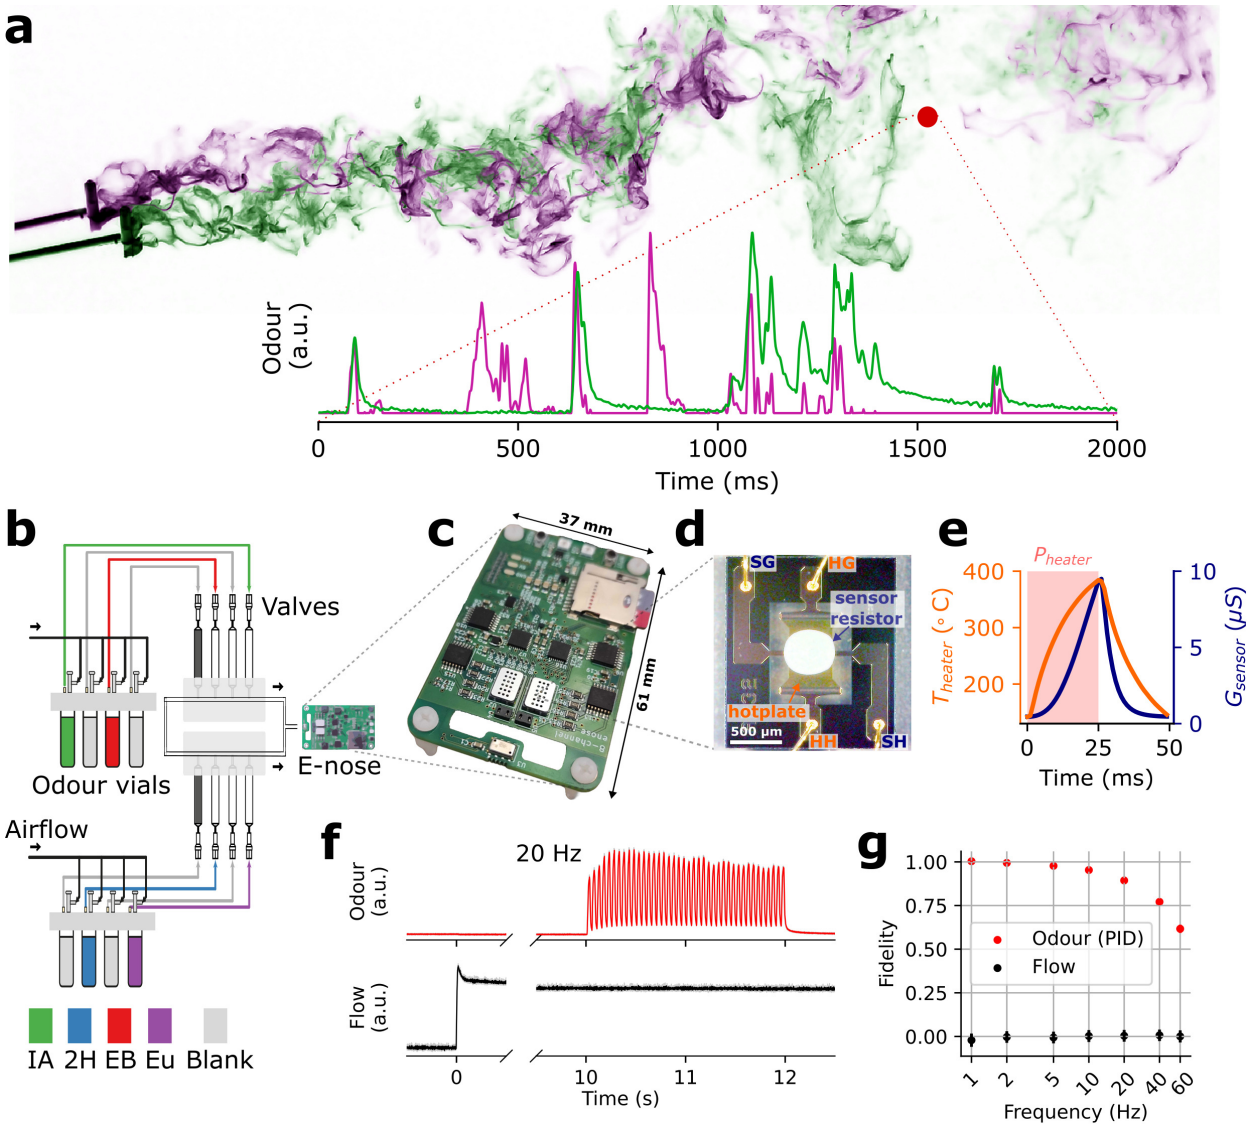

Figure 3.8: Electronic nose and odour delivery system. **a**, Decoding temporal information of odour plumes requires fast sensing. Top: Two sequential $TiCl_4$ smoke plume photographies, shifted and superimposed, kindly provided by Dr. Paul Szyszka. Bottom: Dual-PID recordings of source-separated odour plumes, from Ackels et al. [Ack+21]. Plume and sensor location (red) for illustrative purposes only. **b**, Experimental setup with odour delivery device and electronic nose. Adapted from Ackels et al. [Ack+21]. **c**, Electronic nose circuitry. **d**, Microscopy image of the MiCS-6814 NH3 sensor with its housing removed. **e**, Heater modulation cycle in ambient air. **f**, PID and flow meter traces for a 20 Hz stimulus. Solid / faded (occluded) traces for mean / std. of five trials. **g**, Resulting olfactometer temporal fidelity, for various frequencies. Odourants abbreviations: IA: isoamyl acetate; EB: ethyl butyrate; Eu: cineol; 2H: 2-heptanone; blank: odourless control.

+21; NT22; Men+23]. Aiming to achieve ultra-fast heater cycles, our design couples each sensor with a separate temperature control loop, which samples the temperature and adjusts the hotplate current at high frequency. Fig. 3.8d shows a typical heater modulation cycle in ambient air, where the sensor resistance follows the hotplate temperature in a low-pass fashion. In our experiments, we used two different heater temperature control schemes: one that cycled between low and high temperature values (150 °C and 400 °C), and one at a constant high temperature (400 °C).

To provide odour stimuli to the e-nose, we used an odour delivery system that can reliably present gaseous odour samples with a bandwidth beyond 60 Hz, described earlier [Ack+21; Das+22] and depicted in Fig. 3.8b. The system was based on high-speed micro-valves and incorporated a flow compensation mechanism [Ack+21], ensuring exceptionally high temporal signal fidelity, and constant flow across the stimuli (Fig. 3.8f & 3.8g and Methods). As prototypical, simplistic high-frequency odour stimuli, we used square pulses of different duration and separation times. A set of odourants that resembles smells encountered in nature was considered: Ethyl butyrate (pineapple), isoamyl acetate (banana), cineol (eucalyptus) and 2-heptanone (cheese). The odourants were diluted in odourless mineral oil solvent. Additionally, we used two (identical) pure solvent samples as controls. The odours were presented as singular pulses with varying durations (10 ms - 1 s) and concentrations (20%-100%), and as correlated and anti-correlated odour pulse trains (1 s) at different modulation frequencies (2Hz - 60Hz).

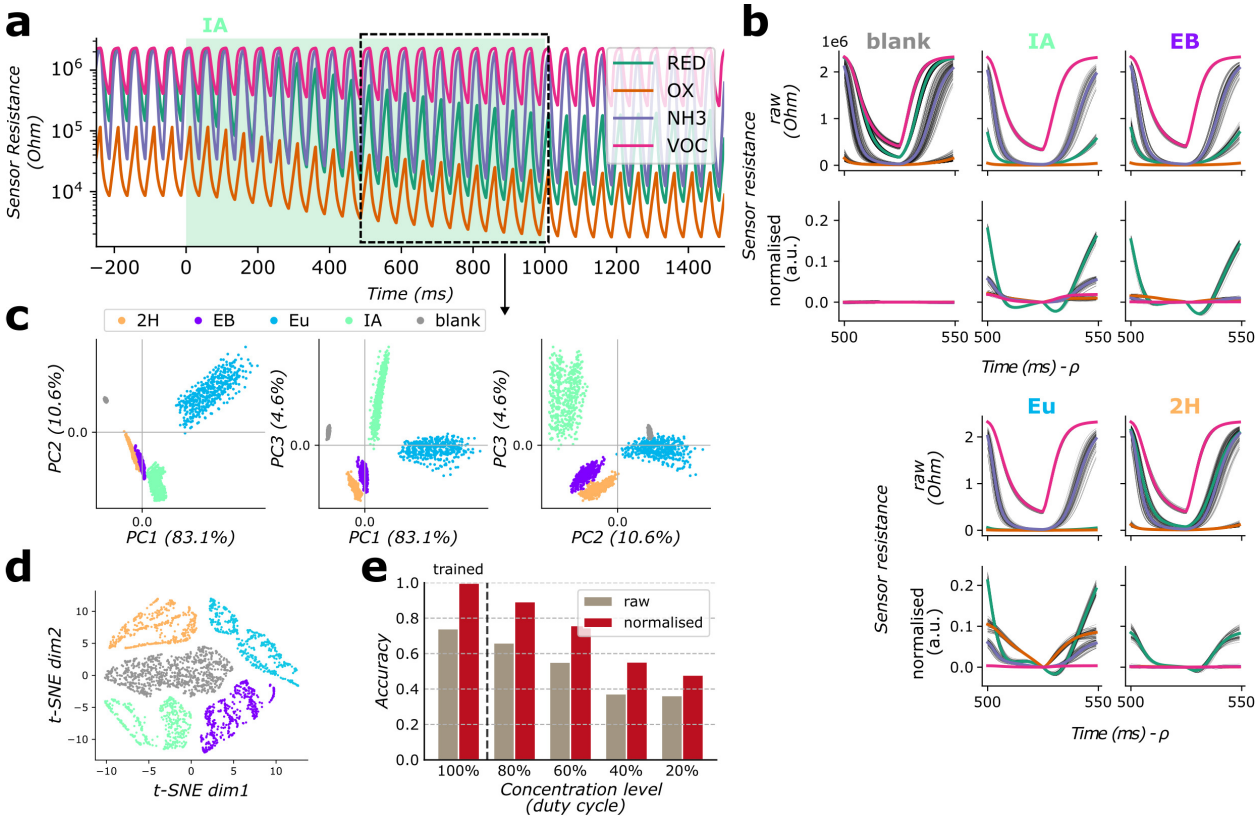

Figure 3.9: Rapid heater modulation enables robust data features. **a**, Sensor resistance of four MOx sensors with 20 Hz hotplate temperature modulation, responding to a 1 s odour pulse of isoamyl acetate (green background). **b**, 50 ms data feature for different gases, selected between 500 ms and 550 ms after odour pulse onset. Raw sensor response (upper) and normalised sensor response (lower, see Methods for normalisation procedure). Time shifted by cycle phase ρ w.r.t. odour onset, for visual guidance only. **c**, Principal component analysis (PCA), explained variance (most left) and projections, and **d**, t-distributed stochastic neighbour embedding (t-SNE) visualisation, for the set of normalised data features extracted between 500 ms and 1000 ms after odour onset. **e**, Accuracy scores for a k-nearest neighbours (k-NN) classifier trained on 50 ms data features from 1000 ms odour pulses at full concentration, and tested on 50 ms features from 1000 ms odour pulses at different concentration levels (tuned by adjusting the duty cycle of the micro-valves).

Rapid heater modulation enables robust data features

While cycling the sensor heater temperature can yield better odour classification results, the cycle duration may be restricting the temporal bandwidth at which a stimulus can be resolved. In recent studies, we tested the effect of 150 ms duty cycles and found evidence for robust data features [Den+22a; DDS22]. In the current work, we leveraged our system's

ability to rapidly modulate the sensor temperature, and cycled the heater temperature between a low step at 150 °C and a high step at 400 °C with a period of 50 ms. Notably, this is orders-of-magnitudes shorter than what had been suggested in previous studies [Men+23]. The resistance of the gas sensing elements closely tracks these changes in operating temperature (Fig. 3.8e & 3.9a), enabling us to extract gas features that are phase-locked with the heater cycles for subsequent analysis and classification. For this purpose, we divided the continuous stream of gas sensor samples into 50 ms chunks aligned with the temperature cycles (Fig. 3.9b, upper row). The 50 ms data features further underwent pre-stimulus baseline normalisation and scaling (Fig. 3.9b, lower; from now on referred to as "normalised data feature"); for details see Methods).

For testing class discriminability and robustness to concentration fluctuations, data features were extracted by sampling four sensors between 500 ms and 1000 ms after the onset of a 1000 ms odour stimulus. Principal Component projections (Fig. 3.9c) and t-distributed stochastic neighbour embeddings (t-SNE) (and 3.9d) show distinct clustering that coincided with odour classes. Further, a k-nearest neighbours (k-NN) classifier was trained on data features of one second long odour pulses at full concentration (100%), and tested on features of one second long odour pulses at various concentrations (20% - 100%). The classification performance results are shown in Fig. 3.9e. Notably, for the normalised data feature, the model provided 100% classification accuracy at the trained concentration level, which remained at (88.7 ± 0.5)% and (81.2 ± 0.6)% when tested on 80% and 60% of the trained concentration level. Accuracy at lower concentration levels dropped significantly but remained well above chance. This is significantly better than what is achieved with the raw data feature (see Fig. 3.9e).

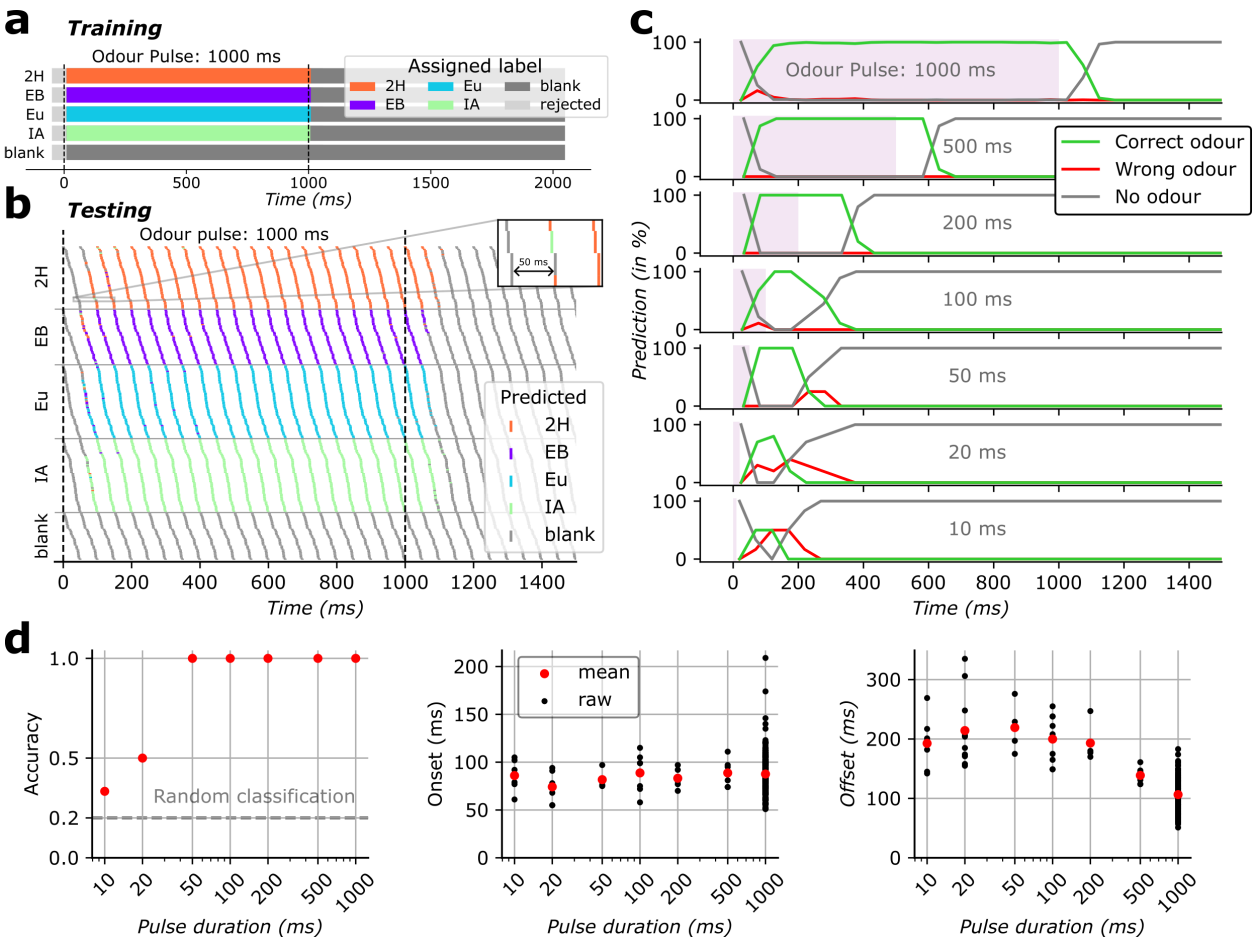

Figure 3.10: Electronic nose can classify short odour pulses based on 50 ms data features. **a**, Feature labels for the training set were phase aligned in relation to odour on- and offset. Features that overlapped with transition periods were not considered for training ("rejected", see Methods for parameters). **b**, Odour stimulus classification over time for odour pulses of various lengths (10 ms - 1000 ms, as predicted by a RBF-kernel SVM classifier trained on 50 ms features from 1000 ms second odour pulses. Shown here are 1000 ms pulses. For visual clarity only, the trials are sorted by odour, and within each odour are sorted by phase w.r.t. stimulus onset. **c**, Classification correctness over time (evaluated via the true odour presence), for different pulse durations. **d**, Test accuracy, onset time and offset time for the prediction over time described in **b** & **c**. Onset and offset were extracted using time-to-first non-'blank' and 'blank' prediction respectively, and shown here with respect to theoretical odour onset and offset.

Time-resolved classification of millisecond odour pulses

In natural settings, odour bouts can be as brief as only milliseconds long. For an agent's successful interaction with the environment, this requires the ability to classify odours fast and robustly. We evaluated the ability of the electronic nose to classify odour pulses of various durations. A Support vector machine (SVM) with Gaussian radial basis function was trained

on 50 ms data features, which were acquired from eight gas sensors throughout a 1000 ms odour stimulus at five concentration levels. Control trials ('blank') were included, obtained during a 1000 ms odourless mineral oil stimulation or immediately after odour pulses. See Fig. 3.10a for a depiction of the labelled features. The trained model was deployed to predict the odour presence over time during exposure to odour stimulations of various durations, ranging from 10 ms to 1000 ms. Fig. 3.10b displays the predicted classes over time on the example of an 1000 ms odour pulse, where Fig. 3.10c summarises the predictions over time for all pulse durations.

From these predictions, the corresponding accuracy, onset times, and offset times were derived and shown in Fig. 3.10d. The classifier attained a 100% accuracy in predicting the correct class for odour pulse durations ranging from 1000 ms down to 50 ms, despite not having been trained on pulses shorter than 1000 ms. Accuracy dropped for 20 ms and 10 ms pulses but remained above chance level. Notably, the classifier accurately and rapidly predicted the recovery of the sensor site, indicating 'no odour' when no odour was present. The time required for the classifier to correctly identify the odour remained relatively consistent across odour pulse durations, with an average value of (87 ms $\pm$ 20 ms). Following odour offset, the classifier robustly predicted 'no odour' within (106 ms $\pm$ 24 ms) for 1000 ms odour pulses. For shorter durations, this time increased inversely proportional to the pulse duration, which can presumably be attributed to the sensor's integration time that may approach or exceed the duration of the short odour pulses.

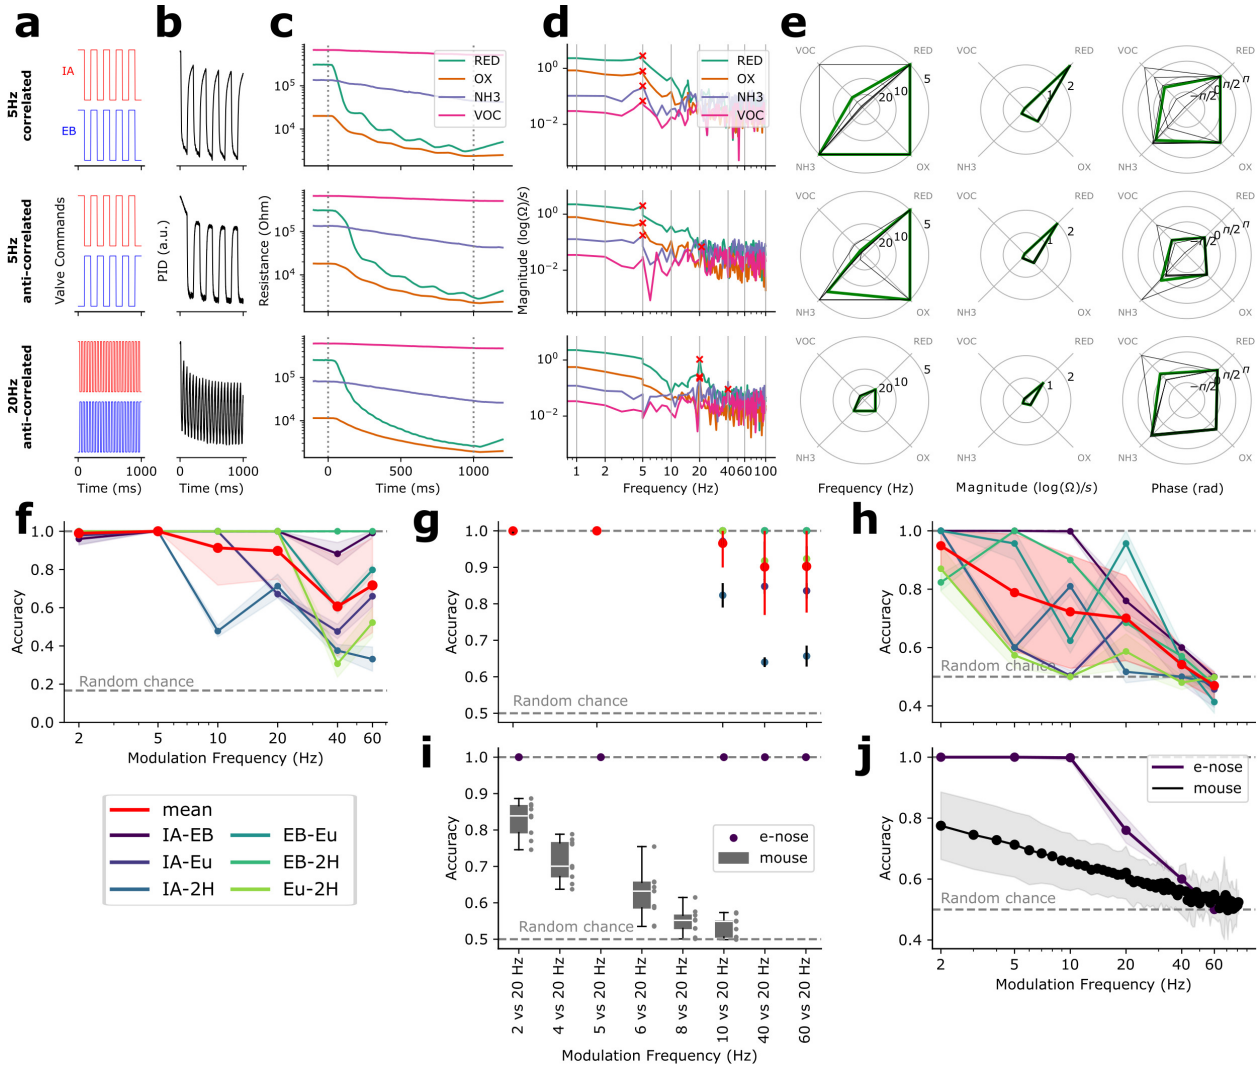

Figure 3.11: Decoding temporal structure of rapidly switching odours. **a,** Odour valve commands. **b,** PID response. **c,** Electronic nose response. **d,** Discrete Fourier Transformation (DFT) of first derivative of the sensor log-resistance. Crosses denote highest-magnitude peaks. DFT bin frequencies were rounded to nearest integers, for visual clarity only. **e,** Feature visualisation frequency, magnitude and phase of the dominant DFT peaks. Thick lines; means of corresponding trials, thinner lines; single trials. **f**, Class-balanced accuracies for modulation frequency classification. **g**, Accuracies for binary modulation frequency classification. **h**, Class-balanced accuracies for binary modulation mode classification (corr. vs. anti-corr.). **i**, Subset of *g* for IA-EB, for mouse performance comparison (described in detail in Ackels et al.[Ack+21]). **j**, Subset of *h* for IA-EB, for mouse performance comparison. Panels *a-e* show representative trials only. For *f-j*, electronic nose accuracy mean and SD (clipped at 1.0) arise from repeated training and testing with different random seeds.

Decoding temporal structure of rapidly switching odours

In the presence of multiple odours, detecting whether the odour encounters are correlated or not can help to infer whether they come from the same source or from separate locations [Hop91]. Further, information about the encounter frequency can give rise to spatial source information [SBH16]. It has been shown that mice can distinguish between correlated and anti-correlated odour pairs reliably up to correlation frequencies of 40 Hz [Ack+21]; a feat that has not yet been matched in robotic systems. Considering performance metrics based on similar tasks, here we explored the ability of the electronic nose to resolve temporal structure of odour stimuli. Rapidly alternating odour pairs were presented at frequencies between 2 Hz and 60 Hz for a duration of one second. We discriminated between two odour pulse trains being either in phase (correlated) or shifted by half a cycle (anti-correlated) (Fig. 3.11a). The resulting odour patterns follow the pulses rapidly (PID recordings in Fig. 3.11b).

While heater modulations lend themselves for extracting phase-locked data features and thus allowing for efficient odour classification, maintaining the sensor heater temperature constant instead allows for analysing the data in continuous time. Particularly when observing repeating patterns or complex temporal dynamics of a stimulus this may be advantageous, as the sensor response can be regarded in its frequency domain. Thus, for the following experiments, we operated four MOx sensors of the electronic nose under a constant hotplate temperature of 400 °C. The sensors responded to the stimuli by dropping their resistance at the pulse-train onset, with the stimulus modulation visually embedded in the response (Fig. 3.11c). We extracted data features by differentiating and logarithmically scaling the raw sensor response, followed by a discrete Fourier transform (DFT) (Fig. 3.11d). For each

sensor, where the maximal magnitude is found, the frequency, the magnitude and the phase were extracted. This yielded a 12-dimensional data feature (Fig. 3.11e). A visual comparison of the features reveals distinct differences between correlated and anti-correlated pulse trains (top vs. middle), as well as between different frequencies (middle vs. bottom).

On those features, ensembles of Random Forest classifiers were trained for three tasks: 1. Decoding the modulation frequency of two odour pulse trains from a set of frequencies, 2. predicting the modulation frequency of two odour pulse trains from pairs of frequencies, and 3. decoding if two odours pulse trains are either correlated or anti-correlated. For the latter two tasks and a subset of the odours, a comparison with the mouse performance as detailed in Ackels et al. [Ack+21] is provided.

The test accuracies for the three tasks and all gas combinations are shown in Fig. 3.11f-h. For task 1, the data recorded with the electronic nose enables nearly perfect frequency classification (Fig. 3.11f) for modulation frequencies up to 5 Hz, then on average decreasing to (0.91 ± 0.20) and (0.90 ± 0.15) for 10 Hz and 20 Hz respectively, and finally dropping to (0.61 ± 0.26) and (0.72 ± 0.25) for 40 Hz and 60 Hz respectively. For the pair-wise frequency classification (task 2, Fig. 3.11g), classification performance is perfect for modulation frequency pairs 2 Hz vs 20 Hz and 4 Hz vs 20 Hz, decreasing to accuracies of (0.97 ± 0.07), (0.90 ± 0.13) and (0.90 ± 0.13) for the pairs 10 Hz vs 20 Hz, 40 Hz vs 20 Hz and 60 Hz vs 20 Hz, respectively. For discriminating correlated vs. anti-correlated pulse trains (task 3, Fig. 3.11h), it appears that the electronic nose, on average, enables high prediction scores of (0.95 ± 0.08), (0.79 ± 0.20), (0.72 ± 0.19) and (0.70 ± 0.14) for modulation frequencies of 2 Hz, 5 Hz, 10 Hz and 20 Hz, respectively. This drops to (0.54 ± 0.06) for 40 Hz and finally to (0.47 ± 0.05) for 60 Hz.

In Ackels et al. [Ack+21], the odour pair isoamyl acetate — ethyl butyrate (IA-EB) has been used to test the discrimination power of fast odour dynamics in mice. In the following, we consider the corresponding subset of the electronic nose recordings and compare them to the named study. For the pair-wise frequency classification (task 2, Fig. 3.11i), the electronic nose classification performance is perfect for all the tested modulation frequency pairs, from 2 Hz vs 20 Hz up to 60 Hz vs 20 Hz. Here, the mouse performed significantly worse — for the pair 2 Hz vs 20 Hz, the mouse accuracy score was (0.83 ± 0.05) and then progressively dropped down to (0.53 ± 0.03) for 10 Hz vs 20 Hz. Finally, considering the results from the phase prediction task (task 3, Fig. 3.11j), it appears that the electronic nose enables perfect prediction scores up to modulation frequencies of 10 Hz, which then steeply drops to an accuracy of (0.76 ± 0.05) for 20 Hz, (0.60 ± 0.00) for 40 Hz, and finally to chance level for 60 Hz. In comparison, the mouse scores (0.78 ± 0.11) at a modulation frequency of 2 Hz, linearly decaying in accuracy down to chance level at around 80 Hz.

In order to validate if the observed performance can be attributed to the odour signal, and not to potential artefacts caused by potential hotplate temperature variations (which may be caused by unnoticed flow fluctuations), we repeated the experiment using the hotplate temperature signal (see Fig. 3.7i), as well as the PID responses (see Fig. 3.7m). In both cases, the analogous feature extraction and classification pipeline was performed, resulting in classification performances as displayed in Fig. 3.7j-l for the hotplate temperature, and Fig. 3.7n-p for the PID responses. The analysis confirmed that there was not enough information in the hotplate temperature response alone to classify the odourants with above chance performance. Further, using the PID response, which should be unaffected by potential flow fluctuations and is commonly used as a ground-truth measurement, nearly perfect accuracy scores were achieved for most gas combinations across the tested tasks.

Discussion

For many tasks and applications in robotics, natural and turbulent environments pose the challenge of highly dynamic and rapidly changing odour concentrations, which demands high temporal resolution odour sampling and processing. Intrigued by the exceptional speed at which animals process and respond to odours, we challenged the limits of artificial olfaction by introducing and evaluating a portable and low-power high-speed electronic nose. For this, we coupled an integrated design of MEMS-based MOx sensors and fast sampling periphery with a set of highly optimised algorithms for control, sensing, and signal processing. To assess its capabilities, the electronic nose was subjected to odour stimuli delivered through a high-bandwidth system. Odours were presented in the form of square pulses with varying durations and concentrations, or as pairs of pulse trains at varying frequencies and phase.

We demonstrated that the electronic nose can successfully infer the odour identity of single-odour pulses down to durations of 10 ms, albeit being trained on 1 s odour pulses only. This was achieved through modulating the sensor temperature with cycle periods of 50 ms, extracting and pre-processing the corresponding sensor response, then training and evaluating a classifier. Further, we demonstrate the electronic nose's ability to predict whether two-odour pulse trains were correlated or anti-correlated up to switching frequencies of 40 Hz; matching or exceeding mice on the equivalent task. For tasks involving determining the odour switching frequency (multi-class and binary), we demonstrate a high performance up to 60 Hz, outperforming mice on equivalent tasks. For this, the sensor heaters were set to provide a constant temperature, which allowed an analysis of the data in the frequency domain.

An interesting discussion arises on *why* fast heater modulations lend themselves so well for fast odour pulse classification. We see two main reasons: 1. Heater modulations generally result in much faster sensor responses compared to isothermal operation [Ver+14]. This is because the sensors' conductive behaviour, influenced by adsorption, desorption, diffusion, and reaction processes, varies substantially with temperature [BW01]. By cycling the sensor through a range of temperatures, transient processes within these phenomena are driven, producing responses that are highly characteristic of a given gas. This creates a more nuanced mapping of the sensor response, capturing a wide range of temperature-dependent behaviours specific to each odour in a short time. 2. From a signal processing perspective, coupling sample windows to the phase of heater modulations creates natural, multi-dimensional units of computation. Each data window begins at the same temperature and follows the same temperature range, akin to active sensing. This approach results in consistent data features, which improves machine learning classifier performance by allowing for a well-separable feature space and thus for robust odour classification.

Further insights on task-specific sensing modes and sampling can be gained from those results. First, data windows that are phase-locked with ultra-short sensor temperature cycles appear to be a sensible choice when given the task of odour classification. The high sampling rate allows for a multi-dimensional data feature (here, 8 sensors times 50 samples per feature), which — together with the pre-stimulus normalisation procedure — successfully captures the odour-specific sensor response. In mammalian sensory neuroscience, the analogy would be the phase-coupling of spike-trains to the inhalation cycle, allowing the spike-timings to encode information about the odour identity, thus suggesting one "sniff" as the unit of olfactory processing [KUM06]. Conversely, if the task is not classification but decoding temporal information about the odour stimuli, such as frequency or correlations, integrating

the information across an artificial time window would limit the temporal resolution, hence recording continuously and without heat modulation might be the better choice. The ability of mammals to access temporal stimulus information at sub-sniff resolution has been demonstrated [Ack+21; Das+22], and shown to be relevant for behavioural tasks [Ack+21]. Those findings may suggest future experiments in which both modes are active simultaneously on separate sensor instances — continuously-sampled constant-temperature and time-integrated temperature-modulated — which could allow for extracting information about the temporal profile and the identity in parallel. Research on insect have suggested dual-pathway olfactory systems [GR10], which may facilitate the simultaneous extraction of odour identity and concentration information [Sch+11]. Such an approach might suggest elegant solutions to the olfactory cocktail party problem [Rok+14].

The proposed experimental approach appears to be well-suited to verify and characterise the temporal capabilities, and in particular the high-frequency properties of the device. Yet—it presents limitations when considering its evaluation with respect to more natural environments. The current setup, characterised by high concentration stimuli (not measured, but likely tens to hundreds of ppms), fast valves, short odour delivery lines and precise flow compensation, does not necessarily replicate the turbulent and intermittent nature of odour plumes encountered in typical environments, where odour packets are less sharp and potentially of lower concentration. This should motivate further experiments that take into account scenarios that resemble natural environments more realistically. When designing such, it will crucial to ensure that the setup allows for both accurately replicating the turbulent and variable nature of odour plumes, and carefully controlling or monitoring the ground-truth odour concentrations. In such, it will be important to use methods that do not interfere with the natural flow characteristics, which e.g. for PID concentration monitoring

is a known challenge [TC22b], and which e.g. Planar-laser-induced-fluorescence (PLIF) monitoring [Cri08] would allow for under certain constraints. Ensuring that these conditions are met will help validating the robustness and adaptability of the electronic nose in real-world environments, thus providing a more comprehensive assessment of its performance. Further, for the odour classification data features to be fully applicable to dense odour stimuli, it may require altering the pre-baseline normalisation, i.e. replacing the fixed pre-stimulus distance with e.g. a moving average. This would also alleviate the short-term drift that can be observed in the high- and low step values of the temperature cycle responses, see e.g. 3.9a.

The proposed technology and its evaluation hold promise for tackling many real-world challenges that require rapid odour sensing. In particular, any instance of olfactory robotic solutions might currently be cut short in terms of performance; as for both UGVs and UAVs, the sensor response time dictates the maximum speed at which the agent can move while still obtaining spatially resolved measurements [BM20a]. Thus, such applications may directly benefit from using the proposed sensor modalities, allowing for faster identification and localisation of odour sources. For instance, a recent work proposed swarms of nano quadcopters performing gas source localisation in indoor environments [Dui+21], and evaluated different search strategies. A decreased latency in detection and classification may not only assist in more efficient source localisation, but also in expanding the use case to multiple odours and more complex outdoor environments. Another recent study proposes odour sensing on drones for wildfire monitoring [Wan+23]. For detecting smoke; vision and gas sensors are fused, however the gas sensor update frequency is just 1 Hz. Given the intermittent and fine-structured nature of odour plumes, an improvement on sensing timescales could reduce false-negatives and aid in gaining critical time in localising the fire.

Beyond robotics, most applications in security still use static and relatively slow sensing platforms, for e.g. the odour-based detection of explosives [TGK20] at airports. At checkpoints, fast and portable electronic noses could replace random spot checks with exhaustive controls, and thus minimise risk further. Further, recent investigations on mammalian olfactory-guided behaviour use head-mounted MOx sensors as control recordings [Tar+21]. Using a high-resolution data acquisition system — particularly one that matches the temporal capabilities of the subject — would allow for better data quality and hence could improve the resulting models.

In a related vein, neuromorphic information processing [Mea90; Ind+11] has seen much traction in recent years, where in particular the reduced latency, power consumption and data bandwidth have enabled highly optimised vision and auditory sensors [LPD08; LD10]. We suggest that the information embedded in millisecond odour packets, together with the sparse and intermittent nature of odour plume encounters, makes the sense of olfaction an ideal candidate for neuromorphic sensing. We foresee that revealing the rapid nature of the sensors will further stimulate this field of research, motivating event-driven and asynchronous odour sampling [PMG13], for MOx sensors [CSN14; Ras+23] and beyond [Lap+22; Wan+24].

In conclusion, our study marks a groundbreaking advancement in electronic olfaction systems, demonstrating the ability to discern odours and decode odour patterns with unprecedented temporal precision in miniaturised low-power settings. Our findings unlock new possibilities for developing robotic systems capable of rapidly and precisely tracking odour plumes in compact and low-power environments, with the potential to transform electronic nose designs and their applications across various domains.

3.3 Conclusion

In Section 3.1, I introduced a novel design for a portable and high-bandwidth electronic nose prototype, which allows for tight control of the MOx sensor operating temperature. I demonstrated that modulating the sensor heaters at a high frequency can lead to robust data features, which in our case were used to train and evaluate a classifier on different natural olfactory scenes.

Building on those results, Section 3.2 describes the introduction and evaluation of a similar but more compact electronic nose. In an extensive data campaign, I used an odour delivery device with exceptionally high signal fidelity for deploying single odour pulses of various durations, as well as dual-odour pulse trains of various frequencies, to the electronic nose. At pulse durations of 10 ms to 1 s, I tested the efficacy of short (50 ms) heater cycles for performing odour classification. I were able to successfully retain the odour identity with above-chance accuracies down to 10 ms (the shortest tested pulse). For the dual-odour pulse trains, I demonstrated the successful encoding of the modulation frequency up to 60 Hz, as well as the phase (0° vs. 180° shift) up to 40 Hz. I compared the electronic nose's performance with that of mice on similar tasks (see [Ack+21]), and demonstrated that the electronic nose matched or surpassed the mice in every instance.

I argued that the two modes, i.e., constant temperature and short heater cycles, could work well in parallel. While the sensors operating at a constant temperature may be leveraged for encoding the temporal profile of an odour stimulus with high temporal precision, the sensors that operate under temperature cycles may provide data features for odour inference. This dual-pathway finds analogies in biology, where the insect olfactory system is thought to simultaneously extract odour identity and concentration information [GR10; Sch+11].

The proposed fast electronic nose technology shows promise for various real-world applications requiring rapid odour sensing. In robotics, particularly for unmanned ground vehicles (UGVs) and unmanned aerial vehicles (UAVs), faster sensor response times enable higher movement speeds while maintaining spatial resolution. This is crucial for tasks like odour source localisation using small drones, which would highly benefit from reduced detection latency. I identified and discussed potential applications, such as gas source detection or wildfire monitoring, in which rapidly sensing and processing odour information could improve detection accuracy and response times. In security, faster electronic noses could enhance explosive detection at airports by enabling more thorough and efficient screening processes. Furthermore, high-resolution odour sensing systems could improve research on mammalian olfactory behaviour by providing better data quality.

Chapter 4

Towards Neuromorphic Olfaction

"Erlaubt ist, was gelingt."

– Max Frisch

In the previous chapter, we have seen that it is possible to reach animal-like temporal capabilities in odour sensing using portable and low-power sensing devices. The notorious limitations of metal oxide gas sensors, namely their slow response and recovery times, were mitigated by using a) MEMS-based sensors with a small form factor, b) fast peripheral components that allow for tight heater temperature control and rapid sensor readout, and c) novel algorithms for control and signal conditioning. It is expected that this line of work will drive further advances in gas sensor technology and odour-guided robotics that emphasise speed, potentially leading to even shorter sensing timescales.

In robotics, both aerial and ground-based platforms typically operate under highly restricted power budgets and computing resources. Sampling and processing data at high rates is expensive, especially considering that capturing the vast olfactory space occurring in nature would require scaling up the number of sensors tremendously. Building a robotic equivalent of a pollinating honeybee, for instance—an agent that performs highly complex odour navigation tasks on single milliwatts—is currently unimaginable.

In this chapter, I explore how olfactory signals may be sampled and processed more efficiently. In particular, I investigate the applicability of neuromorphic computing, a field of research that emulates structures and processing methods found in biological neural systems and promises to be power-efficient by processing data only when necessary.

In the first part, I will elaborate in general terms on the notion of "neuromorphic olfaction," build arguments based on the physical characteristics of odour plumes on why event-driven data processing may benefit artificial olfaction, and critically review the available literature. In the second part, I will demonstrate different approaches for encoding olfactory information using binary events, based on MOx electronic nose responses.

4.1 Neuromorphic principles for machine olfaction

Given the current availability of computational tools, it is relatively straightforward to apply neuromorphic algorithms—such as the send-on-delta event generation or even artificial spiking neural networks—to a dataset of interest. However, whether those principles are applicable and potentially beneficial will strongly depend on the sensor modality and the observable itself. In this section, I explore why machine olfaction may be particularly well-suited for neuromorphic principles, explore the advantages and limitations of previously proposed approaches, and discuss the potential implications thereof.

The remainder of this section is adapted from the following in-review manuscript:

Nik Dennler, Aaron True, André van Schaik, and Michael Schmuker. “Neuromorphic Principles for Machine Olfaction”, 2024

The co-authors—in the following abbreviated as "we"—contributed as follows to this work: A.T. helped drafting parts of the manuscript by providing insights into fluid dynamics of odour plumes. M.S. drafted parts of the manuscript. All co-authors were involved in discussions during different stages of the manuscript draft, and assisted in editing of the final manuscript. My contributions were the following: Proposing initial argument on applying bandlimitation to machine olfaction, literature review, analysis of different data sets, creating visualisation, drafting the manuscript, as well as editing and revising during submission and peer-review.

Abstract

Neuromorphic computing, exemplified by breakthroughs in machine vision through concepts like address-event representation (AER) and send-on-delta sampling, has revolutionised sensor technology, enabling low-latency and high dynamic range sensing with minimal bandwidth. While these advancements are prominent in vision and auditory perception, their potential in machine olfaction remains under-explored. Here, we outline the perspectives for neuromorphic principles in machine olfaction. Considering the physical characteristics of turbulent odour environments, we argue that event-driven signal processing is particularly suited to the inherent properties of olfactory signals. We highlight the lack of bandwidth limitation due to turbulent odorant dispersal, as well as the characteristic temporal and chemical sparsity. Further, we critically review and discuss the available literature on neuromorphic olfaction, particularly event generation and information encoding mechanisms, and event processing schemes like spiking neural networks (SNNs). The application of neuro-

morphic principles may significantly enhance response time and task performance in robotic olfaction, enabling autonomous systems to perform complex tasks in turbulent environments — such as environmental monitoring, odour guided search and rescue operations, and hazard detection.

Introduction

The sensory capabilities of animals still surpass the limits of technology in many domains. It is therefore not surprising that the underlying principles have long inspired devices and algorithms for sensing in machines. Neuromorphic computing is a prime example of this approach [Ind+11; Sch+22]. In machine vision, bio-inspiration and neuromorphic principles have enabled breakthroughs in sensor technology. In particular, the send-on-delta sampling scheme [Mis06a] and the *address-event representation* (AER) concept [Boa00] have enabled low-latency and high dynamic range sensing, requiring little bandwidth to represent sparse signals [LPD08; PMW11]. Today, the groundbreaking developments in neuromorphic technology have culminated in commercially available products and various applications where event-based vision solves problems that are difficult to tackle with conventional frame-based image sensors [Gal+20; GS24]. The principle of dynamic sampling has since been extended to other sensing paradigms, most prominently auditory perception [Liu+10]. Here, we explore what benefits and opportunities event-based processing may deliver in machine olfaction — the artificial sense of smell.

Why neuromorphic olfaction?

To determine whether neuromorphic computing may benefit machine olfaction, we shall examine the conditions under which such mechanisms have proven effective. A compelling argument by Liu et al. [Liu+10] suggests attributing to a signal its degree of bandlimitation. This may indicate the presence of a stable mapping function between the signal and uniformly spaced samples, or alternatively suggest that dynamically acquired samples could be superior in representing the signal. In vision, for instance, natural scenes contain motions at varying velocities, and choosing a constant camera frame rate that matches the fastest observed motion is inefficient [Liu+10]. More fundamentally, high-contrast (Dirac- or step-function-like) scenes are not bandlimited in space, while any moving object is not bandlimited in time. Here, nonuniform and data-driven/dynamic sampling may be favourable. Contrary, audio signals (such as speech) are inherently bandlimited, irrespective of the recording device. The bandwidth of the signal of interest may be determined in advance, thus allowing the selection of a periodic sampling rate that captures the relevant signal with minimal aliasing.

Temporal and spatial sparsity of signals also suggest event-based processing [OF04; Del+10]. When imaging natural scenes, most regions do not change significantly over short periods of time (temporal sparsity) and contain large areas of similar intensity or colour (spatial sparsity). Both biological and neuromorphic vision systems tend to encode information efficiently by emphasising changes and discontinuities rather than uniform areas, which is aligned with the efficient coding hypothesis [Bar61; SO01]. Thus, signal sparsity aids in explaining the effectiveness of event-based processing for visual stimuli, and may indicate applicable regimes for other sensing modalities [Maa15].

What is the situation in olfaction?

Olfaction is not bandlimited

Under almost all relevant natural conditions, odorants are transported by turbulent dispersal. Turbulent flows are characterised by eddies, i.e., vortices of variable size. Large eddies dissipate their power into smaller eddies [Kol41; Obu41], cascading down to the molecular level, where energy is dissipated into heat through friction [Ors73]. This process creates odour structures spanning a range of length scales, producing characteristic temporal odour fluctuations as they are transported across a sensor.

The statistics of these odour fluctuations depend on a) environmental flow conditions, and b) the odour source configuration. Environmental flow conditions set the relative importance of the physical processes acting on odour structures en route to the sensor. This includes fluid dynamic strain, which locally enhances concentration gradients, and molecular diffusion, which diminishes them. The source configuration includes factors like its dimensions, mass flow rate, buoyancy, and proximity to boundaries. Collectively, these factors produce variance in the odour concentration field, where the size L of an odour plume relative to the eddies predicts what fluid dynamic processes drive local concentration fluctuations [Cas+20]. If eddies are larger than L — which is often the case very close to the odour source — they contribute mostly to plume meandering, i.e. the irregular motion of the fluid volume's centre of mass (see Fig. 4.1a). Conversely, if eddies are smaller than L — as encountered at a greater distance from the source — they contribute to turbulent diffusion characterised by shearing, distortion, and expansion of the plume.

The odour concentration fluctuations in turbulent plumes are distributed as a power law [HI89; MM91] and span several orders of magnitude. In particular, it has been postulated and empirically shown [MM91] that for fully developed turbulence regimes (i.e., far away from the source, where the Reynolds number Re is high), the frequency energy spectrum $nS(n)$ exhibits a log-log relationship with the frequency n as

$$nS(n) = a\epsilon^{2/3}\left(\frac{n}{U}\right)^{-2/3}, \tag{4.1}$$

where a and ϵ are a universal constant and the rate of turbulent energy dissipation respectively, and U is the mean wind speed. This implies that for fluctuations in odour concentration, there is a non-vanishing energy contribution even at very high frequencies. Fig. 4.1b displays a collection of power spectral densities, calculated from data of various studies. All the considered datasets exhibit a pronounced log-log power law — collectively covering frequencies that span five orders of magnitude — however, variations in turbulence regime, sensor modality, and post-processing methods may cause deviations from the theoretical $f^{-2/3}$ slope. Therefore, unlike other sensory modalities such as auditory perception, signals derived from the odour space are not subject to narrow bandwidth limitations.

Olfaction is temporally sparse

Intermittency is a defining feature of turbulent odour plumes [MEC92; CVV14]. Single odour packets (also "whiffs" or "bouts") or clusters of odour packets (also "clumps") are separated by periods of no signal (also "blanks"). The timescales of bout durations and inter-bout intervals, as well as bout amplitude (or odour concentration) and bout-per-clump counts, depend on environmental flow conditions — most notably the structure of the mean and fluctuating velocity fields — and the vector between the encounter point and odour source. Relevant

statistical characterisations of whiffs, clumps, and blanks have been postulated and tested [CVV14]. In particular, both the duration of a whiff t_w (the time between the concentration signal crossing a threshold c_{thr} upwards and downwards) and the upcrossing duration t_u (the time between the upward crossings corresponding to two whiffs) are distributed according to power laws:

$$p\left(t_w\right) = \frac{1}{\tau}\left(\frac{t_w}{\tau}\right)^{-3/2} g_w\left(t_w\right), \qquad p\left(t_u\right) = \frac{1}{\tau}\left(\frac{t_u}{\tau}\right)^{-3/2} g_u\left(t_u\right) \tag{4.2}$$

The functions g_w and g_u describe cut-offs that are exponential for large arguments, while τ is a diffusion time constant. Fig. 4.1c displays field recordings by [Yee+95a], demonstrating the $t_x^{-3/2}$ power law, spanning four orders of magnitude for both whiff duration and upcrossing times. While short odour packages at milliseconds or less [CVV14] are possible and most likely, the probability for extended periods of no incoming odour packages is non-vanishing. In such terms, olfactory signals are sparse in time.

Olfaction is chemically sparse

Chemical space is vast. It has been estimated that the number of theoretically synthesisable drug-like molecules, i.e. that obey certain rules regarding oral bioavailability, approaches 10^{60} [Lip+96; Rey15]. Projects like GDB [Rud+12] aim at enumerating as much as possible of this space, currently containing close to 165 billion small molecules. The subset of molecules perceived as odorous is suggested to count around 40 billion [May+22]. To successfully navigate and interact with this vast chemical space, evolution has equipped the animal kingdom with a large number of olfactory receptor neurons. For instance, fruit flies (*Drosophila*) express around 1500 receptors of 60 different types [Vos00], while honeybees have around 60,000 receptors of 160 different types [PG21]. Those numbers pale in compar-

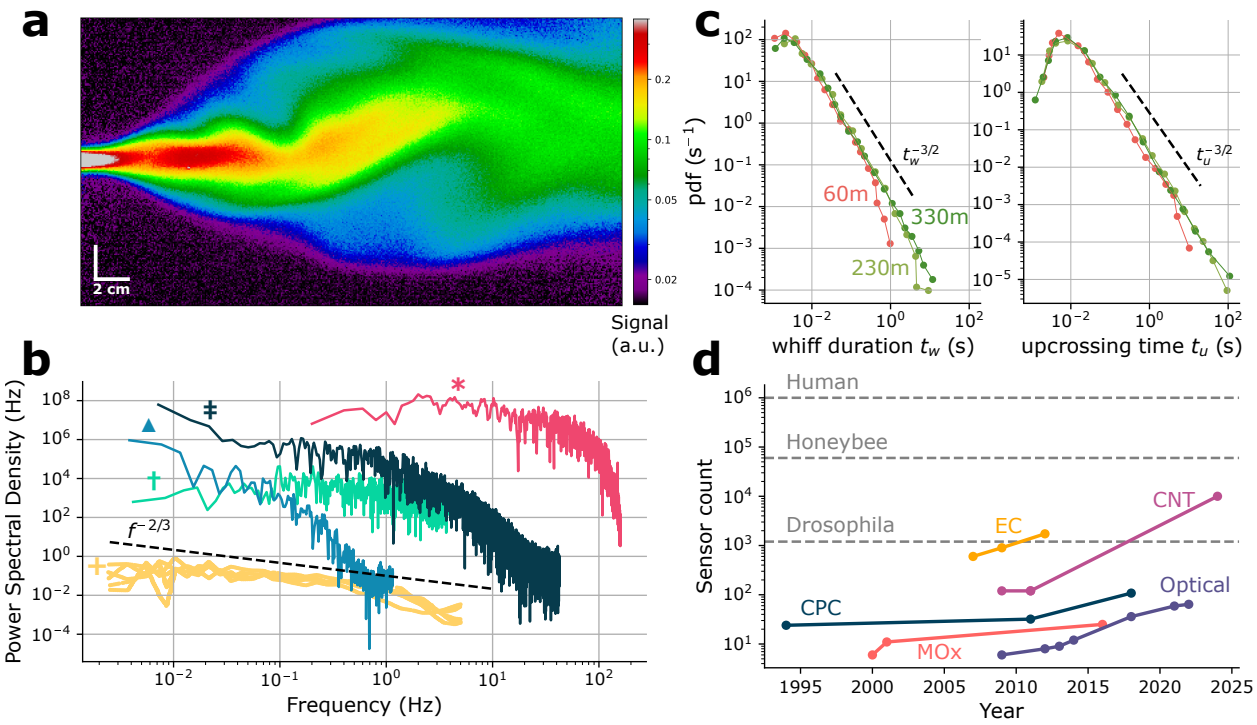

Figure 4.1: **a** Planar-laser-induced fluorescence (PLIF) recording of a meandering odour plume in a wind tunnel. Data from [CMC18]. **b** Energy spectra (variance-normalised) versus frequency, for different studies. Sensor modality, measurement condition, source-to-sensor distance (d), data processing and source as follows: + Photoionisation detector (PID), outdoor field recording, d=200m-300m, filtered data, [MM91]; ▲ E-nose, wind tunnel (WT), d=1.18m, raw data, [Ver+13]; † PLIF, WT, d=20cm, raw data, [CMC18]; * indoor field recording, d=40cm, raw data, [Ack+21]; ‡ E-nose, WT, d=20cm, raw data, [Den+24]. For all raw data recordings, the highest frequency components where strongly contaminated by instrument noise, and cropped out for visual clarity. Dashed line indicates theoretical $-2/3$ slope. **c** Probability density functions of measured whiff duration and upcrossing time, for different distances between source and sensor. Dashed lines indicate theoretical $-3/2$ slopes. PID recordings by [Yee+95a], panel adapted from [CVV14] (CC-BY 3.0). **d** Number of sensors on portable electronic noses of different sensor modalities, as reported in various studies (non-exhaustive, see Table 4.1 for references). Sensor modalities are abbreviated as follows. CPC: Conducting polymer compound; MOx: Metal-Oxide; EC: Electrochemical; CNT: Carbon nanotube. Dashed lines represent total olfactory receptor counts for different animals.

ison to mammals; humans have around a million olfactory receptors of 300 types [Ken20] (postulated to be aligned to key food odourants [Dun+14]), while certain dogs have up to

300 million such receptors [Pho22]. A shared feature of the different olfactory systems is the extreme sparseness in receptor space: Typically, an animal's encounter with an odour evokes only a small subset of the available receptors [Ito+08], often even limited to single receptor types [Dew+18; Bur+22].

In machine olfaction, the term "electronic nose" refers to an array of individually addressable gas or odour sensors. The activation across sensors is often interpreted as a combinatorial code, or odour fingerprint, which correlates with odour identity and concentration [PD82]. Different sensing elements are used, such as electrochemical (EC) sensors, conducting polymer compound (CPC) gas sensors, Metal-Oxide (MOx) gas sensors, optical gas sensors, and carbon-nanotube (CNT) gas sensors. They differ in their capabilities (e.g. sensitivity and response time), constraints (e.g. power consumption, form factor, and sensor drift), and levels of maturity, and are typically selected according to their use cases. In theory, a larger set of sensors with slightly different chemical selectivity would improve the capability of an electronic nose device to discriminate different odorants and mixtures. In line with this postulate, the reported numbers of individual sensing elements embedded on electronic nose devices have been growing steadily (see Fig. 4.1d), most recently surpassing 10,000 sensors [Wan+24]. Analogous to biological olfaction, it has been demonstrated that the activation across large sensor arrays is non-uniform, and considering a spare subset of the available sensors often sufficient [Lap+22; Wan+24].

The odour space is informative, tangible, and suggests event-driven sensing

Intuitively, the notion of turbulence may suggest elusiveness, disorder or even chaos. However, the statistics of odour encounters in turbulent environments encode spatial information of the odour source and its surroundings. In particular, measured bout concentration variations as well as intermittency are indicators of the plume dimensions [FR82b]. The degree of temporal correlation between two encountered odours at a single point can indicate their separating distance [Hop91]. Conversely, analysing the correlation between same-odour encounters at multiple points in space indicates the relative position of the source [Wei+02]. Additionally, several plume features have been shown to reproducibly vary with distance and direction between the sensor and odour source, such as the concentration amplitude and first derivative of a bout [MC04], the degree of intermittency [Rif+14], and the average bout count [SBH16].

While olfaction is temporally sparse, it is not subject to bandlimitations, suggesting that many plume features can be captured reliably only if the plume is sensed and processed fast enough. In fact, stationary measurements of odour concentration fluctuations occur at frequencies exceeding 100 Hz [Yee+95b], while individual odour encounters can last few milliseconds or less [CVV14]. Many animals exhibit remarkable abilities in rapidly detecting and processing short odour stimuli. For instance, insects olfactory receptor neurons' (ORNs) response latency is less than 2 ms [Szy+14; Ege+18], and odour stimuli fluctuations can be resolved at frequencies of over 100 Hz [Szy+14]. This enables them to efficiently track dynamics of fast odour signals [Cri+22], eventually leading to solving complex tasks such as odour source localisation in turbulent environments [BD14; SEE23], or perceptually segregate mixed odours from different sources [Hop91; SEE23; Ack+21]. In machine olfaction, the relatively slow sensor response and recovery times of electronic nose devices have been

prohibitive for many applications. However, recent advances in sensors and processing algorithms have brought their temporal capabilities closer those of animals. For instance, we reported on a miniaturised electronic nose with an odour response time in the millisecond regime [Den+24], outperforming mice in their temporal capabilities on equivalent tasks.

The observed power law in odour concentration extends to small spatial scales, which in turn correspond to very high frequencies. Accurately extracting information requires high temporal resolution, yet efficiency considerations must account for the signals' sparsity in time and sensor activation. We argue that event-driven processing schemes are optimal in this regime. They allow for efficiently handling the observed scale-invariance and heavy-tailed distributions of odour encounters, as they inherently represent information with high temporal precision and the least possible temporal quantisation, while achieving minimal activity in the frequent periods where no odour is present. For instance, a neuromorphic sensor could generate events in response to the onset and offset of odour whiffs, accurately capturing its concentration and duration. During periods of no signal, the system could remain inactive, conserving energy and focusing computational resources on the next significant event. Further efficiency increases could be achieved by processing data from the growing number of sensors activation-driven and in parallel, which neuromorphic algorithms and substrates inherently support.

Neuromorphic odour sampling and signal processing

In neuromorphic vision, a vast majority of studies and applications are based on events that are generated by adaptive threshold crossing (send-on-delta) of a logarithmically scaled brightness signal, and represented via their pixel address $\boldsymbol{x}$ and time t. Most prominently, the dynamic vision sensor (DVS) [LPD08] integrates this in a circuit, and produces events

e_k of the shape

$$e_k(\boldsymbol{x_k}, t_k, p_k), \tag{4.3}$$

where p_k represents the polarity of change, i.e. $p \in [-1, 1]$. For processing, the events are then either represented as an accumulated image frame, or via particular encoding schemes such as time surfaces or voxel grids [Gal+20], or as individual events to be directly processed via spiking neural networks (SNN). Yet, there exists a multitude of other methods for generating, representing and processing events. In the following, we introduce, review, and critically discuss different principles regarding event-based sampling and signal processing for neuromorphic olfaction.

Event generation and information encoding

Biological action potentials

In biological olfaction, similar principles across phyla are observed to generate odour response spikes at the olfactory receptor level. Insects embed olfactory receptor neurons (ORN) in hair-like structures on their antennae, while in mammals ORNs are located in the olfactory epithelium that covers internal parts of the nasal cavity [Kan+00]. In both cases, once an odour molecule has bound to its corresponding receptor [BA91], a signal transduction cascade is initiated, leading to the opening of ion channels that ultimately cause the neuron to depolarise and initiate an action potential (AP). The AP, which can be understood as a binary event, is transmitted to and further processed in the antennal lobe and olfactory bulb, respectively.

Artificial neurons and current injection

Various methods for converting a continuous gas sensor response into binary events have been introduced. One approach (see Fig. 4.2a) involves simulating or emulating the membrane potential V_m of an artificial spiking neuron, and feeding the sensor response as an input current $I(t)$:

$$\frac{dV_m(t)}{dt} = -\frac{V_m(t) - V_{\text{rest}}}{R_m C_m} + \frac{I(t)}{C_m} \tag{4.4}$$

Here, R_m, C_m and V_{rest} are the membrane's resistance, capacitance, and resting potential, respectively. An event is generated whenever V_m crosses a fixed threshold, followed by resetting V_m to its resting potential. Both early and more recent studies in neuromorphic olfaction have leveraged this approach extensively [Pea+05; Koi+06; Koi+07; HK12; SPN14; Dia+16; Jin+17; Han+22; Wan+22b; SMK23] using different neuron models, parameters, and signal conditioning methods. The method is particularly applicable for hardware spiking neural network implementations, as the computing substrate itself can be used to generate the spikes. Conversely, the approach is less suitable for FPGA and digital implementations, as the simulation of differential equations is computationally expensive. Similarly, events can be generated via a biased Poisson process [Pea+01]. This has been utilised to emulate ORN excitations [Jür+21] in neuromorphic olfactory models; however, it could also be used to generate odour response events by biasing the Poisson process towards an average event rate that is proportional to the odour-specific sensor response. All those approaches represent a form of rate code, i.e. the event rate being proportional to the sensor response [Guo+21], or — in the case of multiple sensors — a population code. While offering a noise-resilient and robust encoding, information transmission relies on downstream event integration and tends to be slow.

Latency-representation events

Other information encoding schemes rely on the precise timing of events. For instance, single events may be generated at latencies Δt_{ij} that are proportional to the log-response of sensor j to an odour i, i.e. $\Delta t_{ij} \propto \ln S_j(t)$, which allows considering the relative latency (i.e. the rank) of the different sensors as a concentration-independent odour mapping [Ng+09; NBB11; Che+11; Yam+12]. Similar methods that use transient features instead of the raw signal have been proposed [Has+15; Huo+23]. Albeit being potentially fast and highly sparse, this encoding breaks down at single noise events, and in the reported form is not suited for dynamically changing signals due to the lack of a reset signal. Similar, yet more sophisticated methods have been proposed, where transient features are converted into events of different latencies [Van+20; Yan+23], to be further processed by an SNN.

Change detection events

Another approach — as prominently featured in event-based vision — is using the send-on-delta principle [Mis06a] to generate events that correspond to positive or negative changes in the sensor response [Pea+13; Van+19; Den+22a; Van+22]. In particular, if the signal crosses a reference voltage plus or minus a set threshold, a positive or negative change event is generated and the reference voltage shifted up- or downwards by the threshold. In vision, the signal is typically preconditioned through logarithmic scaling, which increases the dynamical range drastically [LPD08]. In olfaction, however, this may not be necessary, as the transfer function of many gas sensors implicitly scales the encountered odour concentration range.

The send-on-delta method is fast, sparse, and robust, yet as it tracks changes only, it does not provide the absolute signal intensity. In vision, intensity information is often not needed, however is retained in the ATIS pixel design [PMW11]. Besides generating positive and negative change detection events, the mechanism employs an extra channel containing pairs

of events, with a time difference inversely proportional to the absolute signal intensity (see Fig. 4.2a). This concept loosely resembles the dual-pathway observed in the mammalian olfactory bulb [Fuk+12], which has inspired at least one olfaction study [Ras+23] that adapted the ATIS circuit for event-based gas concentration measurements. A sparser and more straightforward event representation is achieved by providing separate event channels for fixed amplitude levels, as demonstrated first for an event-based speech processing audio front end [GS09]. Whenever an amplitude level is crossed upwards or downwards, an event in the corresponding channel is generated (see Fig. 4.2a). This method may be particularly effective whenever the amplitude range of the signal is known beforehand, which is often the case for gas sensors.

Events encoding turbulent plume features

Given the intermittent nature of turbulent odour plumes (see Section 4.1), it may be reasonable to encode temporal features of the plume, such as bouts and blanks, in the event trains. Recent evidence from neuroscience suggests that animals apply similar principles. For instance, large-scale temporal odour features — specifically plume onset, plume offset, and whiff encounter — are encoded in the mammalian olfactory bulb spike output, both at the single-cell and population levels [Lew+24]. These features are critical for odour-guided navigation [Par+16; Kad+22] and source localisation tasks [Rig+22a]. Such an approach has been demonstrated by isolating the bout onset from the sensor response through Kalman filtering, and subsequently generating change events through send-on-delta sampling [DS21]. In a stereo sensing setup, this method allowed for estimating bout velocity by considering the

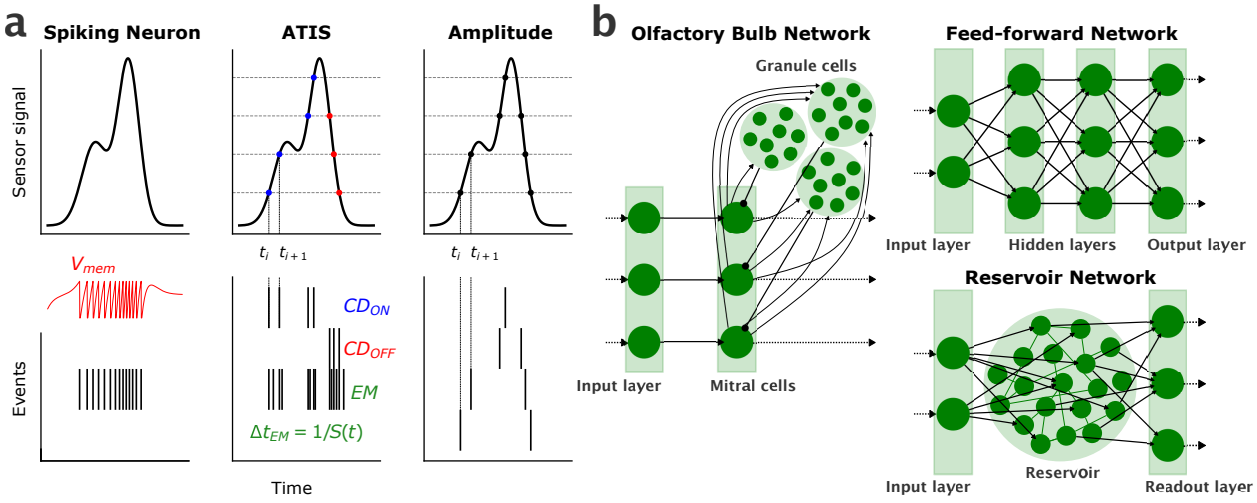

Figure 4.2: **a:** Event generation mechanisms. From left to right: A spiking neuron, described by differential equation for the membrane potential V_{mem} (red, see Eq. (4.4)), generating events whenever V_{mem} crosses a thresholds. ATIS mechanism, producing positive (CD_{ON}) and negative (CD_{OFF}) change detection events whenever an adaptive threshold is crossed, and an exposure measurement (EM) event pair that encodes the signals' absolute intensity $S(t)$ in its time difference Δt_{EM}. The algorithm can be reduced to the familiar send-on-delta mechanism by omitting the EM events. Amplitude crossing mechanism, generating events whenever fixed amplitude levels are crossed, in channel that are specific to those levels. **b:** Spiking neural network architectures. Arrows and point-ends indicate excitatory and inhibitory connections respectively. From left to bottom right: Olfactory bulb network, feed-forward network, reservoir network.

relative event latencies between sensors [DS21]. Other methods that approximate the plume dynamics from the sensor response, such as cascaded filtering [SBH16] or blind deconvolution [MBM19], could be combined with send-on-delta or ATIS mechanisms for generating informative and task-relevant event patterns.

Event-based signal processing algorithms

Biological neural networks

Evolution provides us with effective mechanisms for processing event trains generated by thousands to millions of olfactory receptor neurons (ORN). ORNs that express the same type of OR converge onto spatially separated structures called glomeruli, which are found in

both the mammalian olfactory bulb (OB) and the insect antennal lobe (AL). Interneurons, such as the mammalian granule cells and periglomerular cells, modulate glomerular activity through inhibitory and excitatory interactions, which enhances odour separability and noise resilience. Projection neurons, such as mitral cells, relay the glomerular activity to higher brain centres, which is the olfactory cortex in mammals, and the mushroom bodies and lateral horns in insects.

Spiking neural networks (SNNs)

In artificial olfaction, besides directly considering one of the discussed encoding schemes, such as the average event rate or the order (rank) of incoming events, the typical approach to processing asynchronous events is using spiking neural networks (SNNs). SNNs represent a class of artificial neural networks that mimic biological neural systems more closely than traditional neural networks, such as multilayer perceptrons. The key difference is that SNNs typically incorporate the precise timing of events in their operation and activate subsequent nodes only when a fixed activation threshold is crossed, allowing for asynchronous and data-driven computing [Maa97]. SNNs build the foundation for (non-von-Neumann) neuromorphic computers, which are characterised by their highly parallel operation and collocated processing and memory units [Sch+22].

Bio-inspired SNNs

Various studies have focused on implementing concepts from biological olfaction systems, either in simulation or hardware. Most of these studies concentrate on the mammalian olfactory bulb (OB) and glomeruli models, which perform a form of dimensionality reduction through excitatory and inhibitory dynamics. Early works implemented OB models on hardware VLSI and trained them either via Hebbian learning [Pea+05] or Spike-Timing-

Dependent Plasticity (STDP) [Koi+06; Koi+07], or self-organizing models [Ram+06]. The STDP approach has since been integrated into low-power electronic nose systems [HK12; Yan+23], demonstrating effectiveness in tackling sensor drift [BC19]. A noteworthy SNN implementation of an STDP-based glomerulus model that couples its activity with gamma oscillation promises to perform noise-robust learning from single data instances [IC20]; however, the reported evaluation protocols are not conclusive [DSS23]. Recent efforts focus on few-shot class-incremental learning [Huo+23] or on fusing the olfactory bulb network with an SNN that is driven by visual stimuli [Dai+23]. Other biology-inspired studies choose to implement parts of the insect antennal lobe (AL). Early works implemented simplified AL models on VLSI hardware [Bey+10] and on neuromorphic computing substrates [Pfe+13], which have been extended to more complex models that focus on excitatory-inhibitory dynamics [Jür+21]. Further studies implemented AL models and included learning via STDP [Pea+13], supervised soft-winner-take-all mechanisms [SPN14], and combinations of unsupervised self-organisation processes and supervised synaptic plasticity rules [Dia+16]. A noteworthy contribution has been made by further abstracting the biological networks; implementing a two-layer SNN that performs concentration-invariant odour detection through event synchrony [SMK23].

Feed-forward SNNs

Feed-forward SNNs are more commonly used for neuromorphic vision or auditory perception, as they can simplify model implementation and training procedures by transmitting information in just one direction, i.e. with no recurrent connections. See Fig. 4.2b for an illustration of the concept. In neuromorphic olfaction, similar approaches have been demonstrated. For instance, a two-layer SNN was implemented and trained via STDP on a neuromorphic chip [Van+19], and then used to differentiate between roasted malt samples [Van+22]. Further, a

four-layer SNN was trained from input neuron activation (offline, via backpropagation), then demonstrated to be of practical use on a food science task [Han+22]. While feed-forward SNNs offer simple training, they do not support the recurrent dynamics observed in biological olfactory circuits and hence might be limited in their ability to efficiently deal with olfactory signals.

Reservoir SNNs

A reservoir network, or liquid state machine, is a computing concept that can be implemented using SNNs: An input layer projects onto a reservoir of neurons, typically consisting of sparse, recurrent, and excitatory-inhibitory interconnections using randomised and untrained weights. A readout layer connects to a subset of these neurons, with weights that are trained via supervised learning. See Fig. 4.2 for an illustration. This method allows for casting the input to a spatially and temporally higher-dimensional space [Tan+19; Sch+22]. The concept has inspired reservoir SNNs with cubic network connectivity used for event-based odour data classification [Kas+16; Van+20], as well as implementations based on memristive devices [Wan+21b]. A recent review argues that reservoir computing is particularly well suited to match the physics in photoelectrochemical devices [Abd+23], which may make it an intriguing candidate for emerging optical odour sensing technologies [Lap+22].

Discussion

Based on the physical characteristics of turbulent odour environments, we reason that event-based signal processing appears highly beneficial for machine olfaction. While odour signals are spanned by a space of billions of different odourants, they are presented at fluctuations following power law distributions across multiple orders of magnitude. Event-driven pro-

cessing schemes offer an optimal solution by capturing salient features at high temporal precision, while energy consumption and data bandwidth can be preserved during idle periods or for inactive sensor nodes. This approach aligns well with biological systems' efficiency in handling sparse and dynamic signals.

Despite a key benefit of event-driven processing being the viability of exceptionally high temporal resolution while retaining low data bandwidth, none of the reviewed studies in neuromorphic olfaction have emphasised on speed. This comes at little surprise, as machine olfaction has only recently seen timescales that approach the temporal capabilities of animals [Den+24]. However, the presented methods for asynchronous sampling and processing maintain their demonstrated validity, and could be adapted with little effort to state-of-the-art gas and odour sensors. In particular, the change-detection method and its variants may lend themselves as ideal candidates for producing efficient encoding of rapidly changing sensor dynamics. Alternatively, events that explicitly encode temporal odour plume features may achieve high information density, as the encoding would be tailored to the underlying physical processes. For processing the generated events, asynchronously driven spiking neural networks can be used. Networks with some degrees of recurrency, such as SNN implementations of the antennal lobe or the olfactory bulb, or the more abstract reservoir SNN, may be particularly well suited. Asynchronously driven spiking neural networks can be used to process those events, where networks with some degrees of recurrency, such as biology-inspired networks or reservoir networks, may be particularly well suited.

Neuromorphic olfaction holds promise for tackling many real-world challenges that require rapid and efficient odour sensing. A set of particularly intriguing applications are found in olfactory robotics [LLD06], where unmanned ground or aerial vehicles (UGV / UAV) perform odour source localisation and navigation tasks [KR08; Fra+22]. Examples of such are using smoke-sensing drones for wildfire monitoring [Wan+23], swarm-based gas source localisation in indoor environments [Dui+21], as well as distributed gas discrimination and mapping in emergency response scenarios [Fan+19]. Those tasks rely heavily on sensing the environment fast and efficiently, under restrictive constraints in form factor and power budget. Equipping robots with neuromorphic olfactory sensing and processing capabilities may lead to significant improvements in response time and task performance metrics, enabling operation in complex and turbulent environments, as well as allowing for highly optimised power management strategies.

Supplementary Materials

Year	Count	Sensing principle	Authors
1994	24	Conducting polymer compound (CPC)	Gardner et al. [Gar+94]
2011	32	Conducting polymer compound (CPC)	Ryan et al. [Rya+04]
2013	4096	Conducting polymer compound (CPC)	Marco et al. [Mar+14a]
2011	120	Conducting polymer carbon nanotube (CNT)	Wang et al. [Wan+11]
2024	10000	Conducting polymer carbon nanotube (CNT)	Wang et al. [Wan+24]
2007	600	Electrochemical (EC)	Che Harun et al. [Che+07]
2009	900	Electrochemical (EC)	Che Harun et al. [Che+09]
2012	1728	Electrochemical (EC)	Che Harun et al. [CCG12]
2000	6	Metal Oxide (MOx)	Hong et al. [Hon+00]
2001	11	Metal Oxide (MOx)	O'Connell et al. [OCo+01]
2013	96	Metal Oxide (MOx)	Marco et al. [Mar+14a]
2009	6	Optical	Tang et al. [Tan+09]
2012	8	Optical	Thepudom et al. [The+12]
2021	59	Optical	Maho et al. [Mah+21]
2022	64	Optical	Laplatine et al. [Lap+22]

Table 4.1: Non-exhaustive selection of reported electronic nose sensor count per dye, sorted by sensing principle and year of publication.

4.2 Event-encodings of MOx electronic nose responses

This section loops back to the data campaign described in Section 3.1, where different urban olfactory scenes were measured using a MOx electronic nose prototype and sub-second heater cycles. Here, the collected data was used to investigate the efficacy of different event-encoding schemes. In particular, I generated events from heater cycle data via send-on-delta, then compared the performance of a linear classifier using the rate, the time-to-first-event, the event-rank, and the step-wise reconstructed signal.

The remainder of this section is adapted from the peer-reviewed conference proceedings, with permission from ACM Digital Library:

Nik Dennler†, Damien Drix†, Shavika Rastogi, André van Schaik, and Michael Schmuker († denotes equal contribution). "Rapid Inference of Geographical Location with an Event-based Electronic Nose." In: *The 9th Annual Neuro-Inspired Computational Elements (NICE), 2022, UTSA, USA Issue: 1.* Vol. 1. Association for Computing Machinery, 2022

The co-authors—in the following abbreviated as "we"—contributed as follows to this work: D.D. designed the electronic nose, wrote the acquisition software, collected the data, and drafted the the manuscript methods sections on e-nose design and data collection. S.R. provided feedback and discussion throughout the project. A.v.S. and M.S. supervised the project. All co-authors conceptualised the study and assisted in editing the final manuscript. My contributions were the following: Suggesting the use of short heater cycles for the analysis of the olfactory scene locations, performing the formal analysis — including proposing novel feature extraction methods, testing different event-generation methods, and designing the evaluation pipeline — and drafting the manuscript.

Abstract

Sensory information is crucial for the successful interaction of an agent with its outside world, where its encoding remains an open research question. Particularly resolving rapid fluctuations in odour plumes has received little attention and could benefit from biologically inspired coding schemes. State-of-the-art gas sensors actively modify the sensing site using temperature modulation cycles, which decreases the integration time and increases the discriminability. Yet it remains unclear how much information is present in one cycles' sensor response and how to efficiently sample it. In this work, we propose a novel approach for asynchronous event sampling for gas sensor data, and investigate the effectiveness of different event encoding schemes for solving an inference problem. An multichannel heater-modulated electronic nose is used to record field data at high frequency. Single-cycle sensor conductance windows of 140ms are normalised and a model curve subtracted. Using send-on-delta sampling, on- and off-events are generated and further encoded in either their rate, time-to-first-spike, firing-order or in a reconstructed signal. The different representations are compared by considering the geographical location classification results using linear SVM. Given the small sampling time window, we report a surprisingly high classification accuracy, both for the raw and the reconstructed signal. The high-compression encoding schemes cannot match the reconstructed signal for most cases, hinting at the highly complex temporal dynamics of odour signals. We conclude that heater-modulated gas sensors lend themselves to cycle- and event-based processing, allowing for inference in the sub-second regime. If our work can be extended from distinguishing broad olfactory scenes to recognising individual odorants in turbulent plumes, this would open up a new range of potential use cases for gas sensors, both for traditional gas sensing and for neuromorphic olfaction.

Introduction

Sensory information is crucial for the successful interaction of an agent with its outside world. How sensory stimuli are optimally encoded remains an open research question. Neuro-inspired coding strategies have been successfully applied to many tasks dealing with vision, auditory perception, touch and vibration sensing [LPD08; SL05; Bar+16; Den+21a]. For the sense of olfaction, there still exists a large performance gap between artificial and biological systems [Cov+21]. This is particularly evident if one considers the fast-changing odour distribution caused by air turbulences. The fluctuation frequencies in an odour plume are governed by a power law [MM91] and can carry essential information about the odour source [SBH16].

While mammals can discriminate temporal correlations of rapidly fluctuating odours at frequencies of up to 40 Hz [Ack+21], metal-oxide (MOx) gas sensor based olfactory systems usually have response times that are several orders of magnitude slower [PLT12b]. Methods for improving the response time have been investigated [Gon+11b; DS21].

Latest-generation MOx gas sensors actively modify the sensing site using temperature modulation cycles, which decreases the integration time and increases the class discriminability [Ver+14]. Yet it remains unclear how much information is present in one cycle's sensor response and how to efficiently sample it. Here, contrary to a top-down approach from biological olfaction to neuromorphics [IC20], we propose a data-driven asynchronous event-sampling strategy for state-of-the-art gas sensors, and investigate the effectiveness of different event encoding schemes for solving an inference problem.

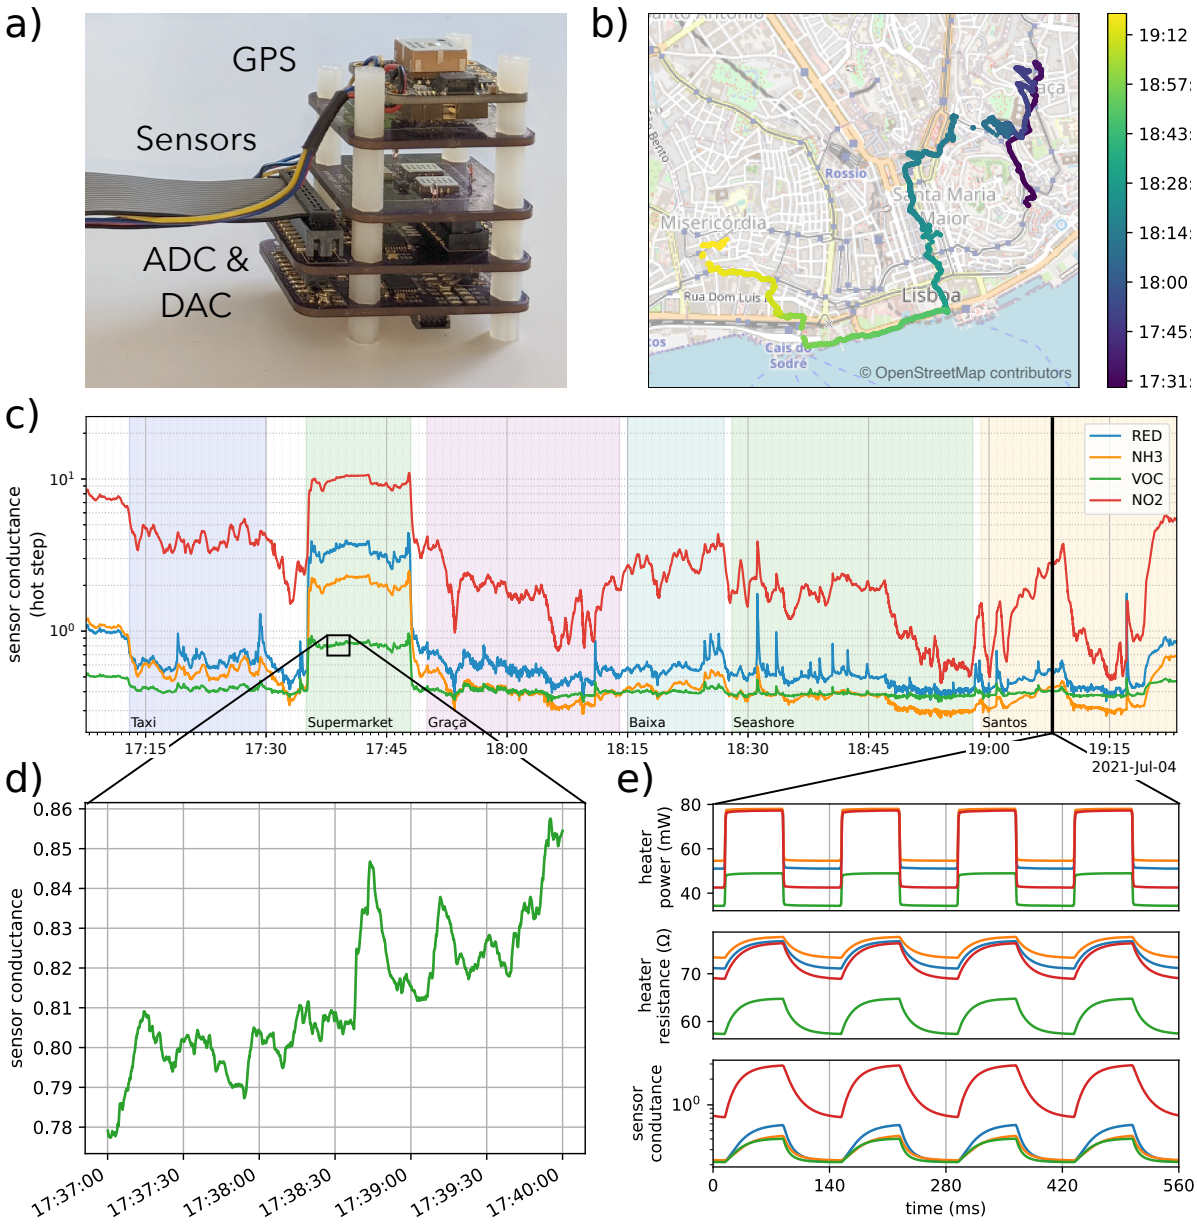

Figure 4.3: a) Portable multichannel e-nose. b) GPS track of the recording. c) MOx sensor data, showing the sensor conductances at the end of the high-temperature step. d) Zooming in reveals structure on a timescale of seconds to minutes. e) Cyclic heater power modulation (top) drives sub-second oscillations in heater temperature (middle, shown as heater resistance) and sensor conductance (bottom).

Methods

We constructed a portable electronic nose equipped with four different MOx gas sensors (*SGX Sensortech* `MiCS5914` & `MiCS4514` dual sensor, *ScioSense* `CCS801`) and recorded the natural olfactory scenes encountered during a walk through the city of Lisbon, Portugal, dividing the dataset into six geographical locations (hereafter labelled 'Taxi', 'Supermarket', 'Graça', 'Baixa', 'Seashore', and 'Santos'). We sampled the conductance of the gas sensing elements at 1 kHz while modulating the heater power with a period of 140 ms, each heater cycle consisting of a high-power step followed by a low-power step. This causes the sensor conductance to oscillate in a way that depends on the heater temperature, the environmental conditions, and the gases present in the sensor cavity (fig. 4.3).

We then investigated whether we could recover the geographical label from the time course of the sensor conductance during a single 140 ms heater cycle. Crucially, each cycle was normalised to the same minimum and maximum values (fig. 4.4a), thus getting rid of the baseline drift that often compromises gas sensor datasets [Ver+12; Den+21b], leaving only the intra-cycle variations to distinguish between different locations. For each normalised cycle, we construct a signal that highlights these intra-cycle variations by subtracting a sensor-specific model curve, which, in our case, is the ensemble average of the normalised cycles across a subset of the data (fig. 4.4b). We apply an algorithm based on *send-on-delta sampling* [Mis06b; VK07] to generate up- and down-events when the signal changes, exploring a range of spike (event) thresholds. These events are then used to compute four different features: the channel-wise event rate (rate code), the channel-wise time-to-first-spike (latency code, [VGT05]), the channel order of first spikes (rank code, [TDV01]), and a signal reconstruction using the reverse *send-on-delta sampling* algorithm (fig. 4.4c). The representations resulting from each of the four encoding strategies are divided into training

and test sets with a ratio of 75% to 25%, by sampling a total of 2000 cycles per class from time-separated bulks as described in [Asa+17]. For each spiking threshold and each encoding strategy, a linear-kernel Support-Vector-Machine (SVM) [CV95] was fitted to the training set and validated on the test set.

Results and Discussion

All four encoding strategies perform better than chance levels (fig. 4.4e). The signal reconstructed from events performs as well as the original intra-cycle signal when the event count is high (82.5±1.0% vs. 84.2±1.2%). Accuracy then decreases as the event threshold increases (fewer events, see fig. 4.4d). While latency code, rank code and rate code representations provide a high signal compression (one value per up- and down channel for each sensor), they are outperformed by the reconstructed curve, until the number of events was reduced to 1% of the original signal's sample count.

Our findings indicate that various olfactory scenes can be distinguished based on the differential time course of sensor conductance during individual sub-second heater cycles, despite a normalisation procedure that removes information about absolute sensor conductance. Furthermore, the full temporal pattern of events within each cycle seems to matter for scene recognition. This hints at a phasic component in the sensor response to a temperature step that contributes to the classification accuracy. Whether this phasic component stems from the temperature-specific reactivity of various gases or from other influences not excluded by our normalisation procedure remains to be investigated.

Conclusion

We propose an event-based sampling scheme to represent cyclic heater-modulated electronic nose data. Asynchronous sampling captured the signal's temporal dynamics better than rate-, rank- or latency-codes, while recognition accuracy degraded gracefully for reduced event counts. Our work paves the way for event-based recognition of natural odour scenes in uncontrolled environments, breaking new ground in neuromorphic gas sensing.

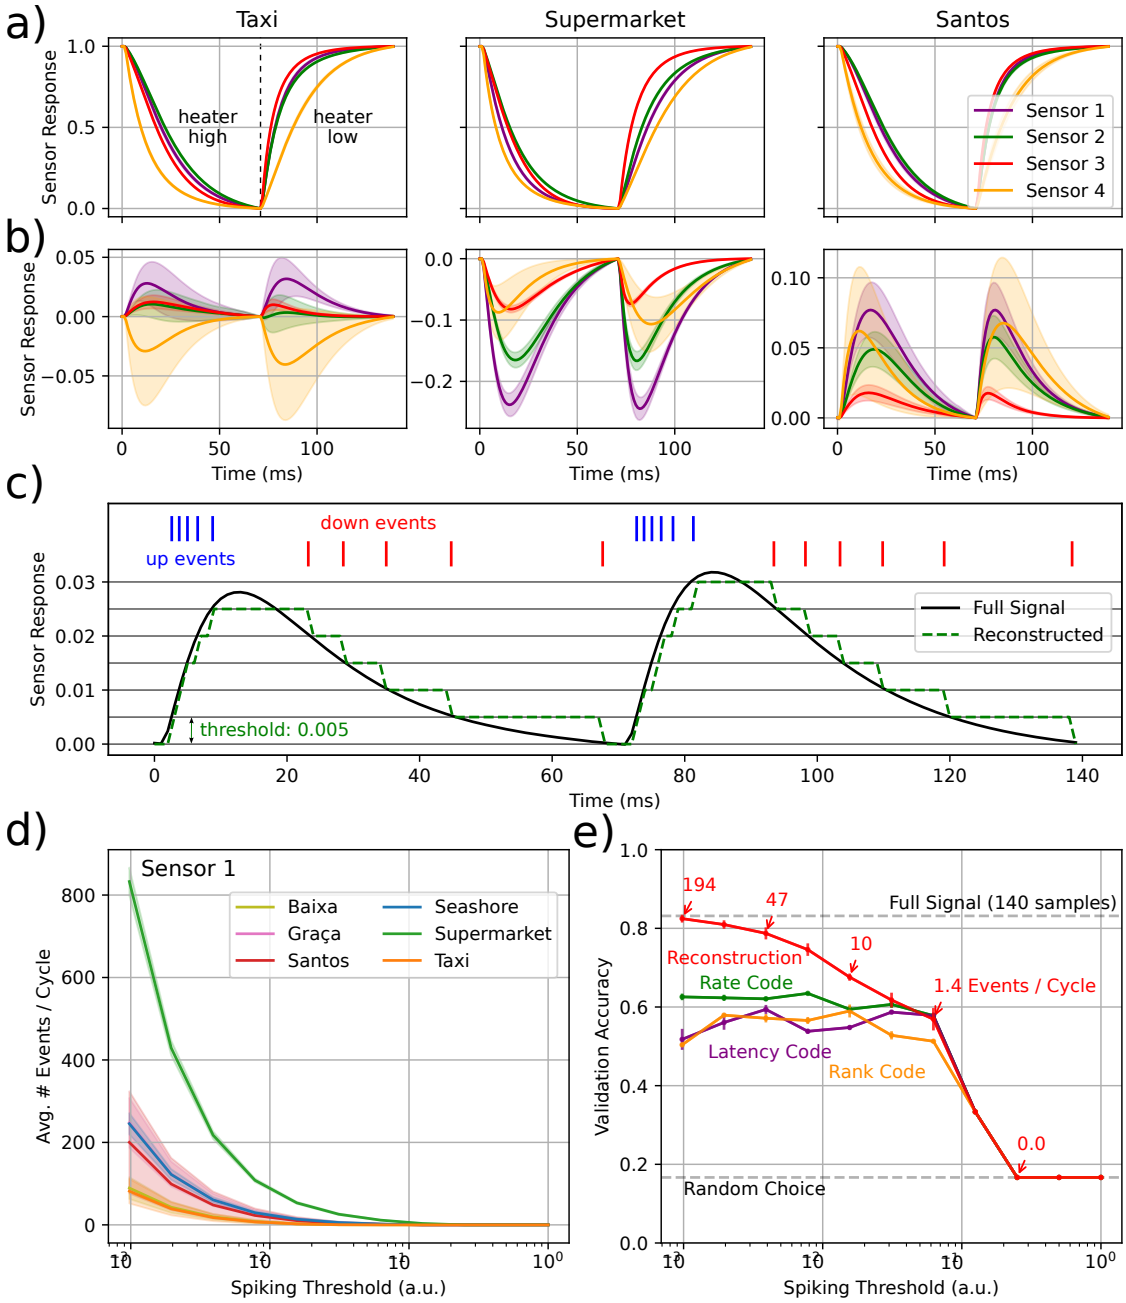

Figure 4.4: a) The normalised sensor conductance oscillates between fixed minima and maxima. b) After subtracting the mean, intra-cycle variations reveal distinct patterns across geographical locations. Solid lines and shading correspond to mean and one standard deviation. c) Example event generation and reconstruction (Sensor 1, location: taxi) d) The average number of events per cycle decreases with increasing spiking threshold, but varies across locations (data for Sensor 1). e) Classification accuracy for different encoding schemes vs. spiking threshold (mean and standard deviation for twelve train/test splits).

4.3 Conclusion

In this chapter, I explored some of the opportunities and premises of applying principles of neuromorphic computing to artificial olfaction, and particularly to fast odour sensing.

In Section 4.1, I initiated a discussion about how the physical characteristics of turbulent odour plumes suggest event-driven sampling. In particular, I noted that olfactory signals are not subject to narrow bandlimitations. This resembles the situation in vision and sets olfaction apart from narrow-bandlimited senses like auditory perception. Further, I noted that temporal signal sparsity is an inherent property of olfaction, with odour bout durations and inter-bout intervals spanning from sub-milliseconds to minutes or above. Therefore, it is either lossy (low rate) or inefficient (high rate) to extract information from the odour plume at a fixed sampling rate. Finally, I observed that the reported number of sensors in electronic nose studies has been growing steadily, most recently surpassing 10,000 sensors. While this increase in dimensionality is reasonable in light of the vast complexity of the odour space, it is expensive to sample a high number of instances at a high sample rate.

These points may be effectively addressed by different methods of event-driven and asynchronous data processing. From such, I identified that variants of the change-detecting send-on-delta algorithm may be particularly suited to generate binary events from the sensor response. These events may be further processed, where spiking neural networks are the obvious candidate whenever low latency is essential. This is the case in many robotic applications, where the time for sampling and processing directly influences the response time and task performance of artificial agents.

In Section 4.2, change-events were generated from electronic nose sensor response curves that are phase-locked to heater cycles. In particular, events were generated by applying a variant of the sigma-delta algorithm to the sensor response. From these events, different event-encodings were constructed— in particular, channel-wise event-rate, time-to-first-spike, first-spike-order, and the step-wise reconstructed signal. The representations were evaluated on an odour signature classification task, where it became evident that none of the signal-compressing methods were able to outperform the reconstruction. This strengthens the argument that directly processing the events, e.g., via a spiking neural network, may be necessary and viable for capturing the temporal dynamics that are characteristic of a given odour presence. However, further investigations will be necessary to elaborate in detail on this.

Chapter 5
Discussion and Conclusion

"Alles Wissen
und alles Vermehren unseres Wissens
endet nicht mit einem Schlußpunkt,
sondern mit einem Fragezeichen."

– Hermann Hesse

5.1 Summary

This thesis explored principles and concepts for achieving high temporal resolution in odour sampling and processing. I began with a review on fast olfaction in both biological and artificial contexts. I recognised that sampling the statistics of turbulent odour plumes requires fast sensing, which is found in both insects and mammals. I discussed the temporal capabilities of different animals, as well as their underlying mechanisms and roles for different tasks. Further, I explored mechanisms of artificial olfaction, with an emphasis on the state-of-the-art in electronic nose technology and its applications.

Initial studies on datasets and algorithms were meant to provide a foundation for our own work on fast odour sensing. However, during this process, I identified—and meticulously described—previously unnoticed limitations in a widely used gas sensor dataset [Ver+13]. Namely, the non-randomised measurement protocol coupled with severe sensor drift rendered the data unusable for most classification benchmarks. Further, I described how a particu-

larly prominent neuromorphic few-shot odour-learning algorithm study [IC20] is affected by this and revealed further limitations thereof. I recognised the need for reliable data and reproducible algorithm evaluations and established a set of best practices for future gas sensor data campaigns.

While these limitations prohibited the usage of the discussed data and algorithms for this thesis, they motivated our own data campaigns. These were performed using a custom-made electronic nose system, which allowed to test different data acquisition and processing methods. In particular, I proposed a novel way of acquiring data features for odour classification, which is based on rapid (50 ms - 150 ms) cycles of the temperature at which the gas sensors are operated. With this, my colleague D.D. recorded two datasets that contain different indoor and outdoor olfactory scenes, which could be robustly distinguished using these data features. Further, I performed an extensive laboratory data campaign, in which I evaluated the electronic nose on a benchmark that has previously been used to shed light on the temporal odour discrimination capabilities of mammals [Ack+21; Das+22]. I showed that, if the sensor temperatures are held constant, it is possible to distinguish correlated odour pulse trains from anti-correlated ones up to modulation frequencies of 40 Hz and determine the frequency up to 60 Hz. This matches and exceeds the capabilities demonstrated in mammals. Further, I showed that, if the sensor temperature is modulated with 50 ms duty cycles, it is possible to classify different odours at pulse widths as low as 10 ms. Both of these results are unprecedented in artificial olfaction and will allow for novel use cases and applications.

Finally, I investigated how to process olfactory signals fast and efficiently. I discussed how principles from the field of neuromorphic computing may be particularly suited for machine olfaction. In general terms, I argued that the physical characteristics of turbulent odour plumes suggest asynchronous sampling and data processing. Further, I critically discussed different methods and algorithms for generating events from gas sensor data, as well as for processing such data. Subsequently, I described how the sensor response of a heater-cycled MOx sensor e-nose may be sampled via asynchronous events, and investigated the efficacy of different event encoding schemes for odour scene classification. Consistent with previous discussions on event-based processing methods in olfaction, I demonstrated that the temporal sensor dynamics may not be captured sufficiently well with a compressed code such as rate-code, order-code, or time-to-first-event, suggesting the use of spiking neural networks for analysing the event trains directly instead.

5.2 Future work

Cycled vs. constant operating temperature for MOx sensors

In Section 3.2, I described how I operated the MOx sensors either at a constant temperature or using heater cycles. I demonstrated that short heater cycles may lend themselves to phase-locked data features that are well suited for odour classification tasks. In contrast, keeping the heater constant allowed for decoding the temporal characteristics of a stimulus, such as pulse-train frequency or the phase between two stimuli. This leads to the hypothesis that, given the MOx sensors I used, these two modes are optimal for the respective tasks;

i.e., heater cycles for odour classification and constant heat for decoding temporal odour structure. To test this, it may be necessary to perform a control experiment, e.g., attempting to decode temporal characteristics using heater cycles, and using the constant heater mode for odour characteristics.

If the hypothesis were supported by such evidence, it may be of interest to combine the two modes on the same device. For instance, a set of sensors could be operated at a constant temperature, thus continuously monitoring the temporal dynamics of the sensor response. Another set of sensors could be operated using heater cycles. Processing the sensor responses of the latter could be triggered by the constant temperature sensors reporting a significant deflection from the baseline, which may correspond to an incoming odour bout. This would preserve computing resources and additionally may reduce false-positive classifications. Going further, a similar protocol may foresee that the heater-cycled sensors are initially in a stand-by mode or even completely turned off, and only turned to cycles whenever a bout is registered. If successful, this may bring substantial power savings, as the classification sensors would not have to be heated during idle times. These savings may allow for scaling the dimensionality of an electronic nose more favourably, as heater-cycled sensors could be added without linearly increasing the device's power consumption. However, MOx sensors are known to require a warm-up time to provide reliable sensor outputs, thus such a protocol may demand further research on adequate feature extraction in the warm-up phase.

Further, to fully understand the potential of the fast electronic nose, extensive field testing should be performed. While the experiments described in Section 3.2 provide detailed insights about the electronic nose's performance under laboratory conditions, it is important to note that, in realistic scenarios, odours do not occur in square pulses and likely at much

lower concentrations than the ones tested. Instead, a thorough evaluation of the measurement protocols and processing algorithms on odour plumes may yield a better estimation of potential task performances. As it is not feasible to control the stimulus directly, it would be necessary to have control measurements in place, such as PID recordings co-located with the electronic nose, or simultaneous Planar-Laser-Induced-Fluorescence (PLIF) concentration mappings.

Towards applications in robotic olfaction

In Section 3.2, I discussed how fast artificial olfaction may allow for tackling many real-world challenges, particularly in the field of mobile olfactory robotics. I noted that any robotic system leveraging olfactory sensing for a given task may benefit from faster sensor modalities, as the sensor response time dictates the maximum speed at which the agent can move while still obtaining spatially resolved measurements. Therefore, future work should focus on integrating the discussed high-speed electronic nose instances with autonomous systems, such as ground robots and drones.

Before placing the electronic nose on a robotic platform, it would be appropriate to experiment with odour exposures that resemble the natural environment more closely than the sharp pulses and pulse trains used in Chapter 4b. Experiments based on a wind tunnel, or an open system using odour sources and e.g. a fan may be considered. The major challenge there would be obtaining either a reproducible and controlled stimulus, and / or verifying the (ground-truth) stimulus without much interference. Further, decorrelating airflow from odour occurence may be an issue. In the following, I briefly elaborate on two possibilities, and their potential limitations:

- Delivering turbulent odour stimuli via a wind tunnel or open setup, then monitoring the concentration via miniPIDs in combination with the electronic nose would be a possible solution. However, it is known that the sensing principle of PIDs is active, i.e. they will disturb the flow [TC22a]. Hence, this approach might lead to unwanted biases in the results, where this effect is stronger at close proximity between PID and electronic nose.

- Another approach is using a passive odour monitoring technique,such as Planar-Laser-Induced-Fluorescence (PLIF) [Cri08; CMC18]. A closed wind-tunnel may deliver turbulent airflow, where fluorescent odours could either be injected or be released inside the tunnel. A laser sheet would be spanned in a 2D plane, in which the passing odour would be made visible. One drawback is that this would only work with the small subset of odours that fluoresces under certain wavelengths, e.g. acetone, which would limit the broad applicability of this approach.

Further, for both ground-based and aerial robotic agents, a key challenge will be overcoming the strong coupling between sensor response and temperature fluctuations, which are inevitably induced by the airflow from agent movement. It may be that the proposed sensor heating and post-processing normalisation scheme, or variants thereof, might address such environmental interferences. Additionally, exploring adaptive sensor placement and protective enclosures could further help mitigate the issue. Yet, future investigations will be necessary to explore such scenarios.

Another promising direction may be to combine fast olfactory sensing with other sensory modalities—such as vision, auditory and tactile systems, accelerometers or anemometers (wind sensors)—on robotic platforms. Such efforts could result in more robust and versatile robotic platforms. In particular, they may be useful for resolving the above-described odour-flow coupling. For instance, the ego motion could be estimated using a vision system or accelerometers, while an anemometer or a flow probe could provide information about the external airflow. This could be used to decorrelate the odour-induced sensor response from environmental effects. Finally, certain applications may directly benefit from multiple sensors. For instance, wildfire monitoring in challenging environments could be performed using olfaction and vision [Wan+23], as the two modalities provide complementary information about a potential fire source, potentially yielding a faster and more robust detection system.

Event-based sensing for MOx electronic nose

Section 4.1 argues that event-driven sensing may benefit machine olfaction, and discusses various ways of generating events from gas sensor data and processing them. Section 4.2 implements a variant of the send-on-delta algorithm to generate events from heater-cycled MOx sensor odour responses. To the best of our knowledge, the described study is the first that addresses this, which indicates a noteworthy gap in the literature. In particular, it is not obvious that send-on-delta is optimal in generating events from heater cycle data features. As the amplitude range can be known beforehand through the feature normalisation procedure used in Section 3.2, amplitude events as described in Section 4.1 may be better suited. This may suggest further studies on the efficacy of different event generation schemes. In particular, a set of different event generation mechanisms could be applied to the laboratory dataset described in Section 3.2, which could then be used for odour classification using

either standard machine learning algorithms (such as support vector machines), or spiking neural networks. The mechanisms could be compared not only with respect to the achieved performance metrics but also considering their compression factor or sampling density, e.g., how many events are required to achieve a desired accuracy.

Similar efforts could be undertaken for MOx sensors operated at a constant temperature. While the literature suggests a wide range of algorithms for asynchronously sampling a continuous odour sensor signal, only a subset may lend themselves to fast sensing. It is not obvious that the ideal sampling algorithm for the constant temperature mode shall be ideal for the cycled heater mode. Complex and multi-pathway event-based olfactory systems may be the outcome of such investigations.

Finally, once a suitable pipeline for event generation and processing has been developed and tested in simulation, further efforts could be undertaken to implement such in hardware. I've contributed to a recent study led by my colleague S. Rastogi, which has investigated an event-based front-end for MOx gas sensors operated at a constant temperature, suggesting a spike-time representation of the gas concentration [Ras+23]. A combination with an analogous hardware front-end for generating events from the heater-cycle MOx sensors, as well as with a spiking neural network that combines the two pathways, could result in a fast and robust neuromorphic olfactory system for the classification and quantisation of different odours.

Event-based sensing beyond MOx electronic nose

This thesis has focused on fast odour sensing with electronic nose technology that leverages MOx gas sensors. While for machine olfaction, these sensors are the obvious candidate due to their temporal characteristics, maturity, and availability, they do have limitations (see Section 1.3 and Section 2.1) and are not the exclusive choice. The following describes the applicability of event-based sensing to two other olfactory sensor modalities, which were explored in the scope of this thesis yet have not been described earlier.

Photonic olfactory sensors

While a range of optical principles have been suggested for machine olfaction (see Fig. 4.1), only recent suggestions promise the scalable integration of sensor arrays in smaller-scale electronic nose technology. A particularly noteworthy candidate emerges from a sequence of studies [Bre+18; Bre+20; Wee+20; Mah+20; Mah+21; Lap+22; Her+22], and has since been commercialised by the company *Aryballe*[1]. The proposed approach [Lap+22] follows the general principle of a Mach-Zehnder interferometer (MZI): A monochromatic light beam is split into two beams. One beam passes through a phase-shift inducing medium, and the other remains at its original phase. After recombination, the interference pattern provides information about the phase shift. For sensing odours, the phase-shift inducing medium is bio-functionalised using organic molecules and short peptides, which react, and change their refractive index, when exposed to particular odours. For increasing sensitivity and reducing signal ambiguity, the authors suggested using a hybrid 120° coherent MZI design, in which the two beams are recombined into three phase-complementary optical outputs instead of

[1] https://aryballe.com/

a single output. See Fig. 5.1a&b for an illustration thereof. The resulting electronic nose combines 64 of such bio-functionalised Mach-Zehnder interferometers (bfMZI) on a silicon nitride platform (see Fig. 5.1c). In their study, the authors report notably high classification performances using seven different VOCs, and a response time in the sub-second range.

The scalability of this technology marks one clear advantage compared to MOx sensors: As the optical outputs are sensed via CMOS camera arrays, increasing the number of single sensor instances may not result in a linear increase in power consumption. This—together with the rapid sensor response time and the arguments provided in Section 4.1—suggests that photonic olfactory sensors may be a suitable candidate for event-based olfactory sensing.

Preliminary results, based on simulating the bfMZI response to an odour pulse, suggest the applicability of event-based sensing. The simulation (see Fig. 5.1d) is based on the information provided in the available literature [Hal+13; Lap+22; Her+22], and approximates the response of the three optical output pathways (p_1-p_3) when encountering an odour-induced phase shift θ. The phase shift may be recovered, up to 2π, as follows:

$$s = \frac{1}{\sqrt{2}} e^{-i\theta} = p_2 - 0.5p_1 - 0.5p_3 + i\frac{\sqrt{3}}{2}(p_1 - p_3) \tag{5.1}$$

$$\theta_{\text{recovered}} = \Re\{i \log(2is) + 2\pi n\} \tag{5.2}$$

From p_1-p_3, asynchronous events were generated using a variant of the send-on-delta sampling algorithm [Mis06a]. Subsequently, the phase shift was recovered from both the regularly sampled signal (here at 20 Hz) and asynchronously sampled events, where the selected threshold ensures a similar number of samples in the considered example. The resulting root-mean-squared-errors (RMSE) between the theoretical phase shift and the recovered signals

were compared for regularly and asynchronously sampled signals. The comparison indicates that, given the provided example, asynchronous samples provide an error that is smaller by a factor of five or more. Unsurprisingly, the difference is significant whenever the signal changes abruptly, with little to no improvements during idle periods, i.e., when fewer or no events are generated. However, confirming whether this apparent advantage of event-based sampling is applicable to real data will require further experiments and analysis, and may be subject to future studies. Further, it is important to note that only a small subset of odours are fluorescent, hence such experiments would be limited in their applicability.

Planar-laser-induced-fluorescence imaging

The introduced and discussed sensor modalities — namely MOx, Optical, EC, CPC and CNT gas sensors, as well as PID — provide means for local, point-source measurements. Other techniques allow for larger-scale odour monitoring. Examples of such are Light Detection and Ranging (LiDAR), Differential Optical Absorption Spectroscopy (DOAS), Schlieren imaging, and Planar Laser-Induced Fluorescence (PLIF). In a collaboration with Dr. Aaron True (AT) and Prof. John Crimaldi (JC) from the University of Boulder Colorado (UCB), we initiated the discussion regarding the applicability of event-based sensing for PLIF.

PLIF is a laser-based visualisation technique that images the distribution of a fluorescent gas or liquid within a plane, thus allowing for non-intrusive measurements of scalar concentrations in fluid flows [Cri08]. The fluorescence illumination $I_{\text{fluorescence}}$ is approximately proportional to the laser power P_{laser} and the concentration of the fluorescent medium C_{medium}:

$$I_{\text{fluorescence}}(t, \mathbf{x}) \propto P_{\text{laser}}(t) \cdot C_{\text{medium}}(t, \mathbf{x}) \tag{5.3}$$

Since the fluorescent effect is induced quasi-instantaneously, the temporal resolution of the technique is merely limited by the camera specification. The experimental setup at UCB is designed for studying slowly moving (low Reynolds number) scalar fluids. However, as discussed in Section 4.1, source-relevant information may be embedded at any spatial or temporal scale, which suggests high sampling rates even for low flow rates. Further, the covered regime is typical for what many aerial and terrestrial animals perform odour-guided navigation tasks in. Recent studies in neuroscience [AH13; Cha+18a] suggest that a natural mechanism for animals to estimate odour concentration ratios is the difference of logarithms:

$$\Delta \log C = \log C_2 - \log C_1 = \log \frac{C_2}{C_1}. \tag{5.4}$$

If combined with Eq. (5.3), this suggests the difference of illumination logarithms as a natural quantity of interest when monitoring fluids in the regime of low Reynolds numbers. Coincidentally, this difference of illumination logarithms is measured by the dynamic vision sensor (DVS), i.e., the most common circuit for event cameras. In such, an event is generated whenever the difference of pixel illumination crosses a fixed threshold. Together, this motivated us to perform a series of experiments combining an event camera and a conventional camera in a standard PLIF setting.

Fig. 5.2a provides an illustration of the experimental setup. A continuous wave (CW) laser is expanded and shaped to a laser sheet of sub-millimetre height. The sheet illuminates a plane in a fluid cell. A vortex field is induced via a magnetic stirrer, and the resulting velocity flow field quantified using particle velocimetry (shown in Fig. 5.2b). A standard frame-based camera (FBC) and an event-based camera (EBC) are equipped with band-pass filters, spatially aligned—using a calibration grid—on the laser plane, and temporally aligned

by tracking the camera shutter triggers. Various experiments have been performed, aiming to characterise the event-camera in a PLIF setting, and extracting insights for future measurement campaigns. As the analysis results are preliminary, they are not further described here.

5.3 Conclusion

During the scope of this thesis, I demonstrated advancements in the field of artificial olfaction, particularly for high temporal resolution odour sampling and processing. By reviewing both biological and artificial olfactory systems, crucial principles for enabling fast and efficient olfaction were identified. The exploration of turbulent odour plume statistics revealed the necessity for rapid sensing mechanisms, which informed the development of electronic nose technologies.

Initial works highlighted critical limitations in existing gas sensor datasets, prompting a set of best practices for data collection and evaluation. Further, limitations in a prominent odour learning algorithm were found and investigated.

These foundations enabled the implementation and evaluation of a custom electronic nose system that leverages rapid temperature cycling of MOx sensors, achieving unprecedented levels of temporal resolution and accuracy in odour classification tasks. In a set of experiments, I have shown that the system can distinguish between odour pulses with modulation frequencies up to 60 Hz and classify odours at pulse widths as short as 10 milliseconds. These findings match or surpass the temporal capabilities demonstrated in mammalian olfaction and will open new avenues for practical applications in various fields, including environmental monitoring and robotic olfaction.

Furthermore, our investigation into neuromorphic computing principles for processing olfactory signals emphasised the advantages of asynchronous sampling and event-driven processing. A demonstration followed for asynchronously sampling heater-cycled MOx sensor data, which allowed for the comparison of different event encoding schemes.

In summary, this thesis has addressed the challenges of high-speed odour detection and classification and further provides a comprehensive framework for future research and development in the field. The methodologies and findings presented here may offer a foundation for advancing the capabilities of artificial olfaction systems, with promising implications for real-world applications and further scientific explorations.

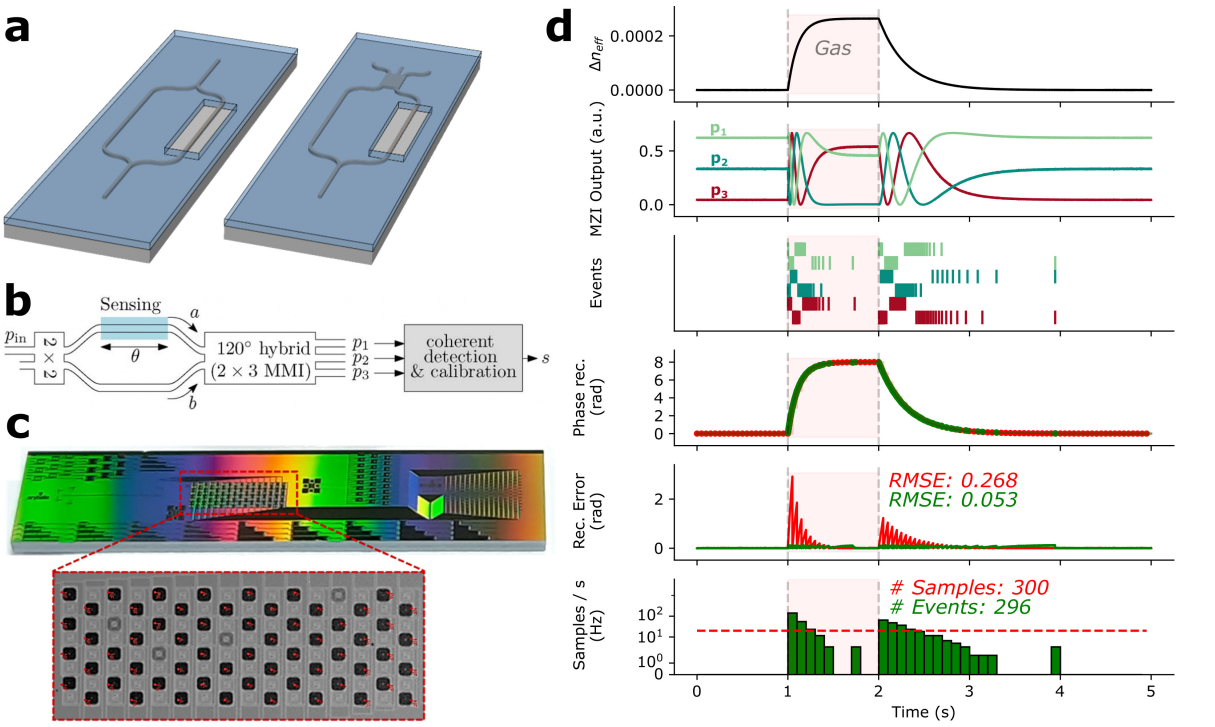

Figure 5.1: a) Bio-functionalised Mach-Zehnder interferometer (MZI). left: standard MZI, right: hybrid 120° coherent MZI. **b)** Measurement principle of hybrid 120° coherent MZI design. **c)** Picture of the *Aryballe* photonic electronic nose die (top), zooming into the sensing arm of the MZI array (bottom). **d)** Simulation study of event-based MZI odour sensor. From top to bottom: odour-induced change in refractive index; resulting outputs of coherent MZI; channel-wise asynchronous events generated via send-on-delta sampling; reconstructed phase-shift from regularly sampled signal (red) and events (green); reconstruction error between actual and reconstructed phase-shift, for both sampling schemes; samples per time, for both sampling schemes. **a)** and **c)** adapted with permission from [Lap+22] (© Optical Society of America), and **b)** adapted with permission from [Hal+13] (© 2013 IEEE. Reprinted, with permission, from R. Halir, L. Vivien, X. Le Roux, D. X. Xu, and P. Cheben. "Direct and sensitive phase readout for integrated waveguide sensors.", IEEE Photonics Journal, 2013).

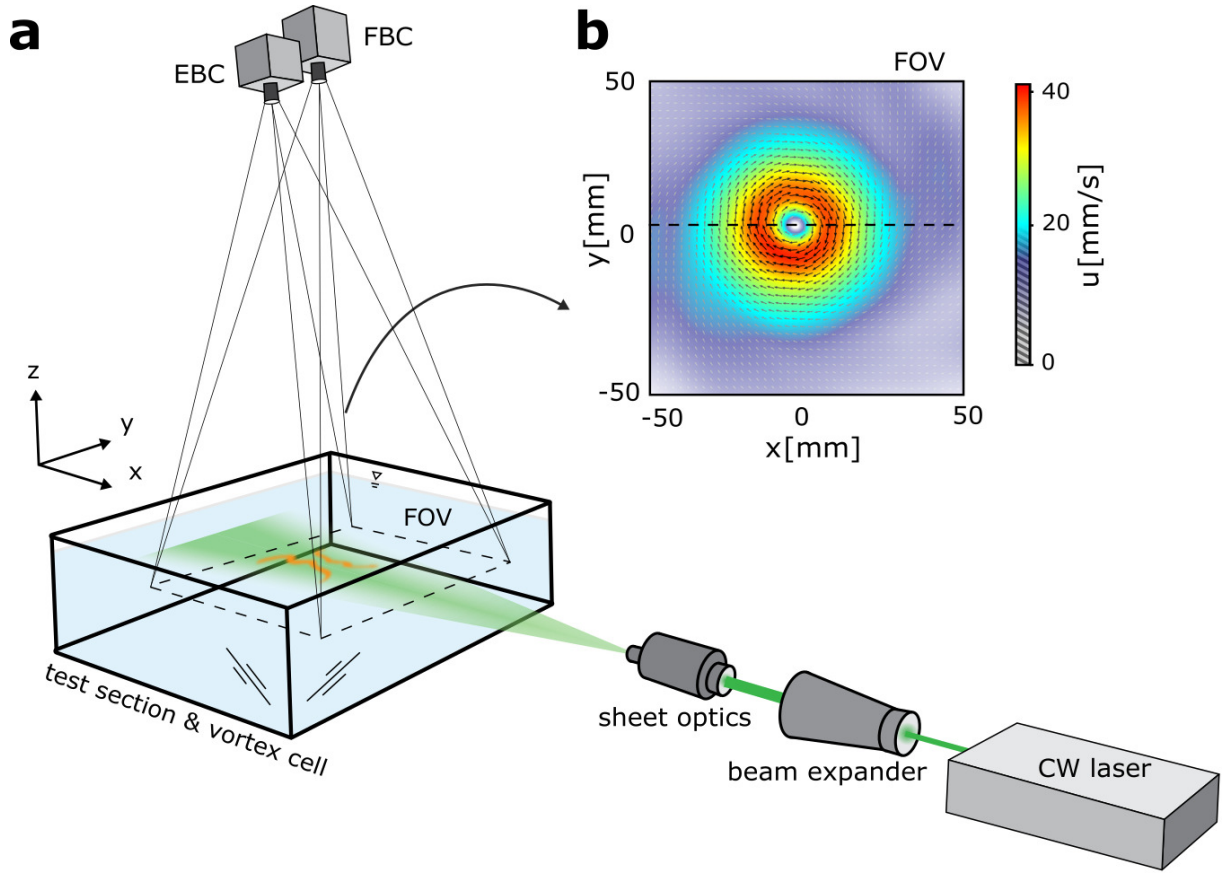

Figure 5.2: a) Experimental setup for simultaneous recording of turbulent flow using frame-based (FBC) and event-based (EBC) vision. **b)** Vortex flow characterisation using particle velocimetry.

Bibliography

[24a] *auroraScientific 200C: miniPID Fast Response Olfaction Sensor*. 2024.

[24b] *Introduction to Photoionization*. 2024.

[24c] *Metal Oxide gas sensor elements*. 2024.

[Abd+23] Gisya Abdi, Lulu Alluhaibi, Ewelina Kowalewska, Tomasz Mazur, Krzysztof Mech, Agnieszka Podborska, Andrzej Sławek, Hirofumi Tanaka, and Konrad Szaciłowski. "Reservoir computing and photoelectrochemical sensors: A marriage of convenience." en. In: *Coordination chemistry reviews* 487 (2023), p. 215155. DOI: 10.1016/j.ccr.2023.215155.

[Abe04] Moshe Abeles. "Time Is Precious." In: *Science* 304.5670 (2004), pp. 523–524. DOI: 10.1126/science.1097725.

[Abr+04] N Abraham, H Spors, A Carleton, T Margrie, T Kuner, and A Schaefer. "Maintaining Accuracy at the Expense of SpeedStimulus Similarity Defines Odor Discrimination Time in Mice." en. In: *Neuron* 44.5 (2004), pp. 865–876. DOI: 10.1016/S0896-6273(04)00753-6.

[Ack+21] Tobias Ackels, Andrew Erskine, Debanjan Dasgupta, Alina Cristina Marin, Tom P.A. Warner, Sina Tootoonian, Izumi Fukunaga, Julia J. Harris, and Andreas T. Schaefer. "Fast odour dynamics are encoded in the olfactory system and guide behaviour." In: *Nature* 593.7860 (2021). Publisher: Springer US, pp. 558–563. DOI: 10.1038/s41586-021-03514-2.

[AGS19] Ismael Araujo, Juan Gamboa, and Adenilton Silva. "Deep learning models for classification of gases detected by sensor arrays of artificial nose." In: *Anais do XVI Encontro Nacional de Inteligência Artificial e Computacional*. ISSN: 0000-0000 event-place: Salvador. Porto Alegre, RS, Brasil: SBC, 2019, pp. 844–855. DOI: 10.5753/eniac.2019.9339.

[AH13] Hannah A. Arnson and Timothy E. Holy. "Robust Encoding of Stimulus Identity and Concentration in the Accessory Olfactory System." en. In: *The journal of neuroscience* 33.33 (2013), pp. 13388–13397. DOI: 10.1523/JNEUROSCI.0967-13.2013.

[AJT20] Naader Alizadeh, Hoda Jamalabadi, and Farnaz Tavoli. "Breath Acetone Sensors as Non-Invasive Health Monitoring Systems: A Review." In: *Ieee sensors journal* 20.1 (2020), pp. 5–31. DOI: 10.1109/JSEN.2019.2942693.

[Álv+18] Efrén Álvarez-Salvado, Angela M Licata, Erin G Connor, Margaret K McHugh, Benjamin Mn King, Nicholas Stavropoulos, Jonathan D Victor, John P Crimaldi, and Katherine I Nagel. "Elementary sensory-motor transformations underlying olfactory navigation in walking fruit-flies." en. In: *Elife* 7 (2018), e37815. DOI: 10.7554/eLife.37815.

[Amo+08] L. Amo, I. Galván, G. Tomás, and J. J. Sanz. "Predator odour recognition and avoidance in a songbird." en. In: *Functional ecology* 22.2 (2008), pp. 289–293. DOI: 10.1111/j.1365-2435.2007.01361.x.

[And+11] Martin N. Andersson, Muhammad Binyameen, Medhat M. Sadek, and Fredrik Schlyter. "Attraction Modulated by Spacing of Pheromone Components and Anti-attractants in a Bark Beetle and a Moth." en. In: *Journal of chemical ecology* 37.8 (2011), pp. 899–911. DOI: 10.1007/s10886-011-9995-3.

[And+20] Melanie J Anderson, Joseph G Sullivan, Timothy K Horiuchi, Sawyer B Fuller, and Thomas L Daniel. "A bio-hybrid odor-guided autonomous palm-sized air vehicle." en. In: *Bioinspiration & biomimetics* 16.2 (2020), p. 026002. DOI: 10.1088/1748-3190/abbd81.

[Are+19] Santiago R Arellano et al. "UAV-based measurements of the high-altitude plume of Manam Volcano." In: *AGU Fall Meeting Abstracts*. Vol. 2019. 2019, EP13A–02.

[Asa+17] Sahar Asadi, Han Fan, Victor Hernandez Bennetts, and Achim J. Lilienthal. "Time-dependent gas distribution modelling." In: *Robotics and autonomous systems* 96 (2017), pp. 157–170. DOI: https://doi.org/10.1016/j.robot.2017.05.012.

[BA91] Linda Buck and Richard Axel. "A novel multigene family may encode odorant receptors: a molecular basis for odor recognition." In: *Cell* 65.1 (1991). Publisher: Cell Press, pp. 175–187.

[Bar+16] Chiara Bartolozzi, Lorenzo Natale, Francesco Nori, and Giorgio Metta. "Robots with a sense of touch." In: *Nature materials* 15.9 (2016), pp. 921–925.

[Bar61] H. B. Barlow. "Possible Principles Underlying the Transformations of Sensory Messages." en. In: *Sensory Communication*. The MIT Press, 1961, pp. 216–234. DOI: 10.7551/mitpress/9780262518420.003.0013.

[Bau+18] Tobias Baur, Caroline Schultealbert, Andreas Schütze, and Tilman Sauerwald. "Novel method for the detection of short trace gas pulses with metal oxide semiconductor gas sensors." en. In: *Journal of sensors and sensor systems* 7.1 (2018), pp. 411–419. DOI: 10.5194/jsss-7-411-2018.

[BBK18] Casey Battaglino, Grey Ballard, and Tamara G. Kolda. "A Practical Randomized CP Tensor Decomposition." In: *Siam journal on matrix analysis and applications* 39.2 (2018). Publisher: Society for Industrial & Applied Mathematics (SIAM), pp. 876–901. DOI: 10.1137/17m1112303.

[BC19] Ayon Borthakur and Thomas A. Cleland. "A Spike Time-Dependent Online Learning Algorithm Derived From Biological Olfaction." en. In: *Frontiers in neuroscience* 13 (2019), p. 656. DOI: 10.3389/fnins.2019.00656.

[BCR00] R. E. Baby, M. Cabezas, and E. N. Walsöe de Reca. "Electronic nose: a useful tool for monitoring environmental contamination." In: *Sensors and actuators b: chemical* 69.3 (2000), pp. 214–218. DOI: https://doi.org/10.1016/S0925-4005(00)00491-3.

[BD14] Floris van Breugel and Michael H. Dickinson. "Plume-Tracking Behavior of Flying Drosophila Emerges from a Set of Distinct Sensory-Motor Reflexes." en. In: *Current biology* 24.3 (2014), pp. 274–286. DOI: 10.1016/j.cub.2013.12.023.

[Ben+12] Victor Hernandez Bennetts, Achim J Lilienthal, Ali Abdul Khaliq, Vıctor Pomareda Sese, and Marco Trincavelli. "Gasbot: A mobile robotic platform for methane leak detection and emission monitoring." en. In: (2012).

[Bey+10] Michael Beyeler, Fabio Stefanini, Henning Proske, Giovanni Galizia, and Elisabetta Chicca. "Exploring olfactory sensory networks: Simulations and hardware emulation." en. In: *2010 Biomedical Circuits and Systems Conference (BioCAS)*. Paphos, Cyprus: IEEE, 2010, pp. 270–273. DOI: 10.1109/BIOCAS.2010.5709623.

[BFC98] T. C. Baker, H. Y. Fadamiro, and A. A. Cosse. "Moth uses fine tuning for odour resolution." en. In: *Nature* 393.6685 (1998), pp. 530–530. DOI: 10.1038/31131.

[BH82] Alexander Borst and Martin Heisenberg. "Osmotropotaxis inDrosophila melanogaster." en. In: *Journal of comparative physiology ? a* 147.4 (1982), pp. 479–484. DOI: 10.1007/BF00612013.

[Bha+10] Vikas Bhandawat, Gaby Maimon, Michael H. Dickinson, and Rachel I. Wilson. "Olfactory modulation of flight in *Drosophila* is sensitive, selective and rapid." en. In: *Journal of experimental biology* 213.24 (2010), pp. 4313–4313. DOI: 10.1242/jeb.053546.

[BJM18] Javier Burgués, Juan Manuel Jiménez-Soto, and Santiago Marco. "Estimation of the limit of detection in semiconductor gas sensors through linearized calibration models." en. In: *Analytica chimica acta* 1013 (2018), pp. 13–25. DOI: 10.1016/j.aca.2018.01.062.

[BJS05] Stacey L Brown, Joby Joseph, and Mark Stopfer. "Encoding a temporally structured stimulus with a temporally structured neural representation." en. In: *Nature neuroscience* 8.11 (2005), pp. 1568–1576. DOI: 10.1038/nn1559.

[BK04] Peter A. Brennan and Eric B. Keverne. "Something in the Air? New Insights into Mammalian Pheromones." en. In: *Current biology* 14.2 (2004), R81–R89. DOI: 10.1016/j.cub.2003.12.052.

[Bla] Andrew R Blaustein. "Sexual Selection and Mammalian Olfaction." en. In: *The americannaturalist* ().

[BM19] Javier Burgues and Santiago Marco. "Wind-Independent Estimation of Gas Source Distance From Transient Features of Metal Oxide Sensor Signals." en. In: *Ieee access* 7 (2019), pp. 140460–140469. DOI: 10.1109/ACCESS.2019.2940936.

[BM20a] Javier Burgués and Santiago Marco. "Environmental chemical sensing using small drones: A review." en. In: *Science of the total environment* 748 (2020), p. 141172. DOI: 10.1016/j.scitotenv.2020.141172.

[BM20b] Javier Burgués and Santiago Marco. "Feature Extraction for Transient Chemical Sensor Signals in Response to Turbulent Plumes: Application to Chemical Source Distance Prediction." en. In: *Sensors and actuators b: chemical* 320 (2020), p. 128235. DOI: 10.1016/j.snb.2020.128235.

[Boa00] K.A. Boahen. "Point-to-point connectivity between neuromorphic chips using address events." en. In: *Ieee transactions on circuits and systems ii: analog and digital signal processing* 47.5 (2000), pp. 416–434. DOI: 10.1109/82.842110.

[Bre+18] Sophie Brenet, Aurelian John-Herpin, François-Xavier Gallat, Benjamin Musnier, Arnaud Buhot, Cyril Herrier, Tristan Rousselle, Thierry Livache, and Yanxia Hou. "Highly-Selective Optoelectronic Nose Based on Surface Plasmon Resonance Imaging for Sensing Volatile Organic Compounds." en. In: *Analytical chemistry* 90.16 (2018), pp. 9879–9887. DOI: 10.1021/acs.analchem.8b02036.

[Bre+20] Sophie Brenet, Jonathan S. Weerakkody, Arnaud Buhot, François-Xavier Gallat, Raphael Mathey, Loïc Leroy, Thierry Livache, Cyril Herrier, and Yanxia Hou. "Improvement of sensitivity of surface plasmon resonance imaging for the gas-phase detection of volatile organic compounds." en. In: *Talanta* 212 (2020), p. 120777. DOI: 10.1016/j.talanta.2020.120777.

[Bur+19] Javier Burgués, Victor Hernández, Achim Lilienthal, and Santiago Marco. "Smelling Nano Aerial Vehicle for Gas Source Localization and Mapping." en. In: *Sensors* 19.3 (2019), p. 478. DOI: 10.3390/s19030478.

[Bur+21a] Javier Burgués, María Deseada Esclapez, Silvia Doñate, and Santiago Marco. "RHINOS: A lightweight portable electronic nose for real-time odor quantification in wastewater treatment plants." en. In: *Iscience* 24.12 (2021), p. 103371. DOI: 10.1016/j.isci.2021.103371.

[Bur+21b] Javier Burgués, María Deseada Esclapez, Silvia Doñate, Laura Pastor, and Santiago Marco. "Aerial Mapping of Odorous Gases in a Wastewater Treatment Plant Using a Small Drone." en. In: *Remote sensing* 13.9 (2021), p. 1757. DOI: 10.3390/rs13091757.

[Bur+22] Shawn D Burton, Audrey Brown, Thomas P Eiting, Isaac A Youngstrom, Thomas C Rust, Michael Schmuker, and Matt Wachowiak. "Mapping odorant sensitivities reveals a sparse but structured representation of olfactory chemical space by sensory input to the mouse olfactory bulb." en. In: *Elife* 11 (2022), e80470. DOI: 10.7554/eLife.80470.

[Bur19] Javier Burgues. "Signal Processing and Machine Learning for Gas Sensors: Gas Source Localization with a Nano-Drone." PhD Thesis. Universitat de Barcelona, 2019.

[BVM19] Javier Burgues, Luis F. Valdez, and Santiago Marco. "High-bandwidth e-nose for rapid tracking of turbulent plumes." en. In: *2019 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*. Fukuoka, Japan: IEEE, 2019, pp. 1–3. DOI: 10.1109/ISOEN.2019.8823158.

[BW01] Nicolae Barsan and Udo Weimar. "Conduction Model of Metal Oxide Gas Sensors." en. In: *Journal of electroceramics* 7 (2001), pp. 143–167.

[Car21] Ring T. Cardé. "Navigation Along Windborne Plumes of Pheromone and Resource-Linked Odors." en. In: *Annual review of entomology* 66.1 (2021), pp. 317–336. DOI: 10.1146/annurev-ento-011019-024932.

[Cas+20] Massimo Cassiani, Matteo B. Bertagni, Massimo Marro, and Pietro Salizzoni. "Concentration Fluctuations from Localized Atmospheric Releases." In: *Boundary-layer meteorology* 177.2-3 (2020), pp. 461–510. DOI: 10.1007/s10546-020-00547-4.

[Cat13] Kenneth C. Catania. "Stereo and serial sniffing guide navigation to an odour source in a mammal." en. In: *Nature communications* 4.1 (2013), p. 1441. DOI: 10.1038/ncomms2444.

[CCG12] F.K. Che Harun, J.A. Covington, and J.W. Gardner. "Mimicking the biological olfactory system: a Portable electronic Mucosa." en. In: *Iet nanobiotechnology* 6.2 (2012), p. 45. DOI: 10.1049/iet-nbt.2010.0032.

[CCP20] Il-Sik Chang, Hee-jo Choi, and Goo-man Park. "UCI Sensor Data Analysis based on Data Visualization." In: *Proceedings of the Korean Society of Broadcast Engineers Conference*. The Korean Institute of Broadcast and Media Engineers, 2020, pp. 21–24.

[Cel+17] Antonio Cellini, Sonia Blasioli, Enrico Biondi, Assunta Bertaccini, Ilaria Braschi, and Francesco Spinelli. "Potential Applications and Limitations of Electronic Nose Devices for Plant Disease Diagnosis." en. In: *Sensors* 17.11 (2017), p. 2596. DOI: 10.3390/s17112596.

[Cha+18a] Ho Ka Chan, Fabian Hersperger, Emiliano Marachlian, Brian H. Smith, Fernando Locatelli, Paul Szyszka, and Thomas Nowotny. "Odorant mixtures elicit less variable and faster responses than pure odorants." en. In: *Plos computational biology* 14.12 (2018). Ed. by Daniel Bush, e1006536. DOI: 10.1371/journal.pcbi.1006536.

[Cha+18b] Ji-Eun Chang, Dae-Sik Lee, Sang-Woo Ban, Jaeho Oh, Moon Youn Jung, Seung-Hwan Kim, SungJoon Park, Krishna Persaud, and Sanghoon Jheon. "Analysis of volatile organic compounds in exhaled breath for lung cancer diagnosis using a sensor system." en. In: *Sensors and actuators b: chemical* 255 (2018), pp. 800–807. DOI: 10.1016/j.snb.2017.08.057.

[Cha+97] S Chandiok, B A Crawley, B A Oppenheim, P R Chadwick, S Higgins, and K C Persaud. "Screening for bacterial vaginosis: a novel application of artificial nose technology." en. In: *Journal of clinical pathology* 50.9 (1997), pp. 790–791. DOI: 10.1136/jcp.50.9.790.

[Che+07] F.K. Che Harun, P.H. King, J.A. Covington, and J.W. Gardner. "Novel gas chromatographic microsystem with very large sensor arrays for advanced odour discrimination." en. In: *2007 IEEE Sensors*. Atlanta, GA, USA: IEEE, 2007, pp. 1361–1363. DOI: 10.1109/ICSENS.2007.4388664.

[Che+09] F.K. Che Harun, J.E. Taylor, J.A. Covington, and J.W. Gardner. "An electronic nose employing dual-channel odour separation columns with large chemosensor arrays for advanced odour discrimination." en. In: *Sensors and actuators b: chemical* 141.1 (2009), pp. 134–140. DOI: 10.1016/j.snb.2009.05.036.

[Che+11] Hung Tat Chen, Kwan Ting Ng, Amine Bermak, Man Kay Law, and Dominique Martinez. "Spike Latency Coding in Biologically Inspired Microelectronic Nose." en. In: *Ieee transactions on biomedical circuits and systems* 5.2 (2011), pp. 160–168. DOI: 10.1109/TBCAS.2010.2075928.

[Che+21] Lu Cheng, Qing-Hao Meng, Achim J Lilienthal, and Pei-Feng Qi. "Development of compact electronic noses: a review." en. In: *Measurement science and technology* 32.6 (2021), p. 062002. DOI: 10.1088/1361-6501/abef3b.

[Chr+22] Dennis V Christensen et al. "2022 roadmap on neuromorphic computing and engineering." en. In: *Neuromorphic computing and engineering* 2.2 (2022), p. 022501. DOI: 10.1088/2634-4386/ac4a83.

[CL19] Jun-Ho Choi and Jong-Seok Lee. "EmbraceNet: A robust deep learning architecture for multimodal classification." In: *Information fusion* 51 (2019). Publisher: Elsevier BV, pp. 259–270. DOI: 10.1016/j.inffus.2019.02.010.

[CLG19] Angelo Cardellicchio, Angela Lombardi, and Cataldo Guaragnella. "Iterative complex network approach for chemical gas sensor array characterisation." In: *The journal of engineering* 2019.6 (2019). Publisher: Institution of Engineering and Technology (IET), pp. 4612–4616. DOI: 10.1049/joe.2018.5125.

[CMC18] Erin G. Connor, Margaret K. McHugh, and John P. Crimaldi. "Quantification of airborne odor plumes using planar laser-induced fluorescence." en. In: *Experiments in fluids* 59.9 (2018), p. 137. DOI: 10.1007/s00348-018-2591-3.

[Cov+21] James A. Covington, Santiago Marco, Krishna C. Persaud, Susan S. Schiffman, and H. Troy Nagle. "Artificial Olfaction in the 21 st Century." en. In: *Ieee sensors journal* 21.11 (2021), pp. 12969–12990. DOI: 10.1109/JSEN.2021.3076412.

[CR16] Jonathan T. Clark and Anandasankar Ray. "Olfactory Mechanisms for Discovery of Odorants to Reduce Insect-Host Contact." en. In: *Journal of chemical ecology* 42.9 (2016), pp. 919–930. DOI: 10.1007/s10886-016-0770-3.

[Cri+22] John Crimaldi, Hong Lei, Andreas Schaefer, Michael Schmuker, Brian H. Smith, Aaron C. True, Justus V. Verhagen, and Jonathan D. Victor. "Active sensing in a dynamic olfactory world." en. In: *Journal of computational neuroscience* 50.1 (2022), pp. 1–6. DOI: 10.1007/s10827-021-00798-1.

[Cri08] J. P. Crimaldi. "Planar laser induced fluorescence in aqueous flows." In: *Experiments in fluids* 44.6 (2008), pp. 851–863. DOI: 10.1007/s00348-008-0496-2.

[CSN14] Elisabetta Chicca, Michael Schmuker, and Martin Nawrot. "Neuromorphic Sensors, Olfaction." en. In: *Encyclopedia of Computational Neuroscience*. New York, NY: Springer New York, 2014, pp. 1–7. DOI: 10.1007/978-1-4614-7320-6_119-2.

[CT65] James W Cooley and John W Tukey. "An algorithm for the machine calculation of complex Fourier series." In: *Mathematics of computation* 19.90 (1965), pp. 297–301.

[CU10] Kevin M. Cury and Naoshige Uchida. "Robust Odor Coding via Inhalation-Coupled Transient Activity in the Mammalian Olfactory Bulb." en. In: *Neuron* 68.3 (2010), pp. 570–585. DOI: 10.1016/j.neuron.2010.09.040.

[CV95] Corinna Cortes and Vladimir Vapnik. "Support-vector networks." In: *Machine learning* 20.3 (1995). Publisher: Springer, pp. 273–297.

[CVV14] Antonio Celani, Emmanuel Villermaux, and Massimo Vergassola. "Odor Landscapes in Turbulent Environments." en. In: *Physical review x* 4.4 (2014), p. 041015. DOI: 10.1103/PhysRevX.4.041015.

[CW08] Ring T. Cardé and Mark A. Willis. "Navigational Strategies Used by Insects to Find Distant, Wind-Borne Sources of Odor." en. In: *Journal of chemical ecology* 34.7 (2008), pp. 854–866. DOI: 10.1007/s10886-008-9484-5.

[Dai+23] Xinyu Dai, Dexuan Huo, Zhanyuan Gao, Jilin Zhang, and Hong Chen. *A visual-olfactory multisensory fusion spike neural network for early fire/smoke detection*. en. 2023. DOI: 10.21203/rs.3.rs-3192562/v1.

[DAm+10] Arnaldo D'Amico, Giorgio Pennazza, Marco Santonico, Eugenio Martinelli, Claudio Roscioni, Giovanni Galluccio, Roberto Paolesse, and Corrado Di Natale. "An investigation on electronic nose diagnosis of lung cancer." en. In: *Lung cancer* 68.2 (2010), pp. 170–176. DOI: 10.1016/j.lungcan.2009.11.003.

[Das+22] Debanjan Dasgupta, Tom P. A. Warner, Andrew Erskine, and Andreas T. Schaefer. "Coupling of Mouse Olfactory Bulb Projection Neurons to Fluctuating Odor Pulses." en. In: *The journal of neuroscience* 42.21 (2022), pp. 4278–4296. DOI: 10.1523/JNEUROSCI.1422-21.2022.

[Dav+18] Mike Davies et al. "Loihi: A Neuromorphic Manycore Processor with On-Chip Learning." In: *Ieee micro* 38.1 (2018), pp. 82–99. DOI: 10.1109/mm.2018.112130359.

[Dav+21] Mike Davies, Andreas Wild, Garrick Orchard, Yulia Sandamirskaya, Gabriel A. Fonseca Guerra, Prasad Joshi, Philipp Plank, and Sumedh R. Risbud. "Advancing Neuromorphic Computing With Loihi: A Survey of Results and Outlook." In: *Proceedings of the ieee* 109.5 (2021), pp. 911–934. DOI: 10.1109/JPROC.2021.3067593.

[Dav21] Mike Davies. *Nice 2021 keynote: lessons from loihi for the future of neuromorphic computing*. https://www.youtube.com/watch?v=-dl1FPrpw1A&t=1463s. Accessed: 2022-11-22. 2021.

[DC11] Teun Dekker and Ring T. Cardé. "Moment-to-moment flight manoeuvres of the female yellow fever mosquito (*Aedes aegypti* L.) in response to plumes of carbon dioxide and human skin odour." en. In: *Journal of experimental biology* 214.20 (2011), pp. 3480–3494. DOI: 10.1242/jeb.055186.

[DCD99] P. Duchamp-Viret, M. A. Chaput, and A. Duchamp. "Odor Response Properties of Rat Olfactory Receptor Neurons." en. In: *Science* 284.5423 (1999), pp. 2171–2174. DOI: 10.1126/science.284.5423.2171.

[DDS22] Damien Drix†, Nik Dennler†, and Michael Schmuker († denotes equal contribution). "Rapid Recognition of Olfactory Scenes with a Portable MOx Sensor System using Hotplate Modulation." en. In: *2022 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*. Aveiro, Portugal: IEEE, 2022, pp. 1–4. DOI: 10.1109/ISOEN54820.2022.9789654.

[Del+10] Tobi Delbruck, Bernabe Linares-Barranco, Eugenio Culurciello, and Christoph Posch. "Activity-driven, event-based vision sensors." en. In: *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. Paris, France: IEEE, 2010, pp. 2426–2429. DOI: 10.1109/ISCAS.2010.5537149.

[Dem+20] Mahmut Demir, Nirag Kadakia, Hope D Anderson, Damon A Clark, and Thierry Emonet. "Walking Drosophila navigate complex plumes using stochastic decisions biased by the timing of odor encounters." en. In: *Elife* 9 (2020), e57524. DOI: 10.7554/eLife.57524.

[Den+21a] Nik Dennler, Germain Haessig, Matteo Cartiglia, and Giacomo Indiveri. "Online Detection of Vibration Anomalies Using Balanced Spiking Neural Networks." In: *2021 ieee 3rd international conference on artificial intelligence circuits and systems, aicas 2021* (2021). arXiv: 2106.00687 ISBN: 9781665419130. DOI: 10.1109/AICAS51828.2021.9458403.

[Den+21b] Nik Dennler, Shavika Rastogi, Jordi Fonollosa, André van Schaik, and Michael Schmuker. "Drift in a popular metal oxide sensor dataset reveals limitations for gas classification benchmarks." In: *Arxiv preprint arxiv:2108.08793* (2021).

[Den+22a] Nik Dennler†, Damien Drix†, Shavika Rastogi, André van Schaik, and Michael Schmuker († denotes equal contribution). "Rapid Inference of Geographical Location with an Event-based Electronic Nose." In: *The 9th Annual Neuro-Inspired Computational Elements (NICE), 2022, UTSA, USA Issue: 1*. Vol. 1. Association for Computing Machinery, 2022.

[Den+22b] Nik Dennler, Shavika Rastogi, Jordi Fonollosa, André van Schaik, and Michael Schmuker. "Drift in a popular metal oxide sensor dataset reveals limitations for gas classification benchmarks." In: *Sensors and Actuators B: Chemical* 361 (2022), p. 131668. DOI: 10.1016/j.snb.2022.131668.

[Den+24] Nik Dennler, Damien Drix, Tom PA Warner, Shavika Rastogi, Cecilia Della Casa, Tobias Ackels, Andreas T Schaefer, André van Schaik, and Michael Schmuker. "High-speed odor sensing using miniaturized electronic nose." In: *Science Advances* 10.45 (2024), eadp1764. DOI: 10.1126/sciadv.adp1764.

[Dew+18] Adam Dewan, Annika Cichy, Jingji Zhang, Kayla Miguel, Paul Feinstein, Dmitry Rinberg, and Thomas Bozza. "Single olfactory receptors set odor detection thresholds." en. In: *Nature communications* 9.1 (2018), p. 2887. DOI: 10.1038/s41467-018-05129-0.

[Dey18] Ananya Dey. "Semiconductor metal oxide gas sensors: A review." en. In: *Materials science and engineering: b* 229 (2018), pp. 206–217. DOI: 10.1016/j.mseb.2017.12.036.

[Di +14] Enrico Di Lello, Marco Trincavelli, Herman Bruyninckx, and Tinne De Laet. "Augmented Switching Linear Dynamical System Model for Gas Concentration Estimation with MOX Sensors in an Open Sampling System." en. In: *Sensors* 14.7 (2014), pp. 12533–12559. DOI: 10.3390/s140712533.

[Di +21] Davide Di Giuseppe, Alexandro Catini, Elisabetta Comini, Dario Zappa, Corrado Di Natale, and Eugenio Martinelli. "Optimizing MOX sensor array performances with a reconfigurable self-adaptive temperature modulation interface." en. In: *Sensors and actuators b: chemical* 333 (2021), p. 129509. DOI: 10.1016/j.snb.2021.129509.

[Di +95] Corrado Di Natale, Santiago Marco, Fabrizio Davide, and Arnaldo D'Amico. "Sensor-array calibration time reduction by dynamic modelling." en. In: *Sensors and actuators b: chemical* 25.1-3 (1995), pp. 578–583. DOI: 10.1016/0925-4005(95)85126-7.

[Dia+16] A. Diamond, M. Schmuker, A. Z. Berna, S. Trowell, and Thomas Nowotny. "Classifying continuous, real-time e-nose sensor data using a bio-inspired spiking network modelled on the insect olfactory system." In: *Bioinspiration and biomimetics* 11.2 (2016). Publisher: IOP Publishing. DOI: 10.1088/1748-3190/11/2/026002.

[Dom+24] Guillem Domènech-Gil, Nguyen Thanh Duc, J. Jacob Wikner, Jens Eriksson, Sören Nilsson Påledal, Donatella Puglisi, and David Bastviken. "Electronic Nose for Improved Environmental Methane Monitoring." en. In: *Environmental science & technology* 58.1 (2024), pp. 352–361. DOI: 10.1021/acs.est.3c06945.

[DS21] Damien Drix and Michael Schmuker. "Resolving fast gas transients with metal oxide sensors." In: *Acs sensors* (2021). DOI: 10.1021/acssensors.0c02006.

[DSS23] Nik Dennler, André van Schaik, and Michael Schmuker. *Limitations in odour recognition and generalisation in a neuromorphic olfactory circuit*. en. arXiv:2309.11555 [cs, eess]. 2023.

[DSS24] Nik Dennler, André van Schaik, and Michael Schmuker. "Limitations in odour recognition and generalization in a neuromorphic olfactory circuit." In: *Nature Machine Intelligence* 6 (2024), pp. 1451–1453. DOI: 10.1038/s42256-024-00952-1.

[Dui+21] Bardienus P. Duisterhof, Shushuai Li, Javier Burgues, Vijay Janapa Reddi, and Guido C. H. E. De Croon. "Sniffy Bug: A Fully Autonomous Swarm of Gas-Seeking Nano Quadcopters in Cluttered Environments." en. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Prague, Czech Republic: IEEE, 2021, pp. 9099–9106. DOI: 10.1109/IROS51168.2021.9636217.

[Dun+14] Andreas Dunkel, Martin Steinhaus, Matthias Kotthoff, Bettina Nowak, Dietmar Krautwurst, Peter Schieberle, and Thomas Hofmann. "Nature's Chemical Signatures in Human Olfaction: A Foodborne Perspective for Future Biotechnology." en. In: *Angewandte chemie international edition* 53.28 (2014), pp. 7124–7143. DOI: 10.1002/anie.201309508.

[Ege+18] Alexander Egea-Weiss, Alpha Renner, Christoph J. Kleineidam, and Paul Szyszka. "High Precision of Spike Timing across Olfactory Receptor Neurons Allows Rapid Odor Coding in Drosophila." en. In: *Iscience* 4 (2018), pp. 76–83. DOI: 10.1016/j.isci.2018.05.009.

[EJM23] Chiara Ercolani, Wanting Jin, and Alcherio Martinoli. "3D Gas Sensing with Multiple Nano Aerial Vehicles: Interference Analysis, Algorithms and Experimental Validation." en. In: *Sensors* 23.20 (2023), p. 8512. DOI: 10.3390/s23208512.

[EM20] Chiara Ercolani and Alcherio Martinoli. "3D Odor Source Localization using a Micro Aerial Vehicle: System Design and Performance Evaluation." en. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Las Vegas, NV, USA: IEEE, 2020, pp. 6194–6200. DOI: 10.1109/IROS45743.2020.9341501.

[Ers+19a] Andrew Erskine, Tobias Ackels, Debanjan Dasgupta, Izumi Fukunaga, and Andreas T. Schaefer. *Mammalian olfaction is a high temporal bandwidth sense*. en. preprint. Neuroscience, 2019. DOI: 10.1101/570689.

[Ers+19b] Andrew Erskine, Thorsten Bus, Jan T Herb, and Andreas T Schaefer. "AutonoMouse: High throughput operant conditioning reveals progressive impairment with graded olfactory bulb lesions." In: *Plos one* 14.3 (2019). Publisher: Public Library of Science San Francisco, CA USA, e0211571.

[Fan+18] Han Fan, Victor Hernandez Bennetts, Erik Schaffernicht, and Achim J. Lilienthal. "A cluster analysis approach based on exploiting density peaks for gas discrimination with electronic noses in open environments." In: *Sensors and actuators b: chemical* 259 (2018). Publisher: Elsevier BV, pp. 183–203. DOI: 10.1016/j.snb.2017.10.063.

[Fan+19] Han Fan, Victor Hernandez Bennetts, Erik Schaffernicht, and Achim Lilienthal. "Towards Gas Discrimination and Mapping in Emergency Response Scenarios Using a Mobile Robot with an Electronic Nose." en. In: *Sensors* 19.3 (2019), p. 685. DOI: 10.3390/s19030685.

[Far+19] Mariana Valente Farraia, João Cavaleiro Rufo, Inês Paciência, Francisca Mendes, Luís Delgado, and André Moreira. "The electronic nose technology in clinical diagnosis: A systematic review." en. In: *Porto biomedical journal* 4.4 (2019), e42. DOI: 10.1097/j.pbj.0000000000000042.

[Fig] Figaro USA, Inc. http://www.figarosensor.com/.

[Fin+10] George F. Fine, Leon M. Cavanagh, Ayo Afonja, and Russell Binions. "Metal Oxide Semi-Conductor Gas Sensors in Environmental Monitoring." en. In: *Sensors* 10.6 (2010), pp. 5469–5502. DOI: 10.3390/s100605469.

[Fin+21] Teresa M Findley et al. "Sniff-synchronized, gradient-guided olfactory search by freely moving mice." en. In: *Elife* 10 (2021), e58523. DOI: 10.7554/eLife.58523.

[Fis+24] Georg K. J. Fischer et al. *Evaluation of a Smart Mobile Robotic System for Industrial Plant Inspection and Supervision*. en. arXiv:2402.07691 [cs]. 2024.

[FKS06] Marion E. Franke, Tobias J. Koplin, and Ulrich Simon. "Metal and Metal Oxide Nanoparticles in Chemiresistors: Does the Nanoscale Matter?" en. In: *Small* 2.1 (2006), pp. 36–50. DOI: 10.1002/smll.200500261.

[FL95] M S Freund and N S Lewis. "A chemically diverse conducting polymer-based "electronic nose"." en. In: *Proceedings of the national academy of sciences* 92.7 (1995), pp. 2652–2656. DOI: 10.1073/pnas.92.7.2652.

[Fon+14] Jordi Fonollosa, Irene Rodríguez-Luján, Marco Trincavelli, Alexander Vergara, and Ramón Huerta. "Chemical Discrimination in Turbulent Gas Mixtures with MOX Sensors Validated by Gas Chromatography-Mass Spectrometry." en. In: *Sensors* 14.10 (2014), pp. 19336–19353. DOI: 10.3390/s141019336.

[Fon+15] Jordi Fonollosa, Sadique Sheik, Ramón Huerta, and Santiago Marco. "Reservoir computing compensates slow response of chemosensor arrays exposed to fast varying gas concentrations in continuous monitoring." en. In: *Sensors and actuators b: chemical* 215 (2015), pp. 618–629. DOI: 10.1016/j.snb.2015.03.028.

[Fon+16] J. Fonollosa, L. Fernández, A. Gutiérrez-Gálvez, R. Huerta, and S. Marco. "Calibration transfer and drift counteraction in chemical sensor arrays using Direct Standardization." In: *Sensors and actuators b: chemical* 236 (2016), pp. 1044–1053. DOI: https://doi.org/10.1016/j.snb.2016.05.089.

[FR82a] J. E. Fackrell and A. G. Robins. "Concentration fluctuations and fluxes in plumes from point sources in a turbulent boundary layer." en. In: *Journal of fluid mechanics* 117 (1982), pp. 1–26. DOI: 10.1017/S0022112082001499.

[FR82b] J. E. Fackrell and A. G. Robins. "The effects of source size on concentration fluctuations in plumes." en. In: *Boundary-layer meteorology* 22.3 (1982), pp. 335–350. DOI: 10.1007/BF00120014.

[Fra+22] Adam Francis, Shuai Li, Christian Griffiths, and Johann Sienz. "Gas source localization and mapping with mobile robots: A review." en. In: *Journal of field robotics* 39.8 (2022), pp. 1341–1373. DOI: 10.1002/rob.22109.

[FSM18] Jordi Fonollosa, Ana Solórzano, and Santiago Marco. "Chemical Sensor Systems and Associated Algorithms for Fire Detection: A Review." en. In: *Sensors* 18.2 (2018), p. 553. DOI: 10.3390/s18020553.

[FT06] Robert Frodl and Thomas Tille. "A High-Precision NDIR $\hbox{CO}_{2}$ Gas Sensor for Automotive Applications." en. In: *Ieee sensors journal* 6.6 (2006), pp. 1697–1705. DOI: 10.1109/JSEN.2006.884440.

[Fuk+12] Izumi Fukunaga, Manuel Berning, Mihaly Kollo, Anja Schmaltz, and Andreas T. Schaefer. "Two Distinct Channels of Olfactory Bulb Output." In: *Neuron* 75.2 (2012), pp. 320–329. DOI: 10.1016/j.neuron.2012.05.017.

[Gal+20] Guillermo Gallego et al. "Event-based vision: A survey." In: *Ieee transactions on pattern analysis and machine intelligence* 44.1 (2020). Publisher: IEEE, pp. 154–180.

[Gam+19] Juan C. Rodriguez Gamboa, Eva Susana Albarracin E, Adenilton J. da Silva, and Tiago A. E. Ferreira. "Electronic nose dataset for detection of wine spoilage thresholds." In: *Data in brief* 25 (2019), p. 104202. DOI: https://doi.org/10.1016/j.dib.2019.104202.

[Gam+21] Juan C. Rodriguez Gamboa, Adenilton J. da Silva, Ismael C. S. Araujo, Eva Susana Albarracin E, and Cristhian M. Duran A. "Validation of the rapid detection approach for enhancing the electronic nose systems performance, using different deep learning models and support vector machines." In: *Sensors and actuators b: chemical* 327 (2021). Publisher: Elsevier BV, p. 128921. DOI: 10.1016/j.snb.2020.128921.

[Gao+11] Gaozhi Xiao et al. "Trace amount formaldehyde gas detection for indoor air quality monitoring." en. In: *2011 IEEE International Instrumentation and Measurement Technology Conference*. Binjiang: IEEE, 2011, pp. 1–4. DOI: 10.1109/IMTC.2011.5944050.

[Gar+10] Julian W Gardner, Prasanta K Guha, Florin Udrea, and James A Covington. "CMOS Interfacing for Integrated Gas Sensors: A Review." en. In: *Ieee sensors journal* 10.12 (2010), pp. 1833–1848. DOI: 10.1109/JSEN.2010.2046409.

[Gar+13] James Gareth, Witten Daniela, Hastie Trevor, and Tibshirani Robert. *An introduction to statistical learning: with applications in R*. New York, NY: Springer, 2013. DOI: https://doi.org/10.1007/978-1-4614-7138-7.

[Gar+94] Julian W. Gardner, Tim C. Pearce, Sharon Friel, Philip N. Bartlett, and Neil Blair. "A multisensor system for beer flavour monitoring using an array of conducting polymers and predictive classifiers." en. In: *Sensors and actuators b: chemical* 18.1-3 (1994), pp. 240–243. DOI: 10.1016/0925-4005(94)87089-6.

[GB21] Kyle Gilman and Laura Balzano. *Grassmannian Optimization for Online Tensor Completion and Tracking with the t-SVD*. _eprint: 2001.11419. 2021.

[Gef+09] Maria N. Geffen, Bede M. Broome, Gilles Laurent, and Markus Meister. "Neural Encoding of Rapidly Fluctuating Odors." en. In: *Neuron* 61.4 (2009), pp. 570–586. DOI: 10.1016/j.neuron.2009.01.021.

[GFS24] Michele Galvani, Sonia Freddi, and Luigi Sangaletti. "Disclosing Fast Detection Opportunities with Nanostructured Chemiresistor Gas Sensors Based on Metal Oxides, Carbon, and Transition Metal Dichalcogenides." en. In: *Sensors* 24.2 (2024), p. 584. DOI: 10.3390/s24020584.

[GGU23] Ethan L. W. Gardner, Julian W. Gardner, and Florin Udrea. "Micromachined Thermal Gas Sensors—A Review." en. In: *Sensors* 23.2 (2023), p. 681. DOI: 10.3390/s23020681.

[Gir+16] David H. Gire, Vikrant Kapoor, Annie Arrighi-Allisan, Agnese Seminara, and Venkatesh N. Murthy. "Mice Develop Efficient Strategies for Foraging and Navigation Using Complex Natural Stimuli." en. In: *Current biology* 26.10 (2016), pp. 1261–1273. DOI: 10.1016/j.cub.2016.03.040.

[Gon+11a] Javier Gonzalez, Javier G. Monroy, Francisco Garcia, and Jose Luis Blanco. "The multi-chamber electronic nose (MCE-nose)." en. In: *2011 IEEE International Conference on Mechatronics*. Istanbul, Turkey: IEEE, 2011, pp. 636–641. DOI: 10.1109/ICMECH.2011.5971193.

[Gon+11b] Javier Gonzalez, Javier G. Monroy, Francisco Garcia, and Jose Luis Blanco. "The multi-chamber electronic nose (mce-nose)." In: *2011 ieee international conference on mechatronics*. 2011, pp. 636–641. DOI: 10.1109/ICMECH.2011.5971193.

[Gor+17] Srinivas Gorur-Shandilya, Mahmut Demir, Junjiajia Long, Damon A Clark, and Thierry Emonet. "Olfactory receptor neurons use gain control and complementary kinetics to encode intermittent odorant stimuli." en. In: *Elife* 6 (2017), e27670. DOI: 10.7554/eLife.27670.

[GR10] C. Giovanni Galizia and Wolfgang Rössler. "Parallel Olfactory Systems in Insects: Anatomy and Function." en. In: *Annual review of entomology* 55.1 (2010), pp. 399–420. DOI: 10.1146/annurev-ento-112408-085442.

[GS09] Robert Gütig and Haim Sompolinsky. "Time-Warp–Invariant Neuronal Processing." en. In: *Plos biology* 7.7 (2009). Ed. by Michael Robert DeWeese, e1000141. DOI: 10.1371/journal.pbio.1000141.

[GS24] Daniel Gehrig and Davide Scaramuzza. "Low-latency automotive vision with event cameras." en. In: *Nature* 629.8014 (2024), pp. 1034–1040. DOI: 10.1038/s41586-024-07409-w.

[GSD16] Leonid Gugel, Yoel Shkolnisky, and Shai Dekel. *Machine olfaction using time scattering of sensor multiresolution graphs*. _eprint: 1602.04358. 2016.

[Guo+21] Wenzhe Guo, Mohammed E. Fouda, Ahmed M. Eltawil, and Khaled Nabil Salama. "Neural Coding in Spiking Neural Networks: A Comparative Study for Robust Neuromorphic Systems." en. In: *Frontiers in neuroscience* 15 (2021), p. 638474. DOI: 10.3389/fnins.2021.638474.

[Hal+13] R. Halir, L. Vivien, X. Le Roux, D. X. Xu, and P. Cheben. "Direct and sensitive phase readout for integrated waveguide sensors." In: *Ieee photonics journal* 5.4 (2013). DOI: 10.1109/JPHOT.2013.2276747.

[Han+22] Joon-Kyu Han, Mingu Kang, Jaeseok Jeong, Incheol Cho, Ji-Man Yu, Kuk-Jin Yoon, Inkyu Park, and Yang-Kyu Choi. "Artificial Olfactory Neuron for an In-Sensor Neuromorphic Nose." en. In: *Advanced science* 9.18 (2022), p. 2106017. DOI: 10.1002/advs.202106017.

[Has+15] Muhammad Hassan, Amine Bermak, Amine Ait Si Ali, and Abbes Amira. "Gas identification with spike codes in wireless electronic nose: A potential application for smart green buildings." en. In: *2015 SAI Intelligent Systems Conference (IntelliSys)*. London, United Kingdom: IEEE, 2015, pp. 457–462. DOI: 10.1109/IntelliSys.2015.7361180.

[HDC06] Elissa A. Hallem, Anupama Dahanukar, and John R. Carlson. "INSECT ODOR AND TASTE RECEPTORS." en. In: *Annual review of entomology* 51.1 (2006), pp. 113–135. DOI: 10.1146/annurev.ento.51.051705.113646.

[Her+12] Victor Hernandez Bennetts, Achim J. Lilienthal, Patrick P. Neumann, and Marco Trincavelli. "Mobile Robots for Localizing Gas Emission Sources on Landfill Sites: Is Bio-Inspiration the Way to Go?" en. In: *Frontiers in neuroengineering* 4 (2012). DOI: 10.3389/fneng.2011.00020.

[Her+22] Cyril Herrier, Nevena Morel, Romain Dubreuil, Jie Liu, Bertrand Gautheron, Loic Laplatine, Tristan Rousselle, and Thierry Livache. "Imaging a smell: from plasmonic-based device to an array of Mach-Zehnder interferometers." In: SPIE-Intl Soc Optical Eng, 2022, p. 9. DOI: 10.1117/12.2622141.

[HI89] Steven R Hanna and Elizabeth M Insley. "Time series analyses of concentration and wind fluctuations." In: *Boundary-layer meteorology* 47.1 (1989), pp. 131–147.

[Hin+23] Patrick Hinsen, Thomas Wiedemann, Dmitriy Shutin, and Achim J. Lilienthal. "Exploration and Gas Source Localization in Advection–Diffusion Processes with Potential-Field-Controlled Robotic Swarms." en. In: *Sensors* 23.22 (2023), p. 9232. DOI: 10.3390/s23229232.

[HK12] Hung-Yi Hsieh and Kea-Tiong Tang. "VLSI Implementation of a Bio-Inspired Olfactory Spiking Neural Network." en. In: *Ieee transactions on neural networks and learning systems* 23.7 (2012), pp. 1065–1073. DOI: 10.1109/TNNLS.2012.2195329.

[HLC19] Michael Hutchinson, Cunjia Liu, and Wen-Hua Chen. "Source term estimation of a hazardous airborne release using an unmanned aerial vehicle." en. In: *Journal of field robotics* 36.4 (2019), pp. 797–817. DOI: 10.1002/rob.21844.

[HLG99] E Llobet Hines, E Llobet, and JW Gardner. "Electronic noses: a review of signal processing techniques." In: *Iee proceedings-circuits, devices and systems* 146.6 (1999). Publisher: IET, pp. 297–310.

[Ho95] Tin Kam Ho. "Random decision forests." In: *Proceedings of 3rd international conference on document analysis and recognition*. Vol. 1. IEEE, 1995, pp. 278–282.

[Hon+00] Hyung-Ki Hong, Chul Han Kwon, Seung-Ryeol Kim, Dong Hyun Yun, Kyuchung Lee, and Yung Kwon Sung. "Portable electronic nose system with gas sensor array and artificial neural network." en. In: *Sensors and actuators b: chemical* 66.1-3 (2000), pp. 49–52. DOI: 10.1016/S0925-4005(99)00460-8.

[Hop91] J J Hopfield. "Olfactory computation and object perception." en. In: *Proceedings of the national academy of sciences* 88.15 (1991), pp. 6462–6466. DOI: 10.1073/pnas.88.15.6462.

[HPB10] Nelika K. Hughes, Catherine J. Price, and Peter B. Banks. "Predators Are Attracted to the Olfactory Signals of Prey." en. In: *Plos one* 5.9 (2010). Ed. by Sean Rands, e13114. DOI: 10.1371/journal.pone.0013114.

[HT19] Aixiang He and Zhenan Tang. "A Novel Gas Identification Method Based on Gabor Spectrogram Using Self-adapted Temperature Modulated Gas Sensors." en. In: *2019 International Conference on Sensing, Diagnostics, Prognostics, and Control (SDPC)*. Beijing, China: IEEE, 2019, pp. 714–717. DOI: 10.1109/SDPC.2019.00135.

[HTK08] D.J. Harvey, Tien-Fu Lu, and M.A. Keller. "Comparing Insect-Inspired Chemical Plume Tracking Algorithms Using a Mobile Robot." en. In: *Ieee transactions on robotics* 24.2 (2008), pp. 307–317. DOI: 10.1109/TRO.2007.912090.

[Huo+23] Dexuan Huo et al. "A Bio-Inspired Spiking Neural Network with Few-Shot Class-Incremental Learning for Gas Recognition." en. In: *Sensors* 23.5 (2023), p. 2433. DOI: 10.3390/s23052433.

[IC20] Nabil Imam and Thomas A. Cleland. "Rapid online learning and robust recall in a neuromorphic olfactory circuit." In: *Nature machine intelligence* 2.3 (2020). Publisher: Springer Science and Business Media LLC, pp. 181–191. DOI: 10.1038/s42256-020-0159-4.

[Ind+11] Giacomo Indiveri et al. "Neuromorphic silicon neuron circuits." In: *Frontiers in neuroscience* 5.MAY (2011), pp. 1–23. DOI: 10.3389/fnins.2011.00073.

[Inta] Intel.Newsroom. *Computers that smell: intel's neuromorphic chip can sniff out hazardous chemicals*. https://newsroom.intel.com/news/computers-smell-intels-neuromorphic-chip-sniff-hazardous-chemicals/. Accessed: 2022-10-24.

[Intb] Intel.Newsroom. *How a computer chip can smell without a nose*. https://newsroom.intel.com/news/how-computer-chip-smell-without-nose/. Accessed: 2022-10-24.

[ISY16] Gaku Imamura, Kota Shiba, and Genki Yoshikawa. "Smell identification of spices using nanomechanical membrane-type surface stress sensors." In: *Japanese journal of applied physics* 55.11 (2016). DOI: 10.7567/JJAP.55.1102B3.

[Ito+08] Iori Ito, Rose Chik-ying Ong, Baranidharan Raman, and Mark Stopfer. "Sparse odor representation and olfactory learning." en. In: *Nature neuroscience* 11.10 (2008), pp. 1177–1184. DOI: 10.1038/nn.2192.

[Jac+15] Lucia F. Jacobs, Jennifer Arter, Amy Cook, and Frank J. Sulloway. "Olfactory Orientation and Navigation in Humans." en. In: *Plos one* 10.6 (2015). Ed. by Matthieu Louis, e0129387. DOI: 10.1371/journal.pone.0129387.

[Jay+23] Viraaj Jayaram, Aarti Sehdev, Nirag Kadakia, Ethan A. Brown, and Thierry Emonet. "Temporal novelty detection and multiple timescale integration drive Drosophila orientation dynamics in temporally diverse olfactory environments." en. In: *Plos computational biology* 19.5 (2023). Ed. by Barbara Webb, e1010606. DOI: 10.1371/journal.pcbi.1010606.

[JDM14] Asa Johannesen, Alison M. Dunn, and Lesley J. Morrell. "Prey aggregation is an effective olfactory predator avoidance strategy." en. In: *Peerj* 2 (2014), e408. DOI: 10.7717/peerj.408.

[Jin+17] Ya-Qi Jing, Qing-Hao Meng, Pei-Feng Qi, Ming Zeng, and Ying-Jie Liu. "Signal processing inspired from the olfactory bulb for electronic noses." en. In: *Measurement science and technology* 28.1 (2017), p. 015105. DOI: 10.1088/1361-6501/28/1/015105.

[JMI21] Tao Jing, Qing-Hao Meng, and Hiroshi Ishida. "Recent progress and trend of robot odor source localization." In: *Ieej transactions on electrical and electronic engineering* 16.7 (2021). Publisher: Wiley Online Library, pp. 938–953.

[JMJ02] Kristine A Justus, John Murlis, and Chris Jones. "Measurement of Odor-Plume Structure in a Wind Tunnel Using a Photoionization Detector and a Tracer Gas." en. In: *Environmental fluid mechanics* 2 (2002), pp. 115–142.

[Jun+23] Gyuweon Jung, Jaehyeon Kim, Seongbin Hong, Hunhee Shin, Yujeong Jeong, Wonjun Shin, Dongseok Kwon, Woo Young Choi, and Jong-Ho Lee. "Energy Efficient Artificial Olfactory System with Integrated Sensing and Computing Capabilities for Food Spoilage Detection." en. In: *Advanced science* 10.30 (2023), p. 2302506. DOI: 10.1002/advs.202302506.

[Jür+21] Anna-Maria Jürgensen, Afshin Khalili, Elisabetta Chicca, Giacomo Indiveri, and Martin Paul Nawrot. "A neuromorphic model of olfactory processing and sparse coding in the Drosophila larva brain." en. In: *Neuromorphic computing and engineering* 1.2 (2021), p. 024008. DOI: 10.1088/2634-4386/ac3ba6.

[Kad+22] Nirag Kadakia, Mahmut Demir, Brenden T. Michaelis, Brian D. DeAngelis, Matthew A. Reidenbach, Damon A. Clark, and Thierry Emonet. "Odour motion sensing enhances navigation of complex plumes." en. In: *Nature* 611.7937 (2022), pp. 754–761. DOI: 10.1038/s41586-022-05423-4.

[Kan+00] Eric R Kandel, James H Schwartz, Thomas M Jessell, Steven Siegelbaum, A James Hudspeth, Sarah Mack, et al. *Principles of neural science*. Vol. 4. McGraw-hill New York, 2000.

[Kas+16] Nikola Kasabov et al. "Evolving spatio-temporal data machines based on the NeuCube neuromorphic framework: Design methodology and selected applications." en. In: *Neural networks* 78 (2016), pp. 1–14. DOI: 10.1016/j.neunet.2015.09.011.

[KBP09] A. Kroll, W. Baetz, and D. Peretzki. "On autonomous detection of pressured air and gas leaks using passive IR-thermography for mobile robot application." en. In: *2009 IEEE International Conference on Robotics and Automation*. Kobe: IEEE, 2009, pp. 921–926. DOI: 10.1109/ROBOT.2009.5152337.

[KC21] Heonmoo Kim and Yosoon Choi. "Self-Driving Algorithm and Location Estimation Method for Small Environmental Monitoring Robot in Underground Mines." en. In: *Computer modeling in engineering & sciences* 127.3 (2021), pp. 943–964. DOI: 10.32604/cmes.2021.015300.

[Ken20] S Saladin Kenneth. *Anatomy & physiology: the unity of form and function*. McGraw Hill, 2020.

[KH20] Tamara G. Kolda and David Hong. "Stochastic Gradients for Large-Scale Tensor Decomposition." In: *Siam journal on mathematics of data science* 2.4 (2020). Publisher: Society for Industrial & Applied Mathematics (SIAM), pp. 1066–1095. DOI: 10.1137/19m1266265.

[Kha+21] Mohammed A. Khallaf, Rongfeng Cui, Jerrit Weißflog, Maide Erdogmus, Aleš Svatoš, Hany K. M. Dweck, Dario Riccardo Valenzano, Bill S. Hansson, and Markus Knaden. "Large-scale characterization of sex pheromone communication systems in Drosophila." en. In: *Nature communications* 12.1 (2021), p. 4165. DOI: 10.1038/s41467-021-24395-z.

[KLS11] Anmo J. Kim, Aurel A. Lazar, and Yevgeniy B. Slutskiy. "System identification of Drosophila olfactory sensory neurons." en. In: *Journal of computational neuroscience* 30.1 (2011), pp. 143–161. DOI: 10.1007/s10827-010-0265-0.

[KLS15] Anmo J Kim, Aurel A Lazar, and Yevgeniy B Slutskiy. "Projection neurons in Drosophila antennal lobes signal the acceleration of odor concentrations." en. In: *Elife* 4 (2015), e06651. DOI: 10.7554/eLife.06651.

[KM74] J. S. Kennedy and D. Marsh. "Pheromone-Regulated Anemotaxis in Flying Moths." en. In: *Science* 184.4140 (1974), pp. 999–1001. DOI: 10.1126/science.184.4140.999.

[Koi+06] T.J. Koickal, A. Hamilton, Tim C. Pearce, S.L. Tan, J.A. Covington, and J.W. Gardner. "Analog VLSI Design of an Adaptive Neuromorphic Chip for Olfactory Systems." en. In: *2006 IEEE International Symposium on Circuits and Systems*. Island of Kos, Greece: IEEE, 2006, pp. 4547–4550. DOI: 10.1109/ISCAS.2006.1693641.

[Koi+07] Thomas Jacob Koickal, Alister Hamilton, Su Lim Tan, James A. Covington, Julian W. Gardner, and Tim C. Pearce. "Analog VLSI Circuit Implementation of an Adaptive Neuromorphic Olfaction Chip." en. In: *Ieee transactions on circuits and systems i: regular papers* 54.1 (2007), pp. 60–73. DOI: 10.1109/TCSI.2006.888677.

[Kol41] Andrey Nikolaevich Kolmogorov. "The local structure of turbulence in incompressible viscous fluid for very large reynolds." In: *Numbers. in dokl. akad. nauk sssr* 30 (1941), p. 301.

[KR08] Gideon Kowadlo and R. Andrew Russell. "Robot Odor Localization: A Taxonomy and Survey." en. In: *The international journal of robotics research* 27.8 (2008), pp. 869–894. DOI: 10.1177/0278364908095118.

[Kro08] Sabine Krofczik. "Rapid odor processing in the honeybee antennal lobe network." en. In: *Frontiers in computational neuroscience* 2 (2008). DOI: 10.3389/neuro.10.009.2008.

[KSB12] Adil Ghani Khan, Manaswini Sarangi, and Upinder Singh Bhalla. "Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling." en. In: *Nature communications* 3.1 (2012), p. 703. DOI: 10.1038/ncomms1712.

[Kul+05] S. Kulinyi, D. Brandszájsz, H. Amine, M. Ádám, P. Fürjes, I. Bársony, and Cs. Dücső. "Olfactory detection of methane, propane, butane and hexane using conventional transmitter norms." en. In: *Sensors and actuators b: chemical* 111-112 (2005), pp. 286–292. DOI: 10.1016/j.snb.2005.06.068.

[KUM06] Adam Kepecs, Naoshige Uchida, and Zachary F. Mainen. "The Sniff as a Unit of Olfactory Processing." en. In: *Chemical senses* 31.2 (2006), pp. 167–179. DOI: 10.1093/chemse/bjj016.

[Lan+99] H P Lang et al. "An artificial nose based on a micromechanical cantilever array." en. In: *Analytica chimica acta* (1999).

[Lap+22] Loic Laplatine et al. "Silicon photonic olfactory sensor based on an array of 64 biofunctionalized Mach-Zehnder interferometers." In: *Optics express* 30.19 (2022), p. 33955. DOI: 10.1364/OE.461858.

[LC16] Christiane Linster and Thomas A Cleland. "Neuromodulation of olfactory transformations." en. In: *Current opinion in neurobiology* 40 (2016), pp. 170–177. DOI: 10.1016/j.conb.2016.07.006.

[LD04] Achim Lilienthal and Tom Duckett. "Experimental analysis of gas-sensitive Braitenberg vehicles." en. In: *Advanced robotics* 18.8 (2004), pp. 817–834. DOI: 10.1163/1568553041738103.

[LD10] Shih Chii Liu and Tobi Delbruck. "Neuromorphic sensory systems." In: *Current opinion in neurobiology* 20.3 (2010). Publisher: Elsevier Ltd, pp. 288–295. DOI: 10.1016/j.conb.2010.03.007.

[Lee+17] Ki-Seong Lee, Sun-Ro Lee, Youngmin Kim, and Chan-Gun Lee. "Deep learning–based real-time query processing for wireless sensor network." In: *International journal of distributed sensor networks* 13.5 (2017). Publisher: SAGE Publications, p. 155014771770789. DOI: 10.1177/1550147717707896.

[Lef20] Melanie Lefkowitz. *Researchers sniff out ai breakthroughs in mammal brains.* https://news.cornell.edu/stories/2020/03/researchers-sniff-out-ai-breakthroughs-mammal-brains. Accessed: 2022-11-22. 2020.

[Lew+24] Suzanne M. Lewis, Lucas M. Suarez, Nicola Rigolli, Nicholas A. Steinmetz, and David H. Gire. *The spiking output of the mouse olfactory bulb encodes large-scale temporal features of natural odor environments*. en. 2024. DOI: 10.1101/2024.03.01.582978.

[Li+14] Anan Li, David H. Gire, Thomas Bozza, and Diego Restrepo. "Precise Detection of Direct Glomerular Input Duration by the Olfactory Bulb." en. In: *The journal of neuroscience* 34.48 (2014), pp. 16058–16064. DOI: 10.1523/JNEUROSCI.3382-14.2014.

[Li+21] Jun Li, Zhiqiang Tao, Yue Wu, Bineng Zhong, and Yun Fu. "Large-Scale Subspace Clustering by Independent Distributed and Parallel Coding." In: *Ieee transactions on cybernetics* 1.c (2021), pp. 1–11. DOI: 10.1109/TCYB.2021.3052056.

[Li+23] Yan Li et al. "Robust odor identification in novel olfactory environments in mice." en. In: *Nature communications* 14.1 (2023), p. 673. DOI: 10.1038/s41467-023-36346-x.

[Lil+01] A. Lilienthal, A. Zell, M. Wandel, and U. Weimar. "Sensing odour sources in indoor environments without a constant airflow by a mobile robot." en. In: *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*. Vol. 4. Seoul, South Korea: IEEE, 2001, pp. 4005–4010. DOI: 10.1109/ROBOT.2001.933243.

[Lip+96] Christopher A Lipinski, Franco Lombardo, Beryl W Dominy, and Paul J Feeney. "Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings." en. In: *Advanced drug delivery reviews* (1996).

[Liu+10] Shih-Chii Liu, Andre Van Schaik, Bradley A. Mincti, and Tobi Delbruck. "Event-based 64-channel binaural silicon cochlea with Q enhancement mechanisms." In: *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. Paris, France: IEEE, 2010, pp. 2027–2030. DOI: 10.1109/ISCAS.2010.5537164.

[Liu+12] Xiao Liu, Sitian Cheng, Hong Liu, Sha Hu, Daqiang Zhang, and Huansheng Ning. "A Survey on Gas Sensing Technology." en. In: *Sensors* 12.7 (2012), pp. 9635–9665. DOI: 10.3390/s120709635.

[Liu+18] Haotian Liu, Li Zhang, King Li, and Ooi Tan. "Microhotplates for Metal Oxide Semiconductor Gas Sensor Applications—Towards the CMOS-MEMS Monolithic Approach." en. In: *Micromachines* 9.11 (2018), p. 557. DOI: 10.3390/mi9110557.

[Liu+20] Annie Liu, Andrew E. Papale, James Hengenius, Khusbu Patel, Bard Ermentrout, and Nathan N. Urban. "Mouse Navigation Strategies for Odor Source Localization." en. In: *Frontiers in neuroscience* 14 (2020), p. 218. DOI: 10.3389/fnins.2020.00218.

[LLD06] Achim J. Lilienthal, Amy Loutfi, and Tom Duckett. "Airborne Chemical Sensing with Mobile Robots." In: *Sensors* 6.11 (2006), pp. 1616–1678. DOI: 10.3390/s6111616.

[Loc+08a] Thomas Lochmatter, Xavier Raemy, Loic Matthey, Saurabh Indra, and Alcherio Martinoli. "A comparison of casting and spiraling algorithms for odor source localization in laminar flow." en. In: *2008 IEEE International Conference on Robotics and Automation*. Pasadena, CA, USA: IEEE, 2008, pp. 1138–1143. DOI: 10.1109/ROBOT.2008.4543357.

[Loc+08b] Thomas Lochmatter, Pierre Roduit, Chris Cianci, Nikolaus Correll, Jacques Jacot, and Alcherio Martinoli. "SwisTrack - A Flexible Open Source Tracking Software for Multi-Agent Systems." en. In: *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Nice: IEEE, 2008, pp. 4004–4010. DOI: 10.1109/IROS.2008.4650937.

[Loi+13] Frédéric Loizeau, Hans Peter Lang, Terunobu Akiyama, Sebastian Gautsch, Peter Vettiger, Andreas Tonin, Genki Yoshikawa, Christoph Gerber, and Nico de Rooij. "Piezoresistive membrane-type surface stress sensor arranged in arrays for cancer diagnosis through breath analysis." In: *2013 IEEE 26th International Conference on Micro Electro Mechanical Systems (MEMS)*. IEEE, 2013, pp. 621–624.

[LPD08] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. "A 128 x 128 120 dB 15 microsecond latency asynchronous temporal contrast vision sensor." In: *Ieee journal of solid-state circuits* 43.2 (2008), pp. 566–576. DOI: 10.1109/JSSC.2007.914337.

[Lu13] Tien-Fu Lu. "Indoor odour source localisation using robot: Initial location and surge distance matter?" en. In: *Robotics and autonomous systems* 61.6 (2013), pp. 637–647. DOI: 10.1016/j.robot.2013.02.002.

[LW71] Michael Levandowsky and David Winter. "Distance between sets." In: *Nature* 234.5323 (1971). Publisher: Nature Publishing Group UK London, pp. 34–35.

[MA91] P. A. Moore and J. Atema. "Spatial Information in the Three-Dimensional Fine Structure of an Aquatic Odor Plume." en. In: *The biological bulletin* 181.3 (1991), pp. 408–418. DOI: 10.2307/1542361.

[Maa15] Wolfgang Maass. "To Spike or Not to Spike: That Is the Question." en. In: *Proceedings of the ieee* 103.12 (2015), pp. 2219–2224. DOI: 10.1109/JPROC.2015.2496679.

[Maa97] Wolfgang Maass. "Networks of spiking neurons: The third generation of neural network models." en. In: *Neural networks* 10.9 (1997), pp. 1659–1671. DOI: 10.1016/S0893-6080(97)00011-7.

[Mah+20] Pierre Maho, Cyril Herrier, Thierry Livache, Guillaume Rolland, Pierre Comon, and Simon Barthelmé. "Reliable chiral recognition with an optoelectronic nose." en. In: *Biosensors and bioelectronics* 159 (2020), p. 112183. DOI: 10.1016/j.bios.2020.112183.

[Mah+21] Pierre Maho, Cyril Herrier, Thierry Livache, Pierre Comon, and Simon Barthelmé. "Real-time gas recognition and gas unmixing in robot applications." en. In: *Sensors and actuators b: chemical* 330 (2021), p. 129111. DOI: 10.1016/j.snb.2020.129111.

[Mai+06] DE Maier, R Hulasare, B Qian, P Armstrong, et al. "Monitoring carbon dioxide levels for early detection of spoilage and pests in stored grain." In: *Proceedings of the 9th International Working Conference on Stored Product Protection*. Vol. 1. 2006, p. 117.

[Mal+99] Bettina Malnic, Junzo Hirono, Takaaki Sato, and Linda B Buck. "Combinatorial receptor codes for odors." In: *Cell* 96.5 (1999). Publisher: Elsevier, pp. 713–723.

[Mam+13] Syed Muhammad Mamduh, Kamarulzaman Kamarudin, Shaharil Mad Saad, Ali Yeon Md Shakaff, Ammar Zakaria, and Abu Hassan Abdullah. "Braitenberg swarm vehicles for odour plume tracking in laminar airflow." en. In: *2013 IEEE Symposium on Computers & Informatics (ISCI)*. Langkawi, Malaysia: IEEE, 2013, pp. 1–6. DOI: 10.1109/ISCI.2013.6612365.

[Mar+14a] S. Marco et al. "A biomimetic approach to machine olfaction, featuring a very large-scale chemical sensor array and embedded neuro-bio-inspired computation." en. In: *Microsystem technologies* 20.4-5 (2014), pp. 729–742. DOI: 10.1007/s00542-013-2020-8.

[Mar+14b] Dominique Martinez, Lotfi Arhidi, Elodie Demondion, Jean-Baptiste Masson, and Philippe Lucas. "Using insect electroantennogram sensors on autonomous robots for olfactory searches." In: *Jove (journal of visualized experiments)* 90 (2014), e51704.

[Mat+19] Daniel Matatagui, Fabio Andrés Bahos, Isabel Gràcia, and María Del Carmen Horrillo. "Portable Low-Cost Electronic Nose Based on Surface Acoustic Wave Sensors for the Detection of BTX Vapors in Air." en. In: *Sensors* 19.24 (2019), p. 5406. DOI: 10.3390/s19245406.

[May+22] Emily J. Mayhew, Charles J. Arayata, Richard C. Gerkin, Brian K. Lee, Jonathan M. Magill, Lindsey L. Snyder, Kelsie A. Little, Chung Wen Yu, and Joel D. Mainland. "Transport features predict if a molecule is odorous." en. In: *Proceedings of the national academy of sciences* 119.15 (2022), e2116576119. DOI: 10.1073/pnas.2116576119.

[MBM19] Martinez, Burgués, and Marco. "Fast Measurements with MOX Sensors: A Least-Squares Approach to Blind Deconvolution." en. In: *Sensors* 19.18 (2019), p. 4029. DOI: 10.3390/s19184029.

[MC04] Paul Moore and John Crimaldi. "Odor landscapes and animal behavior: tracking odor plumes in different physical worlds." en. In: *Journal of marine systems* 49.1-4 (2004), pp. 55–64. DOI: 10.1016/j.jmarsys.2003.05.005.

[MC94] Agenor Mafra-Neto and Ring T. Cardé. "Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths." en. In: *Nature* 369.6476 (1994), pp. 142–144. DOI: 10.1038/369142a0.

[Mea90] Carver Mead. "Neuromorphic Electronic Systems." en. In: *Proceedings of the ieee* 78 (1990).

[MEC92] J Murlis, J S Elkinton, and R T Cardé. "Odor Plumes and How Insects Use Them." en. In: *Annu. rev. entomol.* (1992).

[Men+23] Fanli Meng, Xi Luan, Chunjin Mi, Hanyang Ji, Hongmin Zhu, and Zhenyu Yuan. "Recognition Algorithm for Detection of Precursor Chemicals by Semiconductor Gas Sensor Array Under Dynamic Measurement." In: *Ieee sensors journal* 23.3 (2023), pp. 1818–1826. DOI: 10.1109/JSEN.2022.3232179.

[MG15] Javier Monroy and Javier González-Jiménez. "REAL-TIME ODOR CLASSIFICATION THROUGH SEQUENTIAL BAYESIAN FILTERING." In: *16th International Symposium on Olfaction and Electronic Noses*. 2015.

[MGB12] Javier G. Monroy, Javier González-Jiménez, and Jose Luis Blanco. "Overcoming the Slow Recovery of MOX Gas Sensors through a System Modeling Approach." en. In: *Sensors* 12.10 (2012), pp. 13664–13680. DOI: 10.3390/s121013664.

[MHP97] T.H. Misselbrook, P.J. Hobbs, and K.C. Persaud. "Use of an Electronic Nose to Measure Odour Concentration Following Application of Cattle Slurry to Grassland." en. In: *Journal of agricultural engineering research* 66.3 (1997), pp. 213–220. DOI: 10.1006/jaer.1996.0135.

[Mis06a] Marek Miskowicz. "Send-On-Delta Concept: An Event-Based Data Reporting Strategy." In: *Sensors* 6.1 (2006), pp. 49–63. DOI: 10.3390/s6010049.

[Mis06b] Marek Miskowicz. "Send-on-delta concept: an event-based data reporting strategy." In: *Sensors* 6.1 (2006), pp. 49–63. DOI: 10.3390/s6010049.

[MM91] K. R. Mylne and P. J. Mason. "Concentration fluctuation measurements in a dispersing plume at a range of up to 1000 m." In: *Quarterly journal of the royal meteorological society* 117.497 (1991), pp. 177–206. DOI: 10.1002/qj.49711749709.

[Mon+16] Javier G. Monroy, Esteban J. Palomo, Ezequiel López-Rubio, and Javier Gonzalez-Jimenez. "Continuous chemical classification in uncontrolled environments with sliding windows." In: *Chemometrics and intelligent laboratory systems* 158 (2016). Publisher: Elsevier BV, pp. 117–129. DOI: 10.1016/j.chemolab.2016.08.011.

[Mon+17] Javier Monroy, Victor Hernandez-Bennetts, Han Fan, Achim Lilienthal, and Javier Gonzalez-Jimenez. "GADEN: A 3D Gas Dispersion Simulator for Mobile Robot Olfaction in Realistic Environments." en. In: *Sensors* 17.7 (2017), p. 1479. DOI: 10.3390/s17071479.

[Mor+09] Mélissa M. Moreno, Christiane Linster, Olga Escanilla, Joëlle Sacquet, Anne Didier, and Nathalie Mandairon. "Olfactory perceptual learning requires adult neurogenesis." en. In: *Proceedings of the national academy of sciences* 106.42 (2009), pp. 17980–17985. DOI: 10.1073/pnas.0907063106.

[MRS18] A. Mishra, N. S. Rajput, and D. Singh. "Performance evaluation of normalized difference based classifier for efficient discrimination of volatile organic compounds." In: *Materials research express* 5.9 (2018). Publisher: IOP Publishing, p. 095901. DOI: 10.1088/2053-1591/aad3dd.

[MSA21] Alina Cristina Marin, Andreas T Schaefer, and Tobias Ackels. "Spatial information from the odour environment in mammalian olfaction." en. In: *Cell and tissue research* 383.1 (2021), pp. 473–483. DOI: 10.1007/s00441-020-03395-3.

[Mue+09] Mehmet K. Muezzinoglu, Alexander Vergara, Ramon Huerta, Nikolai Rulkov, Mikhail I. Rabinovich, Al Selverston, and Henry D.I. Abarbanel. "Acceleration of chemo-sensory information processing using transient features." en. In: *Sensors and actuators b: chemical* 137.2 (2009), pp. 507–512. DOI: 10.1016/j.snb.2008.10.065.

[MYS20] Drew Mitchell, Nan Ye, and Hans De Sterck. "Nesterov acceleration of alternating least squares for canonical tensor decomposition: Momentum step size selection and restart mechanisms." In: *Numerical linear algebra with applications* 27.4 (2020). Publisher: Wiley. DOI: 10.1002/nla.2297.

[Nat+15] Brian J. Nathan et al. "Near-Field Characterization of Methane Emission Variability from a Compressor Station Using a Model Aircraft." en. In: *Environmental science & technology* 49.13 (2015), pp. 7896–7903. DOI: 10.1021/acs.est.5b00705.

[NBB11] Kwan Ting Ng, Farid Boussaid, and Amine Bermak. "A CMOS Single-Chip Gas Recognition Circuit for Metal Oxide Gas Sensor Arrays." en. In: *Ieee transactions on circuits and systems i: regular papers* 58.7 (2011), pp. 1569–1580. DOI: 10.1109/TCSI.2011.2143090.

[Neu+12] Patrick Neumann, Sahar Asadi, Achim Lilienthal, Matthias Bartholmai, and Jochen Schiller. "Autonomous Gas-Sensitive Microdrone: Wind Vector Estimation and Gas Distribution Mapping." en. In: *Ieee robotics & automation magazine* 19.1 (2012), pp. 50–61. DOI: 10.1109/MRA.2012.2184671.

[Neu+13a] P.P. Neumann, S. Asadi, V. Hernandez Bennetts, A.J. Lilienthal, and M. Bartholmai. "Monitoring of CCS Areas using Micro Unmanned Aerial Vehicles (MUAVs)." en. In: *Energy procedia* 37 (2013), pp. 4182–4190. DOI: 10.1016/j.egypro.2013.06.320.

[Neu+13b] Patrick P. Neumann, Victor Hernandez Bennetts, Achim J. Lilienthal, Matthias Bartholmai, and Jochen H. Schiller. "Gas source localization with a micro-drone using bio-inspired and particle filter-based algorithms." en. In: *Advanced robotics* 27.9 (2013), pp. 725–738. DOI: 10.1080/01691864.2013.779052.

[Ng+09] Kwan Ting Ng, Bin Guo, Amine Bermak, Dominique Martinez, and Farid Boussaid. "Characterization of a logarithmic spike timing encoding scheme for a 4×4 tin oxide gas sensor array." en. In: *2009 IEEE Sensors*. Christchurch, New Zealand: IEEE, 2009, pp. 731–734. DOI: 10.1109/ICSENS.2009.5398548.

[NHB19] Patrick P. Neumann, Dino Hüllmann, and Matthias Bartholmai. "Concept of a gas-sensitive nano aerial robot swarm for indoor air quality monitoring." en. In: *Materials today: proceedings* 12 (2019), pp. 470–473. DOI: 10.1016/j.matpr.2019.03.151.

[NL02] Alexander A Nikonov and Walter S Leal. "PERIPHERAL CODING OF SEX PHEROMONE AND A BEHAVIORAL ANTAGONIST IN THE JAPANESE BEETLE," en. In: *Journal of chemical ecology* (2002).

[NT20] Satoshi Nakata and Naho Takahara. "Characteristic nonlinear responses of a semiconductor gas sensor to hydrocarbons and alcohols under the combination of cyclic temperature and continuous flow." In: *Sensors and actuators b: chemical* 307 (2020). Publisher: Elsevier, p. 127635.

[NT22] Satoshi Nakata and Naho Takahara. "Distinction of gaseous mixtures based on different cyclic temperature modulations." In: *Sensors and actuators b: chemical* 359 (2022). Publisher: Elsevier, p. 131615.

[Obu41] Alexander Obukhov. "Spectral energy distribution in a turbulent flow." In: *Izv. akad. nauk. sssr. ser. geogr. i. geofiz* 5 (1941), pp. 453–466.

[OCo+01] Manuela O'Connell, Gabriela Valdora, Gustavo Peltzer, and R Martın Negri. "A practical approach for fish freshness determinations using a portable electronic nose." In: *Sensors and actuators b: chemical* 80.2 (2001), pp. 149–154.

[OF04] B Olshausen and D Field. "Sparse coding of sensory inputs." en. In: *Current opinion in neurobiology* 14.4 (2004), pp. 481–487. DOI: 10.1016/j.conb.2004.07.007.

[OMG23] Pepe Ojeda, Javier Monroy, and Javier Gonzalez-Jimenez. *Robotic Gas Source Localization with Probabilistic Mapping and Online Dispersion Simulation*. en. arXiv:2304.08879 [cs]. 2023.

[Ors73] SA Orszag. "Statistical theory of turbulence." In: *Fluid dynamics* (1973), pp. 237–374.

[Pan+18] Rich Pang, Floris Van Breugel, Michael Dickinson, Jeffrey A. Riffell, and Adrienne Fairhall. "History dependence in insect flight decisions during odor tracking." en. In: *Plos computational biology* 14.2 (2018). Ed. by Vikas Bhandawat, e1005969. DOI: 10.1371/journal.pcbi.1005969.

[Par+16] In Jun Park, Andrew M. Hein, Yuriy V. Bobkov, Matthew A. Reidenbach, Barry W. Ache, and Jose C. Principe. "Neurally Encoding Time for Olfactory Navigation." en. In: *Plos computational biology* 12.1 (2016). Ed. by Joseph Ayers, e1004682. DOI: 10.1371/journal.pcbi.1004682.

[Par+21] Minkyu Park, Seulbi An, Jaemin Seo, and Hyondong Oh. "Autonomous Source Search for UAVs Using Gaussian Mixture Model-Based Infotaxis: Algorithm and Flight Experiments." en. In: *Ieee transactions on aerospace and electronic systems* 57.6 (2021), pp. 4238–4254. DOI: 10.1109/TAES.2021.3098132.

[PD82] Krishna Persaud and George Dodd. "Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose." en. In: *Nature* 299.5881 (1982), pp. 352–355. DOI: 10.1038/299352a0.

[Pea+01] Tim Pearce, Paul Verschure, Joel White, and John Kauer. "Robust stimulus encoding in olfactory processing: hyperacuity and efficient signal transmission." In: *Emergent neural computational architectures based on neuroscience: towards neuroscience-inspired computing* (2001). Publisher: Springer, pp. 461–479.

[Pea+02] Tim C. Pearce, Susan S. Schiffman, H. Troy Nagle, and Julian W. Gardner, eds. *Handbook of Machine Olfaction: Electronic Nose Technology*. en. 1st ed. Wiley, 2002. DOI: 10.1002/3527601597.

[Pea+05] Tim C. Pearce, C. Fulvi-Mari, J.A. Covington, F.S. Tan, J.W. Gardner, T.J. Koickal, and A. Hamilton. "Silicon-based Neuromorphic Implementation of the Olfactory Pathway." en. In: *Conference Proceedings. 2nd International IEEE EMBS Conference on Neural Engineering, 2005.* Arlington, Virginia, USA: IEEE, 2005, pp. 307–312. DOI: 10.1109/CNE.2005.1419619.

[Pea+13] Tim C. Pearce, Salah Karout, Zoltán Rácz, Alberto Capurro, Julian W. Gardner, and Marina Cole. "Rapid processing of chemosensor transients in a neuromorphic implementation of the insect macroglomerular complex." en. In: *Frontiers in neuroscience* 7 (2013). DOI: 10.3389/fnins.2013.00119.

[Pea+93] Tim C. Pearce, Julian W. Gardner, Sharon Friel, Philip N. Bartlett, and Neil Blair. "Electronic nose for monitoring the flavour of beers." en. In: *The analyst* 118.4 (1993), p. 371. DOI: 10.1039/an9931800371.

[Pea97a] Tim C. Pearce. "Computational parallels between the biological olfactory pathway and its analogue 'The Electronic Nose': Part I. Biological olfaction." en. In: (1997).

[Pea97b] Tim C. Pearce. "Computational parallels between the biological olfactory pathway and its analogue 'The Electronic Nose': Part II. Sensor-based machine olfaction." en. In: *Biosystems* 41.2 (1997), pp. 69–90. DOI: 10.1016/S0303-2647(96)01660-7.

[Pfe+13] Thomas Pfeil et al. "Six Networks on a Universal Neuromorphic Computing Substrate." en. In: *Frontiers in neuroscience* 7 (2013). DOI: 10.3389/fnins.2013.00011.

[PG21] Marco Paoli and Giovanni C. Galizia. "Olfactory coding in honeybees." en. In: *Cell and tissue research* 383.1 (2021), pp. 35–58. DOI: 10.1007/s00441-020-03385-5.

[PGC07] Tim C. Pearce, Jing Gu, and Eric Chanie. "Chemical Source Classification in Naturally Turbulent Plumes." en. In: *Analytical chemistry* 79.22 (2007), pp. 8511–8519. DOI: 10.1021/ac0710376.

[Pho22] Julianna Photopoulos. "The dogs learning to sniff out disease." In: *Nature* (2022), S10–S11.

[PLT12a] Sepideh Pashami, Achim Lilienthal, and Marco Trincavelli. "Detecting Changes of a Distant Gas Source with an Array of MOX Gas Sensors." en. In: *Sensors* 12.12 (2012), pp. 16404–16419. DOI: 10.3390/s121216404.

[PLT12b] Sepideh Pashami, Achim J Lilienthal, and Marco Trincavelli. "Detecting changes of a distant gas source with an array of mox gas sensors." In: *Sensors* 12.12 (2012), pp. 16404–16419.

[PMG13] Krishna C Persaud, Santiago Marco, and Agustin Gutierrez-Galvez. "Neuromorphic olfaction." In: (2013). Publisher: CRC Press.

[PMW11] Christoph Posch, Daniel Matolin, and Rainer Wohlgenannt. "A QVGA 143 dB Dynamic Range Frame-Free PWM Image Sensor With Lossless Pixel-Level Video Compression and Time-Domain CDS." en. In: *Ieee journal of solid-state circuits* 46.1 (2011), pp. 259–275. DOI: 10.1109/JSSC.2010.2085952.

[Pra+23] Yaqub A. Prabowo, Bambang R. Trilaksono, Egi M. I. Hidayat, and Brian Yuliarto. "Multi-Robot Gas Sources Localization and Mapping Using Adaptive Voronoi-PSO and Bayesian Inference." en. In: *Ieee access* 11 (2023), pp. 135738–135752. DOI: 10.1109/ACCESS.2023.3336560.

[RAH08] Jeffrey A. Riffell, Leif Abrell, and John G. Hildebrand. "Physical Processes and Real-Time Chemical Measurement of the Insect Olfactory Environment." en. In: *Journal of chemical ecology* 34.7 (2008), pp. 837–853. DOI: 10.1007/s10886-008-9490-7.

[Ram+06] Baranidharan Raman, Ping A. Sun, Agustin Gutierrez-Galvez, and Ricardo Gutierrez-Osuna. "Processing of Chemical Sensor Arrays with a Biologically Inspired Model of Olfactory Coding." en. In: *Ieee transactions on neural networks* 17.4 (2006), pp. 1015–1024. DOI: 10.1109/TNN.2006.875975.

[Ram+11] Alejandro R. Garcia Ramirez, Amarilys Lima Lopez, Andy Blanco Rodriguez, Alejandro Duran C. De Albornoz, and Edson Roberto De Pieri. "An infotaxis based odor navigation approach." en. In: *ISSNIP Biosignals and Biorobotics Conference 2011*. Vitoria, Brazil: IEEE, 2011, pp. 1–6. DOI: 10.1109/BRC.2011.5740683.

[Ras+10] Fabio Rastrello, Pisana Placidi, Luca Abbati, Andrea Scorzoni, Enrico Cozzani, Ivan Elmi, Stefano Zampolli, Gian Carlo Cardinali, et al. "Thermal transient measurements of an ultra-low-power MOX sensor." In: *Journal of sensors* 2010 (2010). Publisher: Hindawi.

[Ras+23] Shavika Rastogi, Nik Dennler, Michael Schmuker, and André Van Schaik. "Spike-time encoding of gas concentrations using neuromorphic analog sensory front-end." en. In: *2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. Toronto, ON, Canada: IEEE, 2023, pp. 1–5. DOI: 10.1109/BioCAS58349.2023.10388752.

[RBW08] Frank Röck, Nicolae Barsan, and Udo Weimar. "Electronic nose: current status and future trends." In: *Chemical reviews* 108.2 (2008). Publisher: ACS Publications, pp. 705–725.

[RDR10] Juan Rodríguez, Cristhian Durán, and Adriana Reyes. "Electronic Nose for Quality Control of Colombian Coffee through the Detection of Defects in "Cup Tests"." In: *Sensors* 10.1 (2010), pp. 36–46. DOI: 10.3390/s100100036.

[Reb+14] Michelle R. Rebello, Thomas S. McTavish, David C. Willhite, Shaina M. Short, Gordon M. Shepherd, and Justus V. Verhagen. "Perception of Odors Linked to Precise Timing in the Olfactory System." en. In: *Plos biology* 12.12 (2014). Ed. by Leslie Vosshall, e1002021. DOI: 10.1371/journal.pbio.1002021.

[Ren+19] He Ren, Yanling Zhao, Wu Xiao, and Zhenqi Hu. "A review of UAV monitoring in mining areas: current status and future perspectives." en. In: *International journal of coal science & technology* 6.3 (2019), pp. 320–333. DOI: 10.1007/s40789-019-00264-5.

[Rey15] Jean-Louis Reymond. "The Chemical Space Project." en. In: *Accounts of chemical research* 48.3 (2015), pp. 722–730. DOI: 10.1021/ar500432k.

[RGS17] Georg Raiser, C. Giovanni Galizia, and Paul Szyszka. "A High-Bandwidth Dual-Channel Olfactory Stimulator for Studying Temporal Sensitivity of Olfactory Processing." en. In: *Chemical senses* 42.2 (2017), pp. 141–151. DOI: 10.1093/chemse/bjw114.

[RHB18] Daniel Rüffer, Felix Hoehne, and Johannes Bühler. "New Digital Metal-Oxide (MOx) Sensor Platform." en. In: *Sensors* 18.4 (2018), p. 1052. DOI: 10.3390/s18041052.

[Rif+14] Jeffrey A. Riffell, Eli Shlizerman, Elischa Sanders, Leif Abrell, Billie Medina, Armin J. Hinterwirth, and J. Nathan Kutz. "Flower discrimination by pollinators in a dynamic chemical environment." en. In: *Science* 344.6191 (2014), pp. 1515–1518. DOI: 10.1126/science.1251041.

[Rig+22a] Nicola Rigolli, Nicodemo Magnoli, Lorenzo Rosasco, and Agnese Seminara. "Learning to predict target location with turbulent odor plumes." en. In: *Elife* 11 (2022), e72196. DOI: 10.7554/eLife.72196.

[Rig+22b] Nicola Rigolli, Gautam Reddy, Agnese Seminara, and Massimo Vergassola. "Alternation emerges as a multi-modal strategy for turbulent odor navigation." en. In: *Elife* 11 (2022), e76989. DOI: 10.7554/eLife.76989.

[RMV22] Gautam Reddy, Venkatesh N. Murthy, and Massimo Vergassola. "Olfactory Sensing and Navigation in Turbulent Environments." en. In: *Annual review of condensed matter physics* 13.1 (2022), pp. 191–213. DOI: 10.1146/annurev-conmatphys-031720-032754.

[Rok+14] Dan Rokni, Vivian Hemmelder, Vikrant Kapoor, and Venkatesh N Murthy. "An olfactory cocktail party: figure-ground segregation of odorants in rodents." en. In: *Nature neuroscience* 17.9 (2014), pp. 1225–1232. DOI: 10.1038/nn.3775.

[Rud+12] Lars Ruddigkeit, Ruud Van Deursen, Lorenz C. Blum, and Jean-Louis Reymond. "Enumeration of 166 Billion Organic Small Molecules in the Chemical Universe Database GDB-17." en. In: *Journal of chemical information and modeling* 52.11 (2012), pp. 2864–2875. DOI: 10.1021/ci300415d.

[Rya+04] M.A. Ryan, A.V. Shevade, H. Zhou, and M.L. Homer. "Polymer–Carbon Black Composite Sensors in an Electronic Nose for Air-Quality Monitoring." en. In: *Mrs bulletin* 29.10 (2004), pp. 714–719. DOI: 10.1557/mrs2004.208.

[Sah+13] Debajit Saha, Kevin Leong, Chao Li, Steven Peterson, Gregory Siegel, and Baranidharan Raman. "A spatiotemporal coding mechanism for background-invariant odor recognition." en. In: *Nature neuroscience* 16.12 (2013), pp. 1830–1839. DOI: 10.1038/nn.3570.

[Sar+13] Hamdi Melih Saraoglu, Ali Osman Selvi, M. Ali Ebeoglu, and Cihat Tasaltin. "Electronic Nose System Based on Quartz Crystal Microbalance Sensor for Blood Glucose and HbA1c Levels From Exhaled Breath Odor." en. In: *Ieee sensors journal* 13.11 (2013), pp. 4229–4235. DOI: 10.1109/JSEN.2013.2265233.

[SBH16] Michael Schmuker, Viktor Bahr, and Ramón Huerta. "Exploiting plume structure to decode gas source distance using metal-oxide gas sensors." In: *Sensors and actuators, b: chemical* 235 (2016), pp. 636–646. DOI: 10.1016/j.snb.2016.05.098.

[Sch+08] Julia Schuckel, Shannon Meisner, Päivi H. Torkkeli, and Andrew S. French. "Dynamic properties of Drosophila olfactory electroantennograms." en. In: *Journal of comparative physiology a* 194.5 (2008), pp. 483–489. DOI: 10.1007/s00359-008-0322-6.

[Sch+11] Michael Schmuker, Nobuhiro Yamagata, Martin Paul Nawrot, and Randolf Menzel. "Parallel Representation of Stimulus Identity and Intensity in a Dual Pathway Model Inspired by the Olfactory System of the Honeybee." en. In: *Frontiers in neuroengineering* 4 (2011). DOI: 10.3389/fneng.2011.00017.

[Sch+22] Catherine D. Schuman, Shruti R. Kulkarni, Maryam Parsa, J. Parker Mitchell, Prasanna Date, and Bill Kay. "Opportunities for neuromorphic computing algorithms and applications." en. In: *Nature computational science* 2.1 (2022), pp. 10–19. DOI: 10.1038/s43588-021-00184-y.

[SEE23] Paul Szyszka, Thierry Emonet, and Timothy L Edwards. "Extracting spatial information from temporal odor patterns: insights from insects." en. In: *Current opinion in insect science* 59 (2023), p. 101082. DOI: 10.1016/j.cois.2023.101082.

[Seh+19] Aarti Sehdev, Yunusa G. Mohammed, Tilman Triphan, and Paul Szyszka. "Olfactory Object Recognition Based on Fine-Scale Stimulus Timing in Drosophila." en. In: *Iscience* 13 (2019), pp. 113–124. DOI: 10.1016/j.isci.2019.02.014.

[SGS13] Jacob S. Stierle, C. Giovanni Galizia, and Paul Szyszka. "Millisecond Stimulus Onset-Asynchrony Enhances Information about Components in an Odor Mixture." en. In: *The journal of neuroscience* 33.14 (2013), pp. 6060–6069. DOI: 10.1523/JNEUROSCI.5838-12.2013.

[SH04] Hong T Sun and Peter C Hsi. *Photo-ionization detector and method for continuous operation and real-time self-cleaning*. US Patent 6,734,435. 2004.

[Sha+19] Alexey Shaposhnik, Pavel Moskalev, Elena Sizask, Stanislav Ryabtsev, and Alexey Vasiliev. "Selective Detection of Hydrogen Sulfide and Methane by a Single MOX-Sensor." en. In: *Sensors* 19.5 (2019), p. 1135. DOI: 10.3390/s19051135.

[She04] Gordon M Shepherd. "The Human Sense of Smell: Are We Better Than We Think?" en. In: *Plos biology* 2.5 (2004), e146. DOI: 10.1371/journal.pbio.0020146.

[SHS18] Tizian Schneider, Nikolai Helwig, and Andreas Schütze. "Automatic feature extraction and selection for condition monitoring and related datasets." In: *2018 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*. 2018, pp. 1–6. DOI: 10.1109/I2MTC.2018.8409763.

[Shu+11] Roman Shusterman, Matthew C Smear, Alexei A Koulakov, and Dmitry Rinberg. "Precise olfactory responses tile the sniff cycle." en. In: *Nature neuroscience* 14.8 (2011), pp. 1039–1044. DOI: 10.1038/nn.2877.

[SJR86] Joseph R Stetter, Peter C Jurs, and Susan L Rose. "Detection of hazardous gases and vapors: pattern recognition analysis of data from an electrochemical sensor array." In: *Analytical chemistry* 58.4 (1986). Publisher: ACS Publications, pp. 860–866.

[SL05] André van Schaik and Shih-Chii Liu. "Aer ear: a matched silicon cochlea pair with address event representation interface." In: *2005 ieee international symposium on circuits and systems*. IEEE. 2005, pp. 4213–4216.

[Sme+11] Matthew Smear, Roman Shusterman, Rodney O'Connor, Thomas Bozza, and Dmitry Rinberg. "Perception of sniff phase in mouse olfaction." en. In: *Nature* 479.7373 (2011), pp. 397–400. DOI: 10.1038/nature10521.

[SMK23] Sergey V. Stasenko, Alexey N. Mikhaylov, and Victor B. Kazantsev. "Model of Neuromorphic Odorant-Recognition Network." en. In: *Biomimetics* 8.3 (2023), p. 277. DOI: 10.3390/biomimetics8030277.

[SO01] Eero P Simoncelli and Bruno A Olshausen. "Natural Image Statistics and Neural Representation." en. In: *Annual review of neuroscience* 24.1 (2001), pp. 1193–1216. DOI: 10.1146/annurev.neuro.24.1.1193.

[Sol+22] Ana Solórzano, Jens Eichmann, Luis Fernández, Bernd Ziems, Juan Manuel Jiménez-Soto, Santiago Marco, and Jordi Fonollosa. "Early fire detection based on gas sensor arrays: Multivariate calibration and validation." In: *Sensors and actuators b: chemical* 352 (2022), p. 130961. DOI: https://doi.org/10.1016/j.snb.2021.130961.

[SPN14] Michael Schmuker, Thomas Pfeil, and Martin Paul Nawrot. "A neuromorphic network for generic multivariate data classification." en. In: *Proceedings of the national academy of sciences* 111.6 (2014), pp. 2081–2086. DOI: 10.1073/pnas.1303053111.

[SPP06] Emmanuel Scorsone, Anna Maria Pisanelli, and Krishna C. Persaud. "Development of an electronic nose for fire detection." en. In: *Sensors and actuators b: chemical* 116.1-2 (2006), pp. 55–61. DOI: 10.1016/j.snb.2005.12.059.

[SS00] Boris I. Shraiman and Eric D. Siggia. "Scalar turbulence." en. In: *Nature* 405.6787 (2000), pp. 639–646. DOI: 10.1038/35015000.

[Sta14] R.A.E.S.H. Staff. *The PID Handbook: Theory and Applications of Direct-Reading Photoionization Detectors*. RAE Systems by Honeywell, 2014.

[Ste+05] Marcus C Stensmyr, Susanne Erland, Eric Hallberg, Rita Wallén, Peter Greenaway, and Bill S Hansson. "Insect-like olfactory adaptations in the terrestrial giant robber crab." In: *Current biology* 15.2 (2005). Publisher: Elsevier, pp. 116–121.

[Suh+04] Greg S. B. Suh, Allan M. Wong, Anne C. Hergarden, Jing W. Wang, Anne F. Simon, Seymour Benzer, Richard Axel, and David J. Anderson. "A single population of olfactory sensory neurons mediates an innate avoidance behaviour in Drosophila." en. In: *Nature* 431.7010 (2004), pp. 854–859. DOI: 10.1038/nature02980.

[Sul+15] Regina M. Sullivan, Donald A. Wilson, Nadine Ravel, and Anne-Marie Mouly. "Olfactory memory networks: from emotional learning to social behaviors." en. In: *Frontiers in behavioral neuroscience* 9 (2015). DOI: 10.3389/fnbeh.2015.00036.

[Szy+12] Paul Szyszka, Jacob S. Stierle, Stephanie Biergans, and C. Giovanni Galizia. "The Speed of Smell: Odor-Object Segregation within Milliseconds." en. In: *Plos one* 7.4 (2012). Ed. by Matthieu Louis, e36096. DOI: 10.1371/journal.pone.0036096.

[Szy+14] Paul Szyszka, Richard C. Gerkin, C. Giovanni Galizia, and Brian H. Smith. "High-speed odor transduction and pulse tracking by insect olfactory receptor neurons." en. In: *Proceedings of the national academy of sciences* 111.47 (2014), pp. 16925–16930. DOI: 10.1073/pnas.1412051111.

[Tai+23] István Taisz et al. "Generating parallel representations of position and identity in the olfactory system." en. In: *Cell* 186.12 (2023), 2556–2573.e22. DOI: 10.1016/j.cell.2023.04.038.

[Tak+14] Yoshinori Takei, Yuhei Shimizu, Kazuki Hirasawa, and Hidehito Nanto. "Braitenberg's vehicle-like odor plume tracking robot." en. In: *IEEE SENSORS 2014 Proceedings*. Valencia, Spain: IEEE, 2014, pp. 1276–1279. DOI: 10.1109/ICSENS.2014.6985243.

[Tan+09] ZhongLin Tang, JianHua Yang, Ying Xu, and JunYun Yu. "Towards the Development of a Portable Smell-Seeing Electronic Nose and its Applications in Amine Recognition." en. In: *2009 2nd International Congress on Image and Signal Processing*. Tianjin: IEEE, 2009, pp. 1–5. DOI: 10.1109/CISP.2009.5304113.

[Tan+19] Gouhei Tanaka, Toshiyuki Yamane, Jean Benoit Héroux, Ryosho Nakane, Naoki Kanazawa, Seiji Takeda, Hidetoshi Numata, Daiju Nakano, and Akira Hirose. "Recent advances in physical reservoir computing: A review." en. In: *Neural networks* 115 (2019), pp. 100–123. DOI: 10.1016/j.neunet.2019.03.005.

[Tar+21] Mohammad F. Tariq, Suzanne M. Lewis, Aliena Lowell, Sidney Moore, Jesse T. Miles, David J. Perkel, and David H. Gire. "Using Head-Mounted Ethanol Sensors to Monitor Olfactory Information and Determine Behavioral Changes Associated with Ethanol-Plume Contact during Mouse Odor-Guided Navigation." en. In: *Eneuro* 8.1 (2021), ENEURO.0285–20.2020. DOI: 10.1523/ENEURO.0285-20.2020.

[TC22a] Aaron C True and John P Crimaldi. “Distortion of passive scalar structure during suction-based plume sampling.” In: *Sensors and actuators b: chemical* 367 (2022), p. 132018.

[TC22b] Aaron C. True and John P. Crimaldi. “Distortion of passive scalar structure during suction-based plume sampling.” en. In: *Sensors and actuators b: chemical* 367 (2022), p. 132018. DOI: 10.1016/j.snb.2022.132018.

[TDV01] Simon Thorpe, Arnaud Delorme, and Rufin Van Rullen. “Spike-based strategies for rapid processing.” In: *Neural networks* 14.6 (2001), pp. 715–725. DOI: https://doi.org/10.1016/S0893-6080(01)00083-1.

[Ter+24] Mykhail Tereshkov, Tetiana Dontsova, Bilge Saruhan, and Svitlana Krüger. “Metal Oxide-Based Sensors for Ecological Monitoring: Progress and Perspectives.” en. In: *Chemosensors* 12.3 (2024), p. 42. DOI: 10.3390/chemosensors12030042.

[TGK20] Julio Torres-Tello, Ana V. Guaman, and Seok-Bum Ko. “Improving the Detection of Explosives in a MOX Chemical Sensors Array With LSTM Networks.” In: *Ieee sensors journal* 20.23 (2020), pp. 14302–14309. DOI: 10.1109/JSEN.2020.3007431.

[The+12] Treenet Thepudom, Sumana Kladsomboon, Tawee Pogfay, Adisorn Tuantranont, and Teerakiat Kerdcharoen. “Portable optical-based electronic nose using dual-sensors array applied for volatile discrimination.” en. In: *2012 9th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*. Phetchaburi, Thailand: IEEE, 2012, pp. 1–4. DOI: 10.1109/ECTICon.2012.6254208.

[UM03] Naoshige Uchida and Zachary F Mainen. “Speed and accuracy of olfactory discrimination in the rat.” en. In: *Nature neuroscience* 6.11 (2003), pp. 1224–1229. DOI: 10.1038/nn1142.

[Van+15] Floris Van Breugel, Jeff Riffell, Adrienne Fairhall, and Michael H Dickinson. “Mosquitoes use vision to associate odor plumes with thermal targets.” In: *Current biology* 25.16 (2015). Publisher: Elsevier, pp. 2123–2129.

[Van+19] Anup Vanarse, Adam Osseiran, Alexander Rassau, and Peter Van Der Made. “A Hardware-Deployable Neuromorphic Solution for Encoding and Classification of Electronic Nose Data.” en. In: *Sensors* 19.22 (2019), p. 4831. DOI: 10.3390/s19224831.

[Van+20] Anup Vanarse, Josafath Israel Espinosa-Ramos, Adam Osseiran, Alexander Rassau, and Nikola Kasabov. "Application of a Brain-Inspired Spiking Neural Network Architecture to Odor Data Classification." en. In: *Sensors* 20.10 (2020), p. 2756. DOI: 10.3390/s20102756.

[Van+22] Anup Vanarse, Adam Osseiran, Alexander Rassau, and Peter Van Der Made. "Application of Neuromorphic Olfactory Approach for High-Accuracy Classification of Malts." en. In: *Sensors* 22.2 (2022), p. 440. DOI: 10.3390/s22020440.

[Ver+] Alexander Vergara, Jordi Fonollosa, Jonas Mahiques, Marco Trincavelli, Nikolai Rulkov, and Ramón Huerta. *Gas sensor arrays in open sampling settings data set*. http://archive.ics.uci.edu/ml/datasets/Gas+sensor+arrays+in+open+sampling+settings. Accessed: 2021-08-09.

[Ver+12] Alexander Vergara, Shankar Vembu, Tuba Ayhan, Margaret A. Ryan, Margie L. Homer, and Ramón Huerta. "Chemical gas sensor drift compensation using classifier ensembles." en. In: *Sensors and actuators b: chemical* 166-167 (2012), pp. 320–329. DOI: 10.1016/j.snb.2012.01.074.

[Ver+13] Alexander Vergara, Jordi Fonollosa, Jonas Mahiques, Marco Trincavelli, Nikolai Rulkov, and Ramón Huerta. "On the performance of gas sensor arrays in open sampling systems using Inhibitory Support Vector Machines." en. In: *Sensors and actuators b: chemical* 185 (2013), pp. 462–477. DOI: 10.1016/j.snb.2013.05.027.

[Ver+14] Alexander Vergara, Kurt D. Benkstein, Christopher B. Montgomery, and Steve Semancik. "Demonstration of Fast and Accurate Discrimination and Quantification of Chemically Similar Species Utilizing a Single Cross-Selective Chemiresistor." en. In: *Analytical chemistry* 86.14 (2014), pp. 6753–6757. DOI: 10.1021/ac501490k.

[Ver18] Nico Vervliet. "Compressed sensing approaches to large-scale tensor decompositions." PhD Thesis. KU Leuven, 2018.

[VGT05] Rufin VanRullen, Rudy Guyonneau, and Simon J Thorpe. "Spike times make sense." In: *Trends in neurosciences* 28.1 (2005), pp. 1–4.

[VI99] Emmanuel Villermaux and Claudia Innocenti. "On the geometry of turbulent mixing." en. In: *Journal of fluid mechanics* 393 (1999), pp. 123–147. DOI: 10.1017/S0022112099005674.

[Vic+01] Neil J Vickers, Thomas A Christensen, Thomas C Baker, and John G Hildebrand. "Odour-plume dynamics infuence the brain's olfactory code." en. In: *Nature* 410 (2001).

[VK07] Volodymyr Vasyutynskyy and Klaus Kabitzsch. "Towards comparison of deadband sampling types." In: *2007 ieee international symposium on industrial electronics*. 2007, pp. 2899–2904. DOI: 10.1109/ISIE.2007.4375074.

[VL16] Nico Vervliet and Lieven De Lathauwer. "A Randomized Block Sampling Approach to Canonical Polyadic Decomposition of Large-Scale Tensors." In: *Ieee journal of selected topics in signal processing* 10.2 (2016). Publisher: Institute of Electrical and Electronics Engineers (IEEE), pp. 284–295. DOI: 10.1109/jstsp.2015.2503260.

[Vos00] Leslie B Vosshall. "Olfaction in drosophila." In: *Current opinion in neurobiology* 10.4 (2000). Publisher: Elsevier, pp. 498–503.

[VS07] Leslie B. Vosshall and Reinhard F. Stocker. "Molecular Architecture of Smell and Taste in *Drosophila*." en. In: *Annual review of neuroscience* 30.1 (2007), pp. 505–533. DOI: 10.1146/annurev.neuro.30.051606.094306.

[VTS04] Jean-Philippe Vert, Koji Tsuda, and Bernhard Schölkopf. "A primer on kernel methods." In: *Kernel methods in computational biology* 47 (2004). Publisher: MIT press Cambridge, MA, USA, pp. 35–70.

[VVS07] Massimo Vergassola, Emmanuel Villermaux, and Boris I. Shraiman. "'Infotaxis' as a strategy for searching without gradients." en. In: *Nature* 445.7126 (2007), pp. 406–409. DOI: 10.1038/nature05464.

[Wan+01] M. Wandel, U. Weimar, A. Lilienthal, and A. Zell. "Leakage localisation with a mobile robot carrying chemical sensors." en. In: *ICECS 2001. 8th IEEE International Conference on Electronics, Circuits and Systems (Cat. No.01EX483)*. Malta: IEEE, 2001, pp. 1247–1250. DOI: 10.1109/ICECS.2001.957441.

[Wan+10] Chengxiang Wang, Longwei Yin, Luyuan Zhang, Dong Xiang, and Rui Gao. "Metal Oxide Gas Sensors: Sensitivity and Influencing Factors." en. In: *Sensors* 10.3 (2010), pp. 2088–2106. DOI: 10.3390/s100302088.

[Wan+11] L.C. Wang, K.T. Tang, S.W. Chiu, S.R. Yang, and C.T. Kuo. "A bio-inspired two-layer multiple-walled carbon nanotube–polymer composite sensor array and a bio-inspired fast-adaptive readout circuit for a portable electronic nose." en. In: *Biosensors and bioelectronics* 26.11 (2011), pp. 4301–4307. DOI: 10.1016/j.bios.2011.04.015.

[Wan+21a] Syuan He Wang, Ting I. Chou, Shih Wen Chiu, and Kea Tiong Tang. "Using a Hybrid Deep Neural Network for Gas Classification." In: *Ieee sensors journal* 21.5 (2021), pp. 6401–6407. DOI: 10.1109/JSEN.2020.3038304.

[Wan+21b] Tong Wang, He-Ming Huang, Xiao-Xue Wang, and Xin Guo. "An artificial olfactory inference system based on memristive devices." en. In: *Infomat* 3.7 (2021), pp. 804–813. DOI: 10.1002/inf2.12196.

[Wan+22a] Junhan Wang, Yuezhang Lin, Ruirui Liu, and Jun Fu. "Odor source localization of multi-robots with swarm intelligence algorithms: A review." en. In: *Frontiers in neurorobotics* 16 (2022), p. 949888. DOI: 10.3389/fnbot.2022.949888.

[Wan+22b] Tong Wang, Xiao-Xue Wang, Juan Wen, Zhe-Yuan Shao, He-Ming Huang, and Xin Guo. "A Bio-Inspired Neuromorphic Sensory System." en. In: *Advanced intelligent systems* 4.7 (2022), p. 2200047. DOI: 10.1002/aisy.202200047.

[Wan+23] Lingxiao Wang, Shuo Pang, Mantasha Noyela, Kevin Adkins, Lulu Sun, and Marwa El-Sayed. "Vision and Olfactory-Based Wildfire Monitoring with Uncrewed Aircraft Systems." en. In: *2023 20th International Conference on Ubiquitous Robots (UR)*. Honolulu, HI, USA: IEEE, 2023, pp. 716–723. DOI: 10.1109/UR57808.2023.10202419.

[Wan+24] Chen Wang et al. "Biomimetic olfactory chips based on large-scale monolithically integrated nanotube sensor arrays." In: *Nature electronics* (2024). Publisher: Nature Publishing Group UK London, pp. 1–11.

[WB+95] Greg Welch, Gary Bishop, et al. "An introduction to the Kalman filter." In: (1995). Publisher: Chapel Hill, NC, USA.

[Web+18] Jason Webster, Pratistha Shakya, Eamonn Kennedy, Michael Caplan, Christopher Rose, and Jacob K. Rosenstein. "TruffleBot: Low-Cost Multi-Parametric Machine Olfaction." en. In: *2018 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. Cleveland, OH: IEEE, 2018, pp. 1–4. DOI: 10.1109/BIOCAS.2018.8584767.

[Wee+20] Jonathan S. Weerakkody, Sophie Brenet, Thierry Livache, Cyril Herrier, Yanxia Hou, and Arnaud Buhot. "Optical Index Prism Sensitivity of Surface Plasmon Resonance Imaging in Gas Phase: Experiment versus Theory." en. In: *The journal of physical chemistry c* 124.6 (2020), pp. 3756–3767. DOI: 10.1021/acs.jpcc.9b09973.

[Wei+02] M J Weissburg, D B Dusenbery, H Ishida, J Janata, P J W Roberts, and D R Webster. "A Multidisciplinary Study of Spatial and Temporal Scales Containing Information in Turbulent Chemical Plume Tracking." en. In: *Environmental fluid mechanics* 2 (2002), pp. 65–94.

[Wie+22] Thomas Wiedemann, Marius Schaab, Juan Marchal Gomez, Dmitriy Shutin, Monika Scheibe, and Achim J. Lilienthal. "Gas Source Localization Based on Binary Sensing with a UAV." en. In: *2022 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*. Aveiro, Portugal: IEEE, 2022, pp. 1–3. DOI: 10.1109/ISOEN54820.2022.9789553.

[Wil13a] Alphus Wilson. "Diverse Applications of Electronic-Nose Technologies in Agriculture and Forestry." en. In: *Sensors* 13.2 (2013), pp. 2295–2348. DOI: 10.3390/s130202295.

[Wil13b] Rachel I. Wilson. "Early Olfactory Processing in *Drosophila* : Mechanisms and Principles." en. In: *Annual review of neuroscience* 36.1 (2013), pp. 217–241. DOI: 10.1146/annurev-neuro-062111-150533.

[Wil18] Alphus Wilson. "Application of Electronic-Nose Technologies and VOC-Biomarkers for the Noninvasive Early Diagnosis of Gastrointestinal Diseases." en. In: *Sensors* 18.8 (2018), p. 2613. DOI: 10.3390/s18082613.

[Wu+19] Chunxue Wu, Bobo Ju, Yan Wu, Xiao Lin, Naixue Xiong, Guangquan Xu, Hongyan Li, and Xuefeng Liang. "UAV Autonomous Target Search Based on Deep Reinforcement Learning in Complex Disaster Scene." en. In: *Ieee access* 7 (2019), pp. 117227–117245. DOI: 10.1109/ACCESS.2019.2933002.

[Xin+19] Yuxin Xing, Timothy Vincent, Marina Cole, and Julian Gardner. "Real-Time Thermal Modulation of High Bandwidth MOX Gas Sensors for Mobile Robot Applications." en. In: *Sensors* 19.5 (2019), p. 1180. DOI: 10.3390/s19051180.

[Yam+12] Jaber Al Yamani, Farid Boussaid, Amine Bermak, and Dominique Martinez. "Glomerular Latency Coding in Artificial Olfaction." en. In: *Frontiers in neuroengineering* 4 (2012). DOI: 10.3389/fneng.2011.00018.

[Yan+23] Jia Yan, Biao Wu, Tao Liu, Feiyue Chen, and Shukai Dua. "An SNN-Based Bionic Olfactory Signal Processing Network for Odor Recognition." en. In: *Ieee sensors journal* 23.12 (2023), pp. 13186–13197. DOI: 10.1109/JSEN.2023.3270024.

[Yee+95a] Eugene Yee, R. Chan, P. R. Kosteniuk, G. M. Chandler, C. A. Biltoft, and J. F. Bowers. "Measurements of level-crossing statistics of concentration fluctuations in plumes dispersing in the atmospheric surface layer." In: *Boundary-layer meteorology* 73.1-2 (1995), pp. 53–90. DOI: 10.1007/BF00708930.

[Yee+95b] Eugene Yee, R. Chan, P. R. Kosteniuk, G. M. Chandler, C. A. Biltoft, and J. F. Bowers. "The vertical structure of concentration fluctuation statistics in plumes dispersing in the atmospheric surface layer." en. In: *Boundary-layer meteorology* 76.1-2 (1995), pp. 41–67. DOI: 10.1007/BF00710890.

[YJM20] Zi-Zhen Yang, Tao Jing, and Qing-Hao Meng. "UAV-based Odor Source Localization in Multi-Building Environments Using Simulated Annealing Algorithm." en. In: *2020 39th Chinese Control Conference (CCC)*. Shenyang, China: IEEE, 2020, pp. 3806–3811. DOI: 10.23919/CCC50068.2020.9189425.

[Zho+19a] Peng Zhou, Yi-Dong Shen, Liang Du, and Fan Ye. "Incremental Multi-view Support Vector Machine." In: *Proceedings of the 2019 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, 2019, pp. 1–9. DOI: 10.1137/1.9781611975673.1.

[Zho+19b] Peng Zhou, Yi-Dong Shen, Liang Du, Fan Ye, and Xuejun Li. "Incremental multi-view spectral clustering." In: *Knowledge-based systems* 174 (2019). Publisher: Elsevier BV, pp. 73–86. DOI: 10.1016/j.knosys.2019.02.036.

[Ziy+10] A. Ziyatdinov, S. Marco, A. Chaudry, K. Persaud, P. Caminal, and A. Perera. "Drift compensation of gas sensor array data by common principal component analysis." en. In: *Sensors and actuators b: chemical* 146.2 (2010), pp. 460–465. DOI: 10.1016/j.snb.2009.11.034.

[Ziy+15] Andrey Ziyatdinov, Jordi Fonollosa, Luis Fernández, Agustín Gutierrez-Gálvez, Santiago Marco, and Alexandre Perera. "Bioinspired early detection through gas flow modulation in chemo-sensory systems." en. In: *Sensors and actuators b: chemical* 206 (2015), pp. 538–547. DOI: 10.1016/j.snb.2014.09.001.

[ZLC15] Jie-yong Zhou, Ji-gong Li, and Shi-gang Cui. "A bionic plume tracing method with a mobile robot in outdoor time-varying airflow environment." en. In: *2015 IEEE International Conference on Information and Automation*. Lijiang, China: IEEE, 2015, pp. 2351–2355. DOI: 10.1109/ICInfA.2015.7279679.

[Zul+20] Claudio Zuliani, Joaquim Luque, Claudio Falco, Ethan Gardner, Andrea De Luca, Tim Vincent, Sanjeeb Tripathy, Zeeshan Ali, and Florin Udrea. "Flow Compensated Gas Sensing Array for Improved Performances in Breath-Analysis Applications." en. In: *Ieee sensors letters* 4.3 (2020), pp. 1–4. DOI: 10.1109/LSENS.2020.2974200.

[ZZ23] Zhengxin Zhang and Lixue Zhu. "A Review on Unmanned Aerial Vehicle Remote Sensing: Platforms, Sensors, Data Processing Methods, and Applications." en. In: *Drones* 7.6 (2023), p. 398. DOI: 10.3390/drones7060398.