Temporal filtering network (D9.5.1 – SGA2)

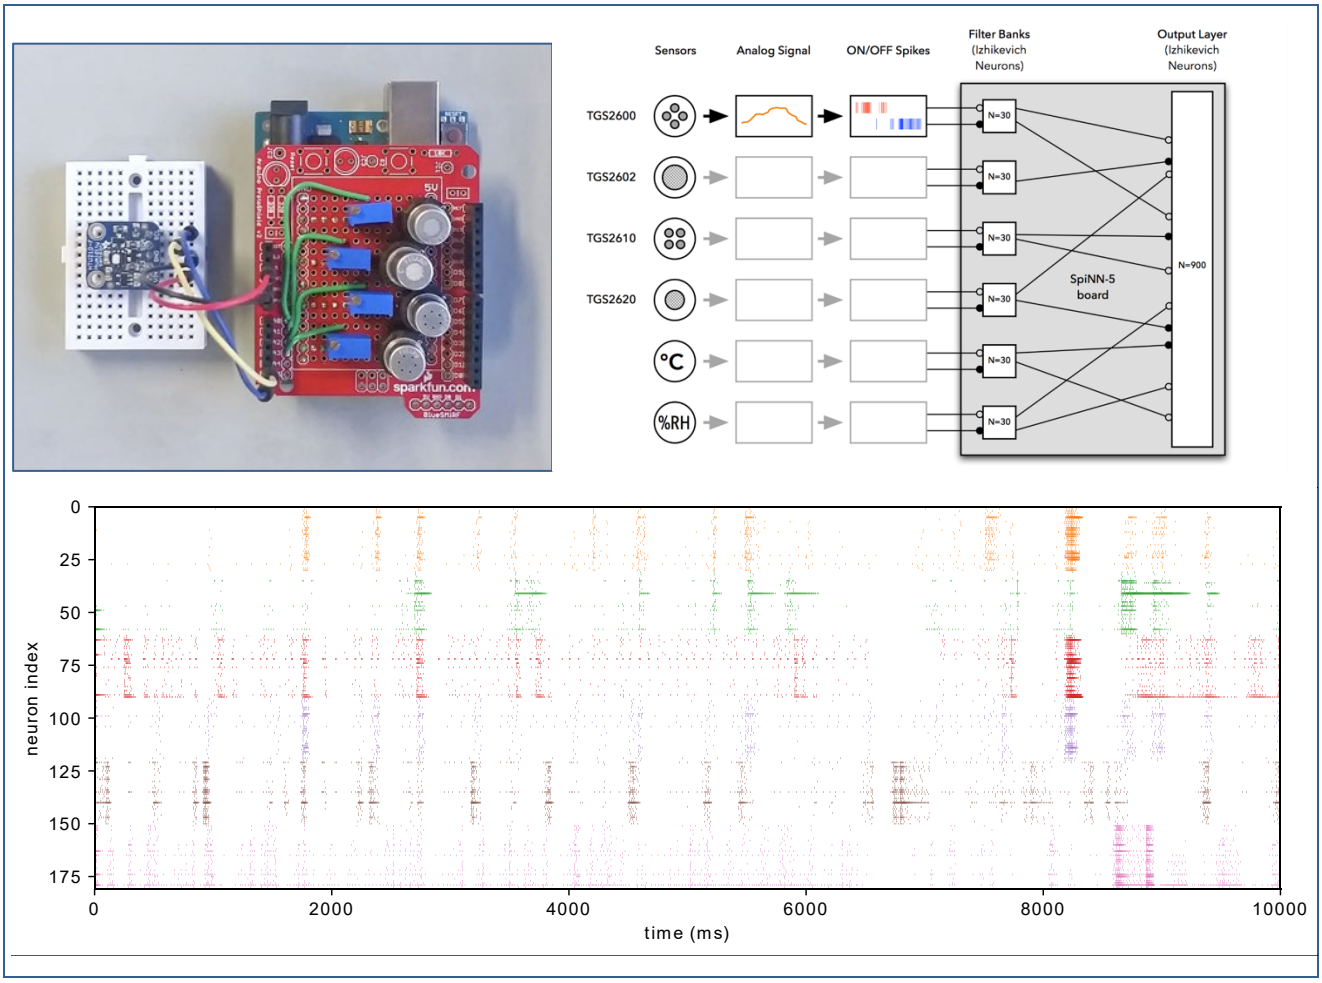

Figure 1: An electronic gas sensor prototype feeds into our neuromorphic olfaction system, that consists of a spiking network for temporal filtering of gas sensor data.

Analogue sensor signals from the sensors are converted to spikes and processed by a spiking network that runs on the SpiNNaker platform.

Project Number:	785907	Project Title:	Human Brain Project SGA2
Document Title:	Temporal filtering network		
Document Filename:	D9.5.1 (D61.1 D93) SGA2 M12 ACCEPTED 190723		
Deliverable Number:	SGA2 D9.5.1 (D61.1, D93)		
Deliverable Type:	Other		
Work Packages:	WP9.5		
Dissemination Level:	PU = Public		
Planned Delivery Date:	SGA2 M12 / 31 Mar 2019		
Actual Delivery Date:	SGA2 M12 / 25 Mar 2019; ACCEPTED 23 Jul 2019		
Author(s):	Damien DRIX, HERTS (P118) Michael SCHMUKER, HERTS (P118)		
Compiled by:	Michael SCHMUKER, HERTS (P118)		
Contributor(s):	Damien DRIX, HERTS (P118), sections 2-5 Michael SCHMUKER, HERTS (P118) section 1		
SciTechCoord Review:	Martin TELEFONT, EPFL (P1)		
Editorial Review:	Guy WILLIS, EPFL (P1)		
Description in GA:	Network for temporal filtering running on SpiNNaker (Task 9.5.1)		
Abstract:	We present a spiking network that is capable of isolating periods of rising sensor signals in a gas sensing system. The system first converts analogue sensor signals into the spiking domain using a temporal difference algorithm, similar to the one used in event-based vision sensors. Temporal filtering is achieved by exploiting neuron-intrinsic dynamics, i.e. bursting and adaptation. A sparsening layer is employed to generate compound features from the input stage. The network is intended to operate as part of a neuromorphic olfaction system that informs gas-based navigation of neurorobotic agents.		
Keywords:	Neuromorphic olfaction, gas sensing, neurorobotic navigation, SpiNNaker.		
Target Users/Readers:	Neuromorphic computing community, electronic nose community, computational neuroscience community, researchers, students, innovators in gas sensing and electronic olfaction.		

Table of Contents

Table of Figures

1. Introduction

Olfactory stimuli have complex temporal profiles. Plumes of volatiles emitted by odour sources are distorted by turbulent processes on their way to the sensory organs. Olfactory neural circuits have evolved to generate stable percepts from rapidly fluctuating intensity profiles encountered in natural environments. Understanding the processing principles of those circuits is a timely question in neuroscience. At the same time, efficient algorithms to rapidly and reliably extract information from turbulent plumes of volatiles are sought after for technical applications involving electronic olfaction, such as environmental monitoring and gas-based robotic navigation.

While the concentration profiles that are produced by point sensors in turbulent gas plumes are the result of a chaotic process, their temporal profiles do contain information about the odour source. We have previously shown in that the distance of a gas source can be decoded from the temporal concentration profile [6]. This method is based on extracting events of rising concentration from the sensor time series, termed “bouts”. Source proximity can be inferred from the number of those bouts detected in a time period (Figure 2).

Here we present a method that uses a network of spiking neurons to perform temporal filtering of gas sensor signals. We use an in-house design of a sensor board that incorporates gas sensing, along with measurements of humidity and temperature. The sensor data is converted to spikes using a bio-realistic algorithm and fed into a purpose-designed network of spiking neurons, running on the SpiNNaker platform, that extracts temporal features from the signal.

Data and code relating to this deliverable are available online [7].

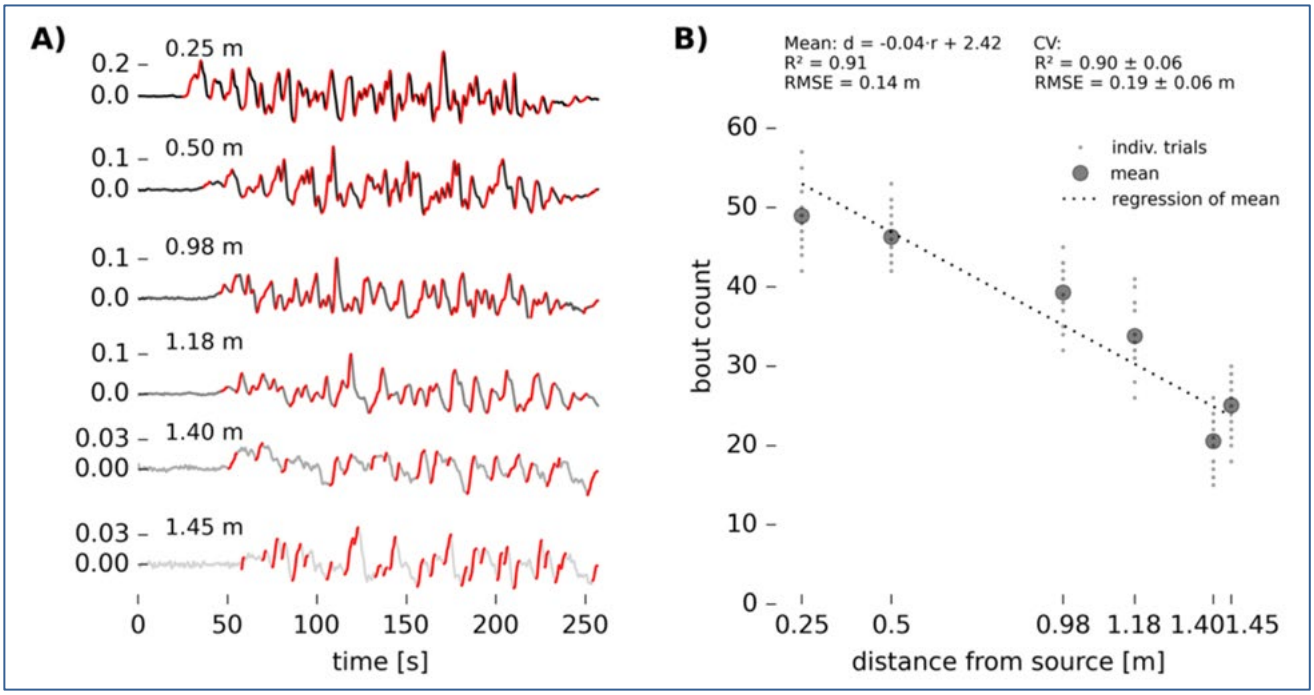

Figure 2: Bout detection and relationship to distance

Reproduced from earlier work [6]. A) Filtered sensor signals showing bouts (positive slopes, in red). B) Source distance is inferred from bout counts.

2. Gas Sensor Board

We designed and built a sensor board consisting of four metal-oxide gas sensors (TGS2600, TGS2602, TGS2610 and TGS2020, Figaro Inc.) together with a digital temperature and humidity sensor (HTU21D-F, TE Connectivity) (Figure 3). The response spectra of the gas sensors covered various compounds including butane, methane, hydrogen, ethanol, ammonia, toluene and carbon monoxide;

they are also highly sensitive to changes in temperature and humidity. Response spectra of different models of sensor overlapped partially, allowing the identity of a sensed gas to be decoded from the combinatorial response patterns, using pattern recognition methods [9].

We used an Arduino Uno microcontroller to measure the voltage drop across a load resistance in series with the sensing elements. Voltage values were sent to a host computer together with timestamps of their acquisition.

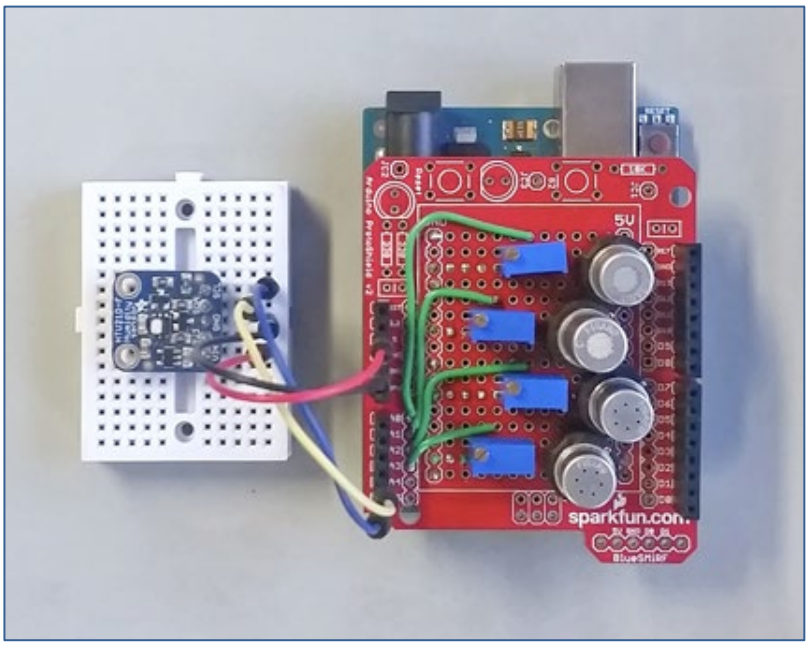

Figure 3: The Gas Sensor Board

Four gas sensors (round metallic protrusions) are visible on the Arduino shield (red), and the temperature/humidity sensor connected via I2C (on the breadboard).

3. Dataset and Signal Conditioning

We recorded a real-world data set for the purpose of evaluating the ability of the recording system and the analysis algorithms to decode different environmental conditions, i.e. various indoor and outdoor scenarios, at varying temperature and humidity levels.

For this purpose, we took our recording rig on a route around the premises of the University of Hertfordshire, that consisted of an indoor segment through the university's corridors and an outdoor segment around the boiler house that provides heating for the complex (Figure 4).

Figure 5 shows the recorded data. Correspondence to specific points along the predefined route is indicated by the numbered time points. Time point “4” denotes the coincidence of a spike in the signal recorded from sensors with a sensitivity for butane with the spatial proximity of the boiler house, around t=600s.

The recorded data was conditioned before further processing with a spiking network. The first step was to remove noise from the acquired signal. The main source of noise in our recordings was quantisation noise generated as an artefact of analogue-to-digital (ADC) conversion. Quantisation noise manifested itself as apparent steps in signal value, caused by the flipping of the least significant bit (LSB). This dominated other sources of noise, due to the fairly low resolution of the Arduino's 10-bit ADC, where the LSB corresponds to a step of 5 mV. We employed a Savitzky-Golay polynomial filter to remove quantisation noise. Future iterations of the sensor board will feature ADCs with higher resolution (24 bit). The expected reduction of quantisation noise will likely cause a stronger emphasis of electronic (stochastic) noise in the signal that can be tackled by a simpler low-pass filter.

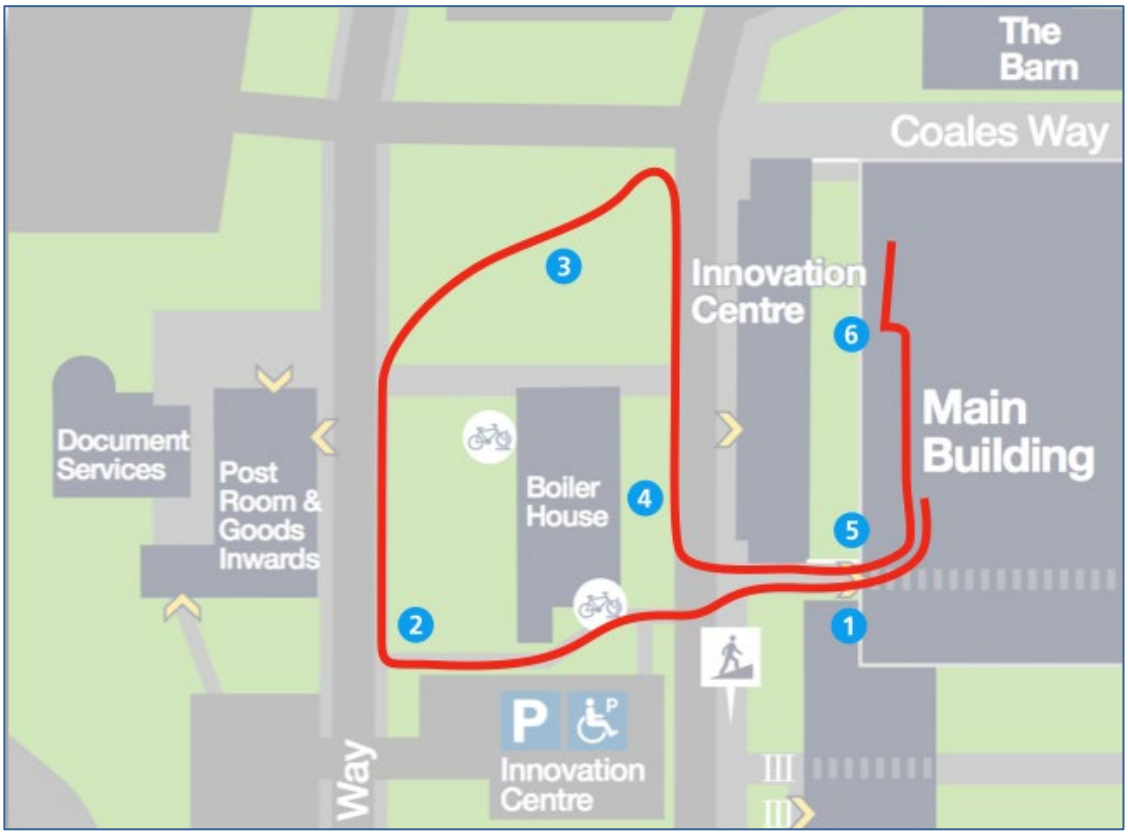

Figure 4: Route taken while recording the dataset.

Numbered locations match those marked in the time series plot (Figure 5).

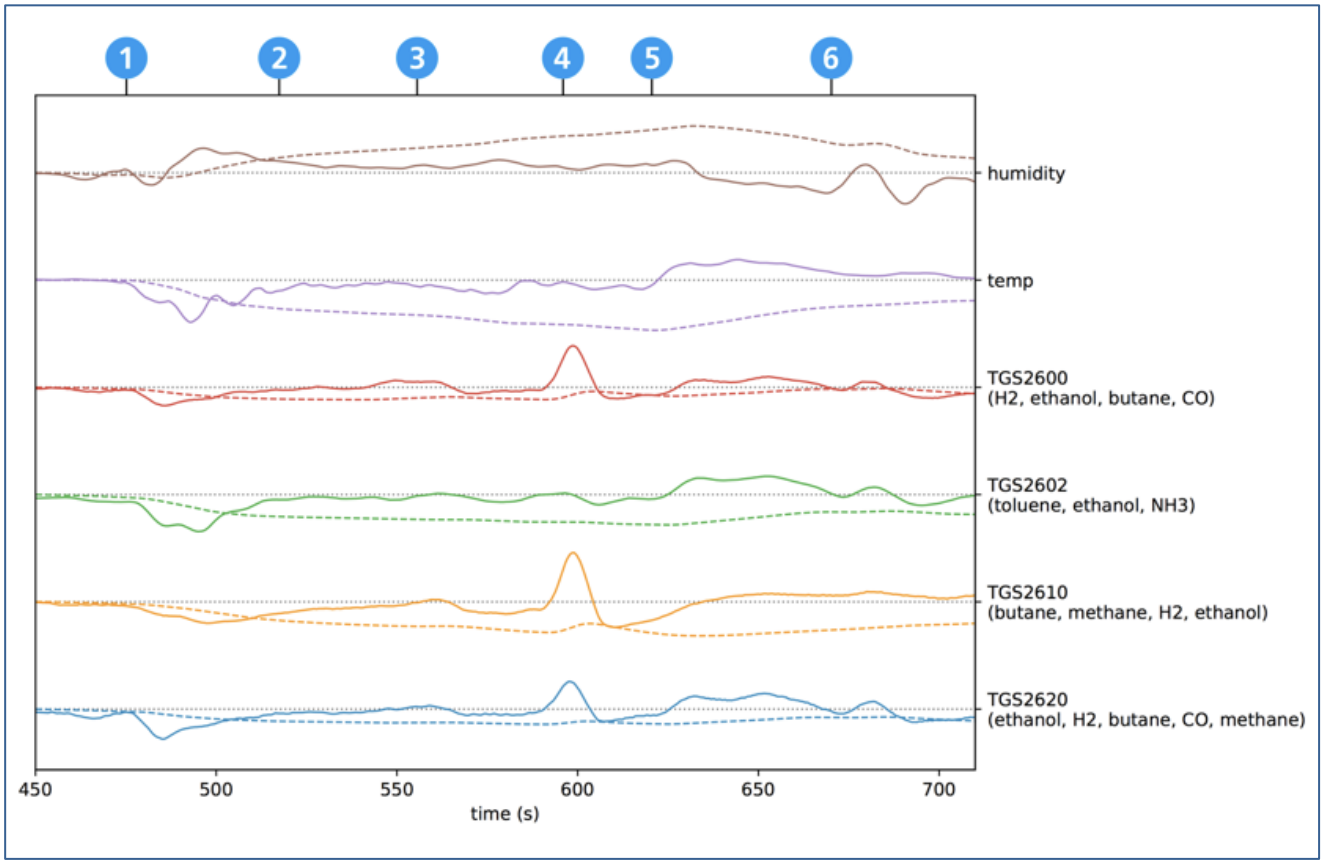

Figure 5: A segment of recorded sensor data.

Dashed lines correspond to the sensor reading (smoothed to remove quantisation artefacts). Solid lines show the derivative of the signal used for further processing.

4. Temporal Filtering Network

The next step was to transform the analogue signals into an event-based representation that we could feed into spiking neurons on a SpiNNaker neuromorphic processor. Our neural network was organised into two layers: a first layer consisting of neural filter banks that extracted temporal features for each sensor, and an output layer that generated sparse responses for various combinations of these features (Figure 6).

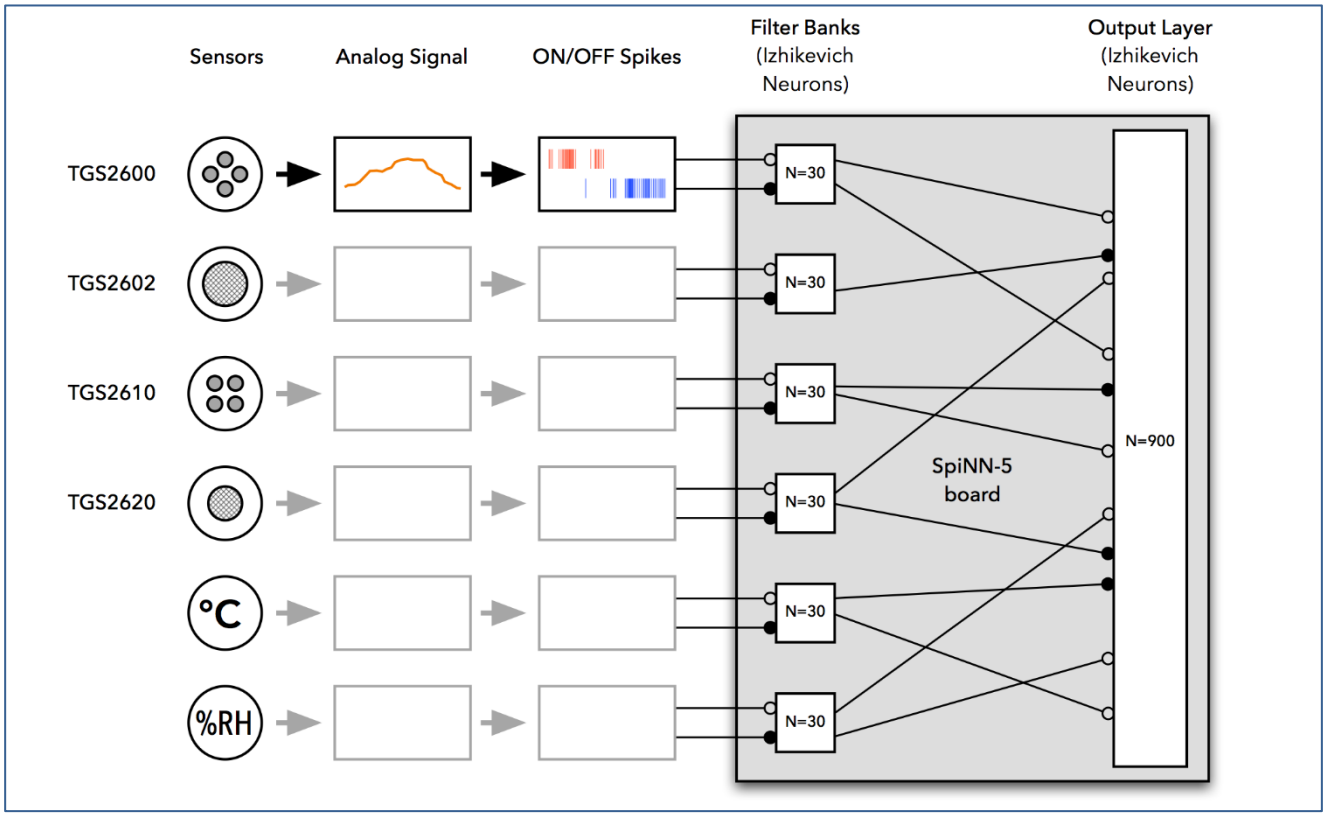

Figure 6: Diagram of the complete temporal filtering system.

4.1 Conversion from analogue signals to spikes

There are various strategies for turning analogue signals into spikes, for instance using the amplitude of the signal to control the rate of a Poisson process. Here, we used another method which involved emitting an event whenever the signal value changes by more than a certain threshold; a strategy similar to the one employed in Delta modulators and in the DVS family of neuromorphic cameras. This yielded a stream of positive (ON) and negative (OFF) pulses that encoded the derivative of the signal and were sparser in time than a rate code, as they only responded to changes (Figure 7).

It should be noted that there was a mismatch between the timescale of metal-oxide gas sensors and the timescale of neurons. The gas sensors' impulse response was estimated to extend over several seconds [5]. The neuron models we employed operated on much faster timescales, on the order of milliseconds. The DVS algorithm created spikes upon changes of the input and the input signal was changing so slowly that, even with thresholds close to the noise level, the resulting inter-spike intervals were much larger than the time constants of neurons and synapses. In consequence, there was no integration of inputs on the neuronal membrane models and therefore the contribution of the neural dynamics to the output vanished.

We currently circumvent this problem by speeding up the recorded data by a factor of approximately 70 to match the timescale of neural integration, compressing 12 minutes of recording into 10 seconds of neural simulation. Another, equally workable solution would be to slow down neuronal time. In the future, sensors with faster response times, such as photo-ionisation detectors might provide

temporal resolution that is more compatible with neuronal timescales, albeit at the expense of gas recognition ability.

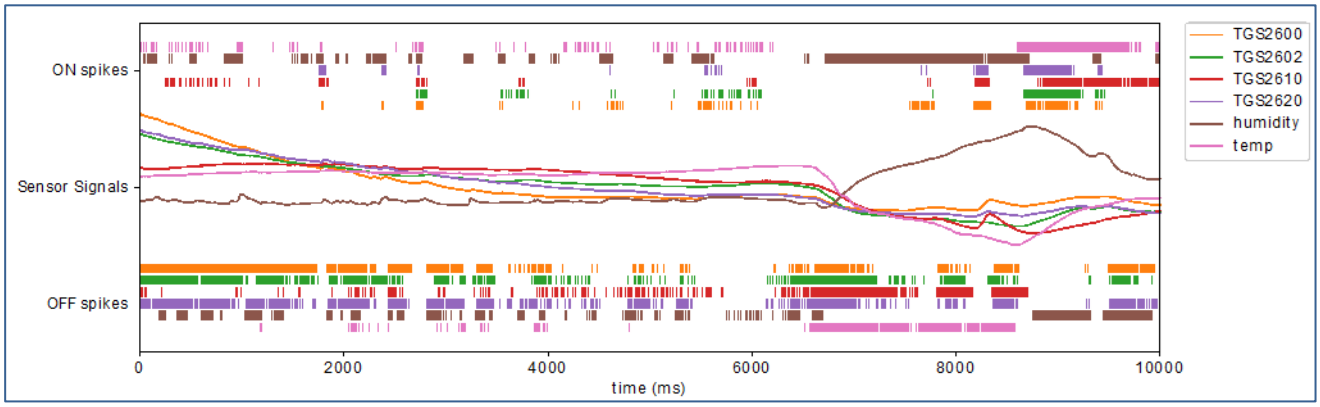

Figure 7: Sensor signals and ON/OFF spike encoding

4.2 Temporal filter banks

Broadly speaking, there are three ways to detect temporal features with spiking neural networks. One can employ delay lines or sliding windows to fold time into an input dimension; exploit recurrent network activity, as in echo state networks; or rely on the intrinsic dynamics of the neurons and synapses.

Bout detection as performed in [6] is based on the detection of periods where the input signal is continuously rising, i.e. has positive slope, and we had already suggested in that study that the bout detection algorithm is consistent with the operation of an adaptive neuron. Kepecs and coworkers have equally described how neurons that fire a short burst of spikes in response to a current step (onset detection) will respond to a current ramp with sustained spiking [3].

Here we exploited intrinsic dynamics of a neuron model to detect positive input slope. We picked the Izhikevich neuron model as one that possesses the required dynamical repertoire and runs efficiently on SpiNNaker. The dynamics of the Izhikevich model are set by five parameters: an adaptation time constant *a*, subthreshold adaptation *b*, after-spike adaptation *d*, after-spike reset *c* and a current bias term *I_bias* which sets the resting potential and baseline excitability of the neuron.

We created six neuronal populations corresponding to six sensor inputs. Each population contained 30 Izhikevich neurons getting their input from a single sensor. The ON and OFF spike trains targeted excitatory and inhibitory conductance-based synapses, which in turn drove the membrane potentials of the neurons. Within each filter bank population, the synaptic parameters were uniform while the five neuron parameters were each drawn from a distribution that was tuned to generate diverse dynamics and neural responses. The result was that some of the neurons exhibited regular spiking, where they fired with a frequency proportional to the input, whereas others fired bursts of spikes in response to the onset of excitation, or in response to the offset of inhibition (post-inhibitory rebound).

Figure 8 shows the output spikes of the filter banks and highlights the effect of these various dynamics. During the gas signal spike shortly after t=8,000 ms, some neurons responded only to the rising slope (black bracket) and function as bout detectors, while others responded throughout the positive area (red bracket). Some neurons with post-inhibitory rebound also responded to the positive slopes of disinhibitory bouts (purple square), even though there was no excitatory input to that filter bank at this point.

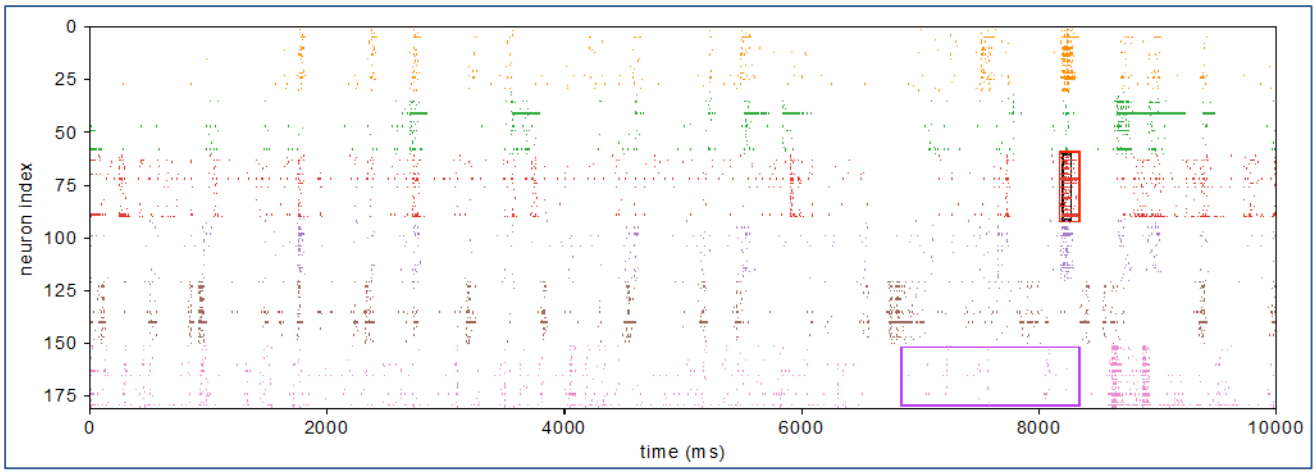

Figure 8: Output of the filter banks (colour-coded by sensor)

4.3 Sparse Fan-out Layer

The final processing step involved expanding the output of the six filter banks into a sparse, higher-dimensional fan-out layer. This mirrored a configuration often found in the olfactory system of insects and invertebrates, where the fan-out layer is followed by a fan-in layer which also receives reinforcement signals and performs associative learning.

The fan-out layer used the same neuron and synapse parameters as for the filter banks, but had sparse feedforward connections: each filter bank neuron made excitatory connections to approximately five randomly chosen fan-out neurons, and inhibitory connections to another five.

This resulted in a sparse code where each neuron spiked in response to specific combination of temporal features from each sensor (Figure 9). This sparse code supported learning and inference for pattern recognition. This result is the basis for the next phase of the project, where we will assess the network's performance in inferring gas identity and source proximity from temporal dynamics of olfactory stimuli.

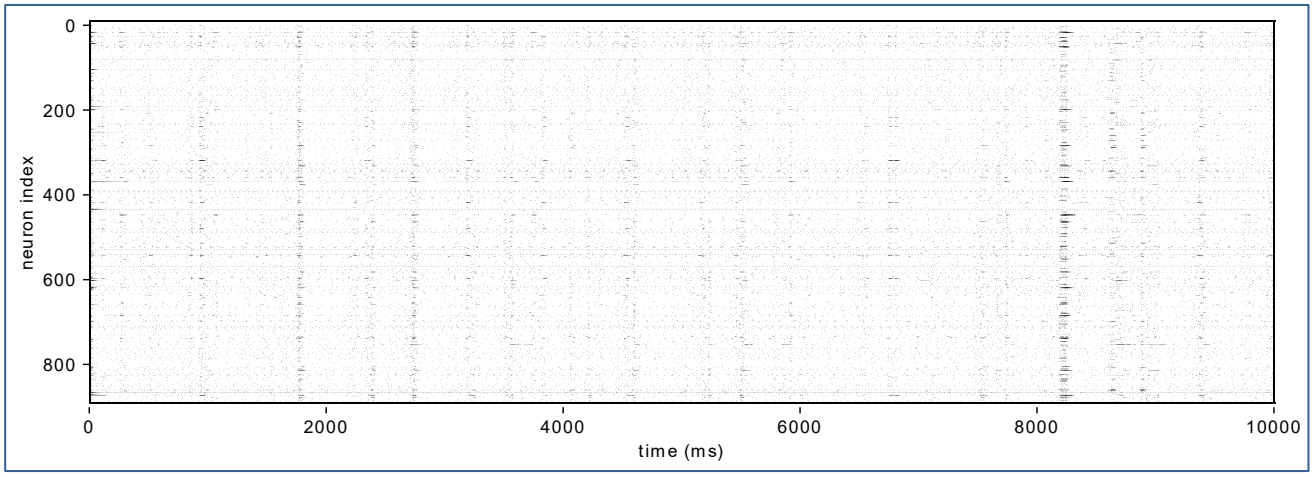

Figure 9: Output of the sparse fan-out layer

5. Conclusions and Future Steps

We have designed and implemented a neuromorphic olfaction system, consisting of of analogue gas sensors and a spiking network running on SpiNNaker, which performs temporal filtering of gas sensor input. The spiking network extracts features from the gas readings that are potentially helpful for gas-based navigation.

Gas-based navigation is an area of intense research due to its relevance in safety-critical operations, such as disaster management and environmental monitoring, where visual navigation alone is limited. A typical gas-based navigation task is, for example, to locate casualties on disaster sites. This is currently often performed by specially trained rescue dogs. Robotic applications could help reduce the risks to animals and their handlers, as well as improving availability and reducing costs. Another scenario where robotic solutions for gas-based navigation are highly relevant is for locating toxic gas leaks on industrial sites, flammable methane sources on landfill sites, or emerging wildfires. Our work aims to leverage event-based signal processing, machine learning and AI to help solve these technological challenges.

The next step in the project is to extract information from the output of the temporal filter network. We will achieve this using spiking networks performing pattern recognition and running on neuromorphic hardware [2]. Our long-term goal is to build a neuromorphic olfaction system that infers the identity and distance of gas sources, thus informing gas-based navigation algorithms for neurorobotics.

Concrete ongoing work focuses on better characterising the temporal features extracted by our network and preparing our setup for tasks such as learning and navigation. To work around the limitations of the ADC in our current sensor board, we are making a new prototype using a higher-resolution, 24-bit ADC that will be capable of a higher sampling rate as well (Figure 10).

Further work is currently investigating various spiking network architectures for supervised learning and inference, including LSNNs from the Maass group [1].

We are now working to apply our results in gas-based navigations scenarios, in conjunction with robotics groups. The first step is to simulate robotic gas-based navigation using the GADEN framework that is available as a ROS package [4]. Unfortunately, a voucher application to support this technology on the Neurorobotics platform was declined, but we are now seeking alternative routes outside the HBP to progress in this direction.

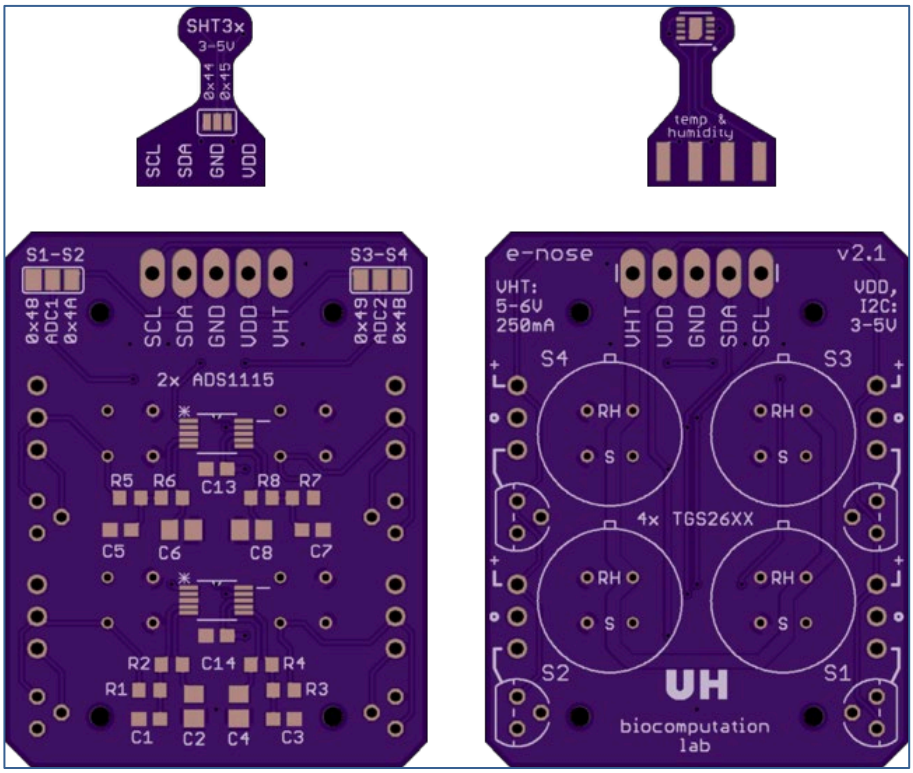

Figure 10: Printed Circuit Board (PCB) for the new e-nose sensor.

6. References

[1] Bellec, G., Salaj, D., Subramoney, A., Legenstein, R. A., and Maass, W. (2018). Long short-term memory and Learning-to-learn in networks of spiking neurons. *arXiv* cs.NE. (HBP publication ID P1449)

[2] Diamond, A., Nowotny, T., and Schmuker, M. (2016). Comparing Neuromorphic Solutions in Action: Implementing a Bio-Inspired Solution to a Benchmark Classification Task on Three Parallel-Computing Platforms. *Front. Neurosci.* 9, 20. doi:10.3389/fnins.2015.00491. (HBP publication ID P1765)

[3] Kepecs, A., Wang, X.-J., and Lisman, J. (2002). Bursting neurons signal input slope. *J. Neurosci.* 22, 9053-9062.

[4] Monroy, J., Hernandez-Bennetts, V., Fan, H. Lilienthal, A., and Gonzalez-Jimenez, J. (2017). GADEN: A 3D Gas Dispersion Simulator for Mobile Robot Olfaction in Realistic Environments. *MDPI Sensors* 17, 1479.

[5] Pashami, S., Lilienthal, A., and Trincavelli, M. (2012). Detecting Changes of a Distant Gas Source with an Array of MOX Gas Sensors. *Sensors* 12, 16404-16419. doi:10.3390/s121216404.

[6] Schmuker, M., Bahr, V., and Huerta, R. (2016). Exploiting plume structure to decode gas source distance using metal-oxide gas sensors. *Sensors and Actuators B: Chemical* 235, 636-646. doi:10.1016/j.snb.2016.05.098.

[7] Schmuker, M., and Drix, D. (2019). Code and data for HBP deliverable SGA2 D9.5.1 (D61.1, D93). Zenodo. http://doi.org/10.5281/zenodo.2602891 .

[8] Schmuker, M., Pfeil, T., and Nawrot, M. P. (2014). A neuromorphic network for generic multivariate data classification. *Proc. Natl. Acad. Sci. U.S.A.* 111, 2081-2086. doi:10.1073/pnas.1303053111.

[9] Vergara, A., Fonollosa, J., Mahiques, J., Trincavelli, M., Rulkov, N., and Huerta, R. (2013). On the performance of gas sensor arrays in open sampling systems using Inhibitory Support Vector Machines. *Sensors and Actuators B: Chemical* 185, 462-477. doi:10.1016/j.snb.2013.05.027.