Finding Bacteria in Blood: Scaling A Hardware-Driven Neuromorphic Solution for Real-World E-Nose Applications

***Abstract*—Recent developments in neuromorphic engineering have sparked tremendous interest in implementing these bio-inspired approaches for artificial olfactory systems. The inherent properties of applying neuromorphic processing for machine olfaction, such as massively parallel computing with sparse spike-based representations at minimal power requirements, present a promising solution for a wide range of practical applications. This work presents a hardware-based low-power neuromorphic solution that can be seamlessly integrated as a pattern recognition engine for an electronic nose system. Furthermore, results based on the implementation of the proposed approach for real-time classification of different bacteria in the blood are presented in this paper.**

Keywords—neuromorphic; olfaction; spiking neural networks; event-based applications; electronic noses

I. Introduction (*Heading 1*)

Artificial olfactory systems, or electronic noses (e-noses), has numerous applications, ranging from quality control in the food industry to minimally-invasive medical diagnosis. Biological olfactory systems detect volatile organic compounds (VOCs) using olfactory receptor neurons (ORNs) that relay a sparse spike-encoded signal through higher-order neurons to the brain, where decoding and recognition of odors takes place [1]. E-noses mimic this biological sensory system using chemical sensors and a pattern recognition engine [2]. While these systems have shown promising results in laboratory and controlled environments, scaling their applicability to real-world scenarios is currently limited by factors such as lack of portability, power efficiency, mediocre performance, and substantial latency to reach a reliable classification [2]–[4].

Machine learning has largely improved upon traditional e-nose systems, but this has been at the cost of using hand-crafted features and being driven by statistical approaches [3] or computationally-expensive deep-learning methods [2]. Whilst implementing these methods can potentially deliver improved accuracy, there is a further trade-off of portability and power efficiency [5].

Recently, bioinspired neuromorphic event-based approaches have shown promise in artificial olfaction. Neuromorphic olfaction aims to adopt the biological olfactory system's low power, event-driven, and highly accurate nature. Recent studies in this domain have presented results using spiking neural networks with low power implementations. However, these studies have primarily focused on incorporating a high degree of bio-realism, resulting in complex architectures that are impractical for real-world applications [6].

What is currently lacking is a bio-inspired approach that incorporates the advantage of the biological system using power-efficient neuromorphic hardware. This paper proposes a neuromorphic approach for processing e-nose data that can be seamlessly deployed on a neuromorphic-system-on-chip (NSoC). The proposed system comprises a data-to-spike encoder to generate events from continuous multi-variate sensory data and a spiking neural network (SNN)-based classifier. The primary aim of this work is to leverage the inherent properties of an event-based neuromorphic approach to develop a powerful inference engine capable of delivering highly accurate classification with low-power consumption and minimal latency.

II. Methods and Materials

A. Dataset

The work described in this paper used the bacteria in blood dataset collected as a part of the Mednose project at Örebro University. The Mednose project aimed to use an artificial olfactory system to identify different bacteria species in blood by detecting volatile organic compounds (VOCs) released at the early stages of bacterial incubation.

The data was collected using the NST 3220 Emission Analyzer, an electronic nose system developed by Applied Sensors (Linköping, Sweden) [2], [7]. This sampling system comprises 10 MOS and 12 MOSFET sensors combined in an array. Experiments were conducted to identify ten different bacteria species in blood samples. Altogether, 1200 data samples of the ten bacteria were collected, 120 for each species. A single data sample consisted of a 5-minute measurement, wherein a baseline response of the sensors for a reference gas was measured for the first 10 seconds, followed by 30 seconds of exposure to the target compounds to record the VOCs, and finally, a 260 seconds recovery phase allowing the sensors to return to baseline. Data acquisition was performed at a sampling rate of 2Hz, resulting in 601 data points in each sample, referring to changes in resistance when exposed to the target compounds. A detailed description of the e-nose experiments and the sampling protocols can be found in [7].

B. *Pre-Processing and Event Encoding*

Pre-processing in the form of baseline cancellation followed by normalization using local min-max scaling was performed on each data sample. The classifier model proposed in this study was developed using Akida SNN, which optimally processes event-encoded data represented in the address event representation (AER) format, a de-facto standard protocol in neuromorphic systems to represent sparse event-based information. Hence, one of the critical aspects of this study was to determine the most effective strategy to encode the continuous multi-variate sensor responses into event-based data.

Two different techniques were investigated to encode sensor responses into event-based data: firstly, we implement address event representation for olfaction (AERO) [8], an extension of the AER format for olfactory systems that has, in recent literature, exhibited its capabilities at encoding key features within the sensor responses. AERO uniformly discretizes the sensor responses into amplitude bins, where, at each time-step for a sensor, an event (binary 1) is recorded in the corresponding bin, resulting in a binarized spiking representation of the data. This technique generates a three-dimensional tensor that contains information about the sensor index in the array, time-step and approximate magnitude of the event. Fig. 1 is a reconstruction of the data from AERO encoding.

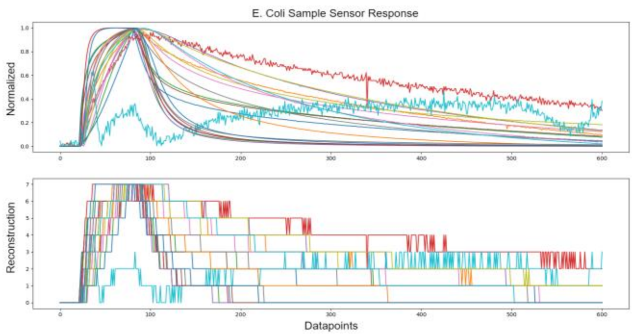

Fig. 1. Normalized sensor response and reconstructed sensor response after AERO encoding for an E. Coli bacteria sample.

The second encoding technique was based on the Step Forward (SF) algorithm. SF is a bio-inspired threshold-based approach with an adaptive baseline. SF generates spiking responses when there is a change in the information. It has been demonstrated that SF can encode fast-changing signals efficiently with minimum reconstruction error, and its implementation has been proven for machine olfaction applications using SNN-based classifiers [9]. The algorithmic implementation of this technique is detailed in [10].

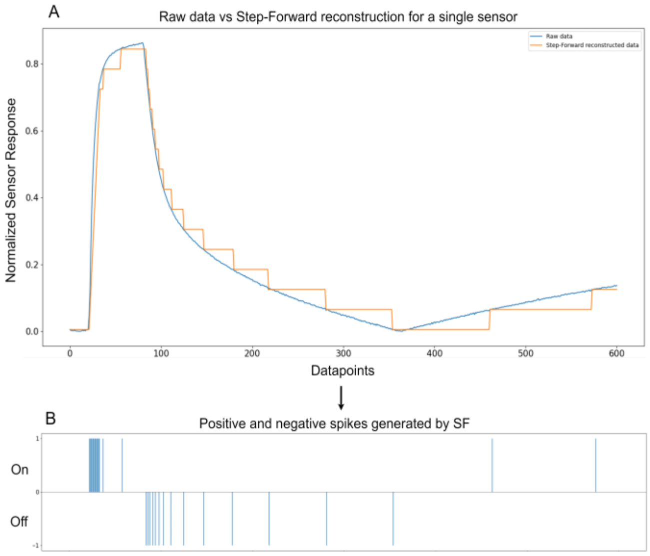

Fig. 2. **A.** Reconstruction of the response curve using the Step-Forward algorithm. **B.** On and off spikes generated by the SF algorithm.

In this case, the resulting three-dimensional tensor contains information about the sensor, time-step and polarity of an event. The polarity of the event determines whether the sensor reading registered above or below the baseline by more than the threshold. A sample reconstruction of the sensor reading can be seen in Fig. 2, which represents information being retained by SF encoding.

It was hypothesized that the full 5-minute (601 data points) sample may not be required to achieve a reliable classification and that once all the transient features had been captured, the remainder of the recovery period would be redundant. The number of data points used in the samples was varied to test this hypothesis.

III. SNN-Based Classifier

The events generated from the encoder were propagated through a two-layer SNN for training and classification. The SNN-based classifier comprised an input and a processing layer. These layers were composed of integrate-and-fire neurons, and the SNN was trained using a modified, low-bit-width and naturally homeostatic spike-time dependent plasticity (STDP) learning algorithm. Through this STDP implementation, the network learns to identify repeating patterns in the sensor responses.

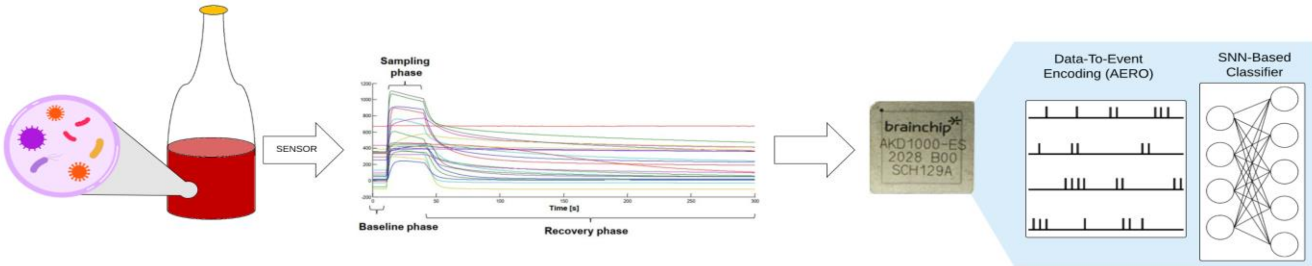

Fig. 3. Neuromorphic data processing hardware pipeline for olfactory sensor data with an SNN-based classifier.

When in inference mode, neurons that have learned these unique patterns will respond to the input data through changes in their potentials. A winner-take-all (WTA) [11] logic is applied, and the class of the winning neuron is identified to be provided as an output label to the user to classify the input sample. The effects of training the network on weight changes can be seen in Fig. 4, where this system enables the SNN to learn to represent the features of data samples.

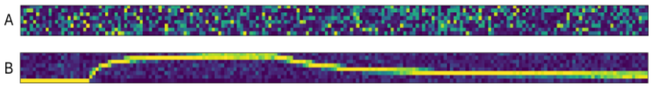

Fig. 4. Neuron receptive fields in the SNN model for sensor 5's response when exposed to an E. Coli in blood sample — **A.** Before training, and **B.** After training.

The network was configured using a grid-search parameter optimization to optimize the following hyper-parameters: the number of neurons per class in the fully connected layer; the number of non-zero weights to be used per neuron (i.e., the number of connections for each neuron); and the learning competition between neurons. As a result, the SF encoded model had 2385 neurons per class, 1884 weights were used per neuron, and a learning competition of 0.478 was set. The AERO encoded model used 120 neurons per class, 8400 weights, and a learning competition of 0.25. The other parameters — initial plasticity, minimum plasticity, and plasticity decay — had minimal impact on classification performance.

Before deploying the model on the Akida NSoC, the model was first validated in the python-based MetaTF development environment, which emulates the hardware capabilities of Akida. An overview of the proposed system is shown in Fig. 3 . The entire function of pre-processing and classification can be mapped on the NSoC.

IV. Classification Performance Analysis

One of the main objectives of this study was to determine if a reliable classification could be achieved without the requirement of an entire sampling frame as an input to the SNN model. In order to achieve this, an increasing window approach was used, with an evaluation frame size of 25 data points to classify the continuous e-nose sensor responses. The data-to-event encoder transformed the e-nose data within the evaluation frame into an event-based format before propagating it through the SNN model for classification. The performance of the SNN-based classifier was evaluated using a randomized stratified 3-fold cross-validation. The classification accuracy trend observed over an increasing number of data points is shown in Fig. 5.

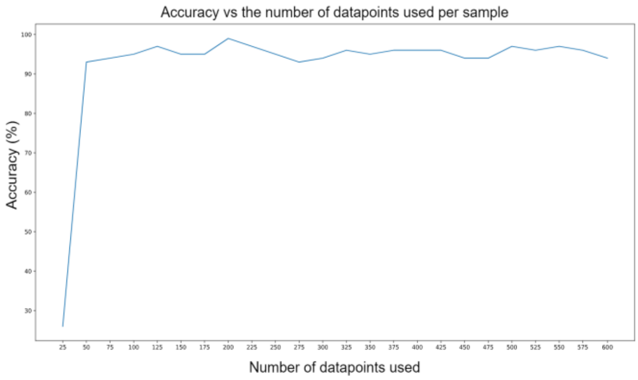

Fig. 5. Classification accuracy vs number of data points used for the SNN model.

Both encoding techniques resulted in high classification accuracy, as shown in Table 1. These results indicate that transient features critical to distinguishing between classes were present within the first 200 data points. Furthermore, it can be inferred that the sampling time for a reliable classification could be significantly reduced, and the extended recovery response did not directly contribute to the result.

TABLE 1

Classification Performance		
Encoding algorithm	***Data points per sample***	***Mean classification accuracy using 3-fold cross validation (%)***
AERO	200	97.33
	600	95.08
Step Forward	200	97.42
	600	96.17

While the classification performance for both data-to-event encoding techniques was broadly similar, it is important to note that using SF-based event encoding significantly reduced the input dimensions of the data, resulting in an efficient encoding of the input space. However, in this case, the trade-off was the pre-processing requirements, which might further affect the overall processing latency and result in a model with higher parameters.

The SNN-based classifiers were deployed on the Akida NSoC to validate their performance on the neuromorphic hardware. The testing environment consisted of a PC with an Intel i7 10700 CPU and an Akida NSoC (AKD1000). This system is limited by CPU performance and the PCIe bus; however, enhanced performance may be observed when Akida IP is implemented directly near-sensor for processing. As anticipated, the hardware-implemented SNN classifiers' performance parameters were similar to those achieved using the MetaTF chip emulator platform. The recorded power estimation measurements corresponded solely to power consumed by the Akida neural fabric in performance mode (at 300 MHz clock speed). In this case, the dynamic power consumption of the AERO classifier model was observed to be 24.5 mW ± 0.5 mW with a throughput of 181 inferences per second. The dynamic efficiency of the classifier model, which is measured as energy consumed per inference, was observed to be 135 µJ/inference. The classifier model implementing SF encoding consumed an average of 25 mW ± 3 mW with a throughput of 31.5 inferences per second, thus resulting in a dynamic efficiency of 822 µJ/inference. The drop in performance in terms of throughput and dynamic efficiency for the SF encoder-based classifier results from higher neural resources utilized by the model. These results further reinforce that the hardware-based SNN classifier can be effectively used as a low-power real-time pattern recognition engine for e-nose devices.

V. Conclusion

Developing a real-time and power-efficient hardware-based implementation of a pattern-recognition engine has been of paramount importance for e-nose systems. This paper presents an application of an event-based neuromorphic approach for the encoding and classification of e-nose data. The proposed solution consists of a data-to-event encoder that implements AERO and SF-based event encoding of continuous multi-variate sensor array responses and a two-layer SNN-based classifier that implements bio-inspired learning to identify and classify repeating patterns in the event-encoded sensor responses. When applied to the MedNose dataset, the SNN-based classifier identified ten different bacteria types with a classification accuracy of 97.42%. The system outperforms the accuracy of previous implementations, which require a 20-sample window and achieve 96.2% classification accuracy [2], whilst only requiring a single sample for classification. The SNN model was validated on the Akida NSoC and recorded a dynamic power consumption of 24.5 mW and a throughput of 181 inferences per second. Such a system could significantly speed up disease diagnosis in the real world whilst achieving state-of-the-art classification performance.

References

[1] T. C. Pearce, S. S. Schiffman, H. T. Nagle, and J. W. Gardner, Handbook of machine olfaction: electronic nose technology. John Wiley & Sons, 2006.

[2] M. Längkvist and A. Loutfi, “Unsupervised feature learning for electronic nose data applied to bacteria identification in blood,” 2011.

[3] J. A. Covington, J. W. Gardner, A. Hamilton, T. C. Pearce, and S. L. Tan, “Towards a truly biomimetic olfactory microsystem: an artificial olfactory mucosa,” IET Nanobiotechnol., vol. 1, no. 2, pp. 15–21, 2007.

[4] H.-Y. Hsieh and K.-T. Tang, “VLSI implementation of a bio-inspired olfactory spiking neural network,” IEEE Trans. Neural Netw. Learn. Syst., vol. 23, no. 7, pp. 1065–1073, 2012.

[5] A. Vanarse, A. Osseiran, and A. Rassau, “An Investigation into Spike-Based Neuromorphic Approaches for Artificial Olfactory Systems,” Sens. Basel, vol. 17, no. 11, Nov. 2017, doi: 10.3390/s17112591.

[6] A. Diamond, M. Schmuker, A. Z. Berna, S. Trowell, and T. Nowotny, “Classifying continuous, real-time e-nose sensor data using a bio-inspired spiking network modelled on the insect olfactory system,” Bioinspir. Biomim., vol. 11, no. 2, p. 026002, 2016.

[7] M. Trincavelli, S. Coradeschi, A. Loutfi, B. Söderquist, and P. Thunberg, “Direct identification of bacteria in blood culture samples using an electronic nose,” IEEE Trans. Biomed. Eng., vol. 57, no. 12, pp. 2884–2890, 2010.

[8] A. Vanarse, A. Osseiran, A. Rassau, and P. van der Made, “A hardware-deployable neuromorphic solution for encoding and classification of electronic nose data,” Sensors, vol. 19, no. 22, p. 4831, 2019.

[9] A. Vanarse, J. I. Espinosa-Ramos, A. Osseiran, A. Rassau, and N. Kasabov, “Application of a Brain-Inspired Spiking Neural Network Architecture to Odor Data Classification,” Sensors, vol. 20, no. 10, p. 2756, 2020.

[10] B. Petro, N. Kasabov, and R. M. Kiss, “Selection and Optimization of Temporal Spike Encoding Methods for Spiking Neural Networks,” IEEE Trans. Neural Netw. Learn. Syst., vol. 31, no. 2, pp. 358–370, Feb. 2020, doi: 10.1109/TNNLS.2019.2906158.

[11] P. Ferré, F. Mamalet, and S. J. Thorpe, “Unsupervised Feature Learning With Winner-Takes-All Based STDP,” Front. Comput. Neurosci., vol. 12, 2018, Accessed: Feb. 17, 2022. [Online]. Available:https://www.frontiersin.org/article/10.3389/fncom.2018.00024