



A Spike Time-Dependent Online Learning Algorithm Derived From Biological Olfaction

Ayon Borthakur^{1*} and Thomas A. Cleland²

¹ Computational Physiology Laboratory, Field of Computational Biology, Cornell University, Ithaca, NY, United States,

² Computational Physiology Laboratory, Department of Psychology, Cornell University, Ithaca, NY, United States

We have developed a spiking neural network (SNN) algorithm for signal restoration and identification based on principles extracted from the mammalian olfactory system and broadly applicable to input from arbitrary sensor arrays. For interpretability and development purposes, we here examine the properties of its initial feedforward projection. Like the full algorithm, this feedforward component is fully spike timing-based, and utilizes online learning based on local synaptic rules such as spike timing-dependent plasticity (STDP). Using an intermediate metric to assess the properties of this initial projection, the feedforward network exhibits high classification performance after few-shot learning without catastrophic forgetting, and includes a *none of the above* outcome to reflect classifier confidence. We demonstrate online learning performance using a publicly available machine olfaction dataset with challenges including relatively small training sets, variable stimulus concentrations, and 3 years of sensor drift.

Keywords: SNN, online learning, olfaction, STDP, local learning, spike time coding

OPEN ACCESS

Edited by:

Emre O. Neftci,
University of California, Irvine,
United States

Reviewed by:

Hesham Mostafa,
University of California, San Diego,
United States
Thomas Nowotny,
University of Sussex, United Kingdom

*Correspondence:

Ayon Borthakur
ab2535@cornell.edu

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 02 February 2019

Accepted: 07 June 2019

Published: 27 June 2019

Citation:

Borthakur A and Cleland TA (2019) A
Spike Time-Dependent Online
Learning Algorithm Derived From
Biological Olfaction.
Front. Neurosci. 13:656.
doi: 10.3389/fnins.2019.00656

INTRODUCTION

Convolutional networks have enabled tremendous progress in image recognition. However, analogous problems in high-dimensional modalities that lack the two-dimensional internal structure of visual images are not well-addressed by these networks, and the development of brain-mimetic network-based signal identification strategies in such modalities has lagged. This is unfortunate, as there are innumerable applications for such classifiers, including medical screening, genomics, and machine olfaction. Among these, machine olfaction methods have been directly inspired by the mammalian and insect olfactory systems—highly structured and well-studied biological networks that learn rapidly and non-iteratively, utilize local learning rules, resist catastrophic forgetting, can identify and learn new classes of odors (i.e., that do not map to existing representations), and can robustly identify signals of interest in the presence of strong interference. We studied the mammalian olfactory system in order to extract computational principles and algorithms that could underlie its unmatched ability to identify and classify genuinely high-dimensional signals under a variety of challenging conditions.

Most current research effort in machine olfaction is devoted to sensor development, including technologies such as multi-chamber metal oxide semiconductor (MOS) sensors (Gonzalez et al., 2011), high-density polymer sensors (Beccherelli et al., 2010), molecularly imprinted MOS and polymer sensors (Shi et al., 1999; Iskierko et al., 2016; Zhang et al., 2017), and surface acoustic wave sensors (Länge et al., 2008). In an effort to mimic properties of the biological system, there even have been efforts to develop sensors based on G protein-coupled receptor proteins bound to

carbon nanotube transistors (Liu et al., 2006). In contrast, there has been relatively little effort spent mining the post-sensory networks of the olfactory system for clues to its unmatched performance, despite a broad understanding that biological odorant receptors are neither particularly specific nor particularly sensitive to odor stimuli. Rather, the power of the biological olfactory system derives from the concerted effects of the large numbers and diversity of its sensors, and by its post-sensory signal processing in the olfactory bulb and related cortices. These core principles inform recent developments in neuromorphic olfaction (Persaud et al., 2013; Schmuker et al., 2015), and have been highlighted in contemporary artificial systems work based on the similarly-structured olfactory system of insects (Schmuker et al., 2014; Mehta et al., 2017; Diamond et al., 2019).

We here present a spiking neural network (SNN)-based online learning algorithm, based on principles and motifs derived from the mammalian olfactory system, that can accurately classify noisy high-dimensional signals into categories that have been dynamically defined by few-shot learning. In order to better interpret the basis for the algorithm's capabilities, the present work focuses entirely on the properties of the first feedforward projection, omitting the spike timing-based feedback loop that forms the core network of the full OB model (Imam and Cleland, 2019). Glomerular-layer processing is represented here by two preprocessing algorithms, whereas plasticity for rapid learning is embedded in subsequent processing by the external plexiform layer (EPL) network. Information in the EPL network is mediated by patterns of spike timing with respect to a common clock corresponding to the biological gamma rhythm, and learning is based on localized spike timing-based synaptic plasticity rules. The algorithm is implemented in PyTorch for GPU computation, but designed for later implementation on state-of-the-art neuromorphic computing hardware (Davies et al., 2018); the initial version of the complete attractor model has been implemented on Intel Loihi (Imam and Cleland, 2019). We here demonstrate the interim performance of the feedforward algorithm using a well-established machine olfaction dataset with distinct challenges including multiple odorant classes, variable stimulus concentrations, physically degraded sensors, and substantial sensor drift over time.

CORE PRINCIPLES

The network is based on the architecture of the mammalian olfactory bulb (reviewed in Cleland, 2014; Nagayama et al., 2014). Primary olfactory sensory neurons (OSNs) express a single odorant receptor type from a family of hundreds (depending on animal species). The axons of OSNs that express the same receptor type converge to a common location on the surface of the olfactory bulb (OB), forming a mass of neuropil called a glomerulus. Each glomerulus thus is associated with exactly one receptor type, and serves as the basis for an OB *column*. The profile of glomerular activation levels across the hundreds of receptor types (~400 in humans, ~1,200 in rats and mice) that are activated by a given odorant constitutes a high-dimensional vector of sensory input (Zaidi et al., 2013). Within this first (*glomerular*) layer of the OB, a number of preprocessing

computations also are performed, including a high-dimensional form of contrast enhancement (Cleland and Sethupathy, 2006) and an intricate set of computations mediating a type of global feedback normalization that enables concentration tolerance (Cleland et al., 2012). The cellular and synaptic properties of this layer also begin the process of transforming stationary input vectors into spike timing-based representations discretized by 30–80 Hz gamma oscillations (Kashiwadani et al., 1999; Li and Cleland, 2017). The EPL, which constitutes the deeper computational layer of the OB, comprises a matrix of reciprocal interactions between principal neurons activated by sensory input (mitral cells; MCs) and inhibitory interneurons (granule cells; GCs). Computations in this layer depend on fine-timescale spike timing (Lepousez and Lledo, 2013) and odor learning (Lepousez et al., 2014; Mandairon et al., 2018), and modify the information exported from the OB to its follower cortices.

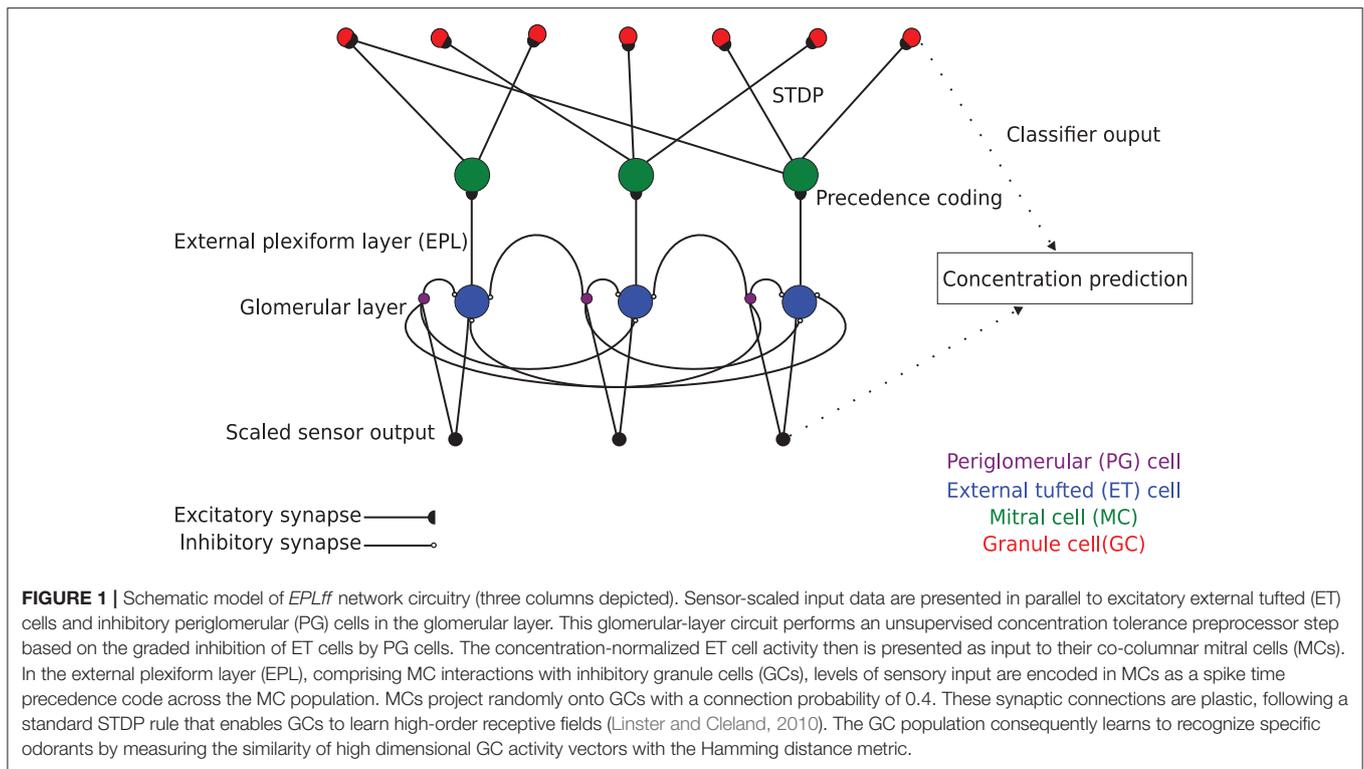
Chemical sensing in machine olfaction is similarly based upon combinatorial coding (Persaud and Dodd, 1982); specificity is achieved by combining the responses of many poorly-selective sensors. In the present algorithm, networks were defined with a number of columns such that each column received input from one type of sensor in the connected input array. Columns each comprised one external tufted (ET) cell and one periglomerular (PG) cell to mediate glomerular-layer preprocessing, and one MC and a variable number of GCs to mediate EPL odorant learning and classification (**Figure 1**; see section Online Learning). Sensory input was preprocessed by the ET and PG cells of the glomerular layer (for concentration tolerance), and then delivered as excitation to the array of MCs, which generated action potentials. Each MC synaptically excited a number of randomly determined GCs drawn from across the entire network, whereas activated GCs synaptically inhibited the MC in their home column. Importantly, for present purposes, these inhibitory feedback weights were all reduced to zero to disable the feedback loop and EPL attractor dynamics, enabling study of the initial feedforward transformation based on excitatory synaptic plasticity alone. During learning, the excitatory synapses followed a STDP rule that systematically altered their weights, thereby modifying the complex receptive fields of recipient GCs in the service of odor learning. In the present study, in lieu of the modified spike timing of the MC ensemble that characterizes the output of the full model (Imam and Cleland, 2019), the binary vector describing GC ensemble activity in response to odor stimulation (0: non-spiking GC; 1: spiking GC) served as the processed data for classification. Because we here report the capacities of the initial feedforward projection of preprocessed data onto the GC interneuron array within the EPL—an initial transformation that sets the stage for ongoing dynamics not discussed herein—we refer to our present method as the *EPLff* network algorithm.

MATERIALS AND METHODS

Data Preprocessing

Sensor Scaling

We defined a set of preprocessing algorithms, any or all of which could be applied to a given data set to prepare it for efficient analysis by the core algorithm. The first of these,



sensor scaling, is applied to compensate for heterogeneity in the scales of different sensors—for example, an array comprising a combination of 1.8V and 5V sensors. One simple solution is to scale the responses of each sensor by the maximum response of that sensor. Let $x_1, x_2, x_3, \dots, x_n$ be the responses of n sensors to a given odor and $s_1, s_2, s_3, \dots, s_n$ be the maximum response values of those sensors. Then, $\frac{x_1}{s_1}, \frac{x_2}{s_2}, \frac{x_3}{s_3}, \dots, \frac{x_n}{s_n}$ represent the sensor-scaled responses. The maximum sensor response vector S could be predetermined (as in sensor voltages), or estimated using a model validation set. Here, we defined S using the model validation set (10% of Batch 1 data; see section Dataset) and utilized the same value of S for scaling all subsequent learning and inference data (see section Sensor Drift). This preprocessing algorithm becomes particularly useful when analyzing data from arbitrary or uncharacterized sensors, or from arrays of sensors that have degraded and drifted non-uniformly over time.

Unsupervised Concentration Tolerance

Concentration tolerance is a critical feature of mammalian as well as insect olfaction (Cleland and Sethupathy, 2006; Cleland et al., 2012; Serrano et al., 2013). Changes in odorant concentration evoke non-linear effects in receptor activation patterns that are substantial in magnitude and often indistinguishable from those based on changes in odor quality. Distinguishing concentration differences from genuine quality differences appears to rely upon multiple coordinated mechanisms within olfactory bulb circuitry (Cleland et al., 2012), but the most important of these is a global inhibitory feedback mechanism instantiated in the deep

glomerular layer (Cleland et al., 2007; Banerjee et al., 2015). The consequence of this circuit is that MC spike rates are not strongly or uniformly affected by concentration changes, and the overall activation of the olfactory bulb network remains relatively stable. We implemented this concentration tolerance mechanism as the graded inhibition of external tufted cells (ET) by periglomerular cell (PG) interneurons in the OB glomerular layer (**Figure 1**)—a mechanism based upon recent experimental findings in which ET cells serve as the primary gates of MC activation (Gire et al., 2012; Banerjee et al., 2015)—and tested its importance empirically on machine olfaction data sets. This concentration tolerance mechanism facilitates recognition of odor stimuli even when they are encountered at concentrations on which the network has not been trained; moreover, once an odor has been identified, its concentration can be estimated based on the level of feedback that the network delivers in response to its presentation. This preprocessing step requires no information about input data labels, and greatly facilitates few-shot learning.

Input from each sensor was delivered directly to PG and ET interneurons associated with the column corresponding to that sensor, and the resulting PG cell activity was delivered via graded synaptic inhibition onto all ET cells within all columns in the network. ET cells in turn then synaptically excited their corresponding, cocolumnar MCs (**Figure 1**). The approximate outcome of this preprocessor algorithm is as follows: given that $x_1^{ET}, x_2^{ET}, x_3^{ET}, \dots, x_n^{ET}$ denote the responses of ET cells to odor inputs (prior to their inhibition by PG cells), and $x_1^{PG}, x_2^{PG}, x_3^{PG}, \dots, x_n^{PG}$ denote the analogous responses of

PG interneurons to these same inputs, the resulting input to MC somata from ET cells following their PG-mediated lateral inhibition will be

$$\frac{x_1^{ET}}{\sum x^{pg}}, \frac{x_2^{ET}}{\sum x^{pg}}, \frac{x_3^{ET}}{\sum x^{pg}}, \dots, \frac{x_n^{ET}}{\sum x^{pg}} \quad (1)$$

A version of this algorithm has been implemented using spiking networks on IBM TrueNorth neuromorphic hardware (Imam et al., 2012).

Core Algorithm

Cellular and Synaptic Models

We modeled the MCs and GCs as leaky integrate-and-fire neurons with an update period of 0.01 ms. The evolution of the membrane potential v of MCs and GCs over time was described as

$$\tau \frac{dv}{dt} = -v + IR \quad (2)$$

where $\tau = r_m c_m$ was the membrane time constant and r_m and c_m denote the membrane resistance and capacitance, respectively. For MCs, the input current I corresponded to sensory input received from ET cells (after preprocessing by the ET and PG neurons of the glomerular layer; **Figure 1**), whereas for GCs, I constituted the total synaptic input from convergent presynaptic MCs. In GCs, the parameter R was set to equal r_m , whereas in MCs it was set to r_m/r_{shunt} , where r_{shunt} was the oscillatory shunting inhibition of the gamma clock (described below). When $v \geq v_{th}$, where v_{th} denotes the spike threshold, a spike event was generated and v was reset to 0. The total excitatory current to GCs was modeled as

$$I = g_w(E_n - v) \quad (3)$$

where E_n was the Nernst potential of the excitatory current (+70 mV), v was the GC membrane potential, and $g_w = \sum_{i=1}^n w_i g_{max} \frac{\tau_1 \tau_2}{\tau_1 - \tau_2} (e^{-\frac{t-t_i}{\tau_1}} - e^{-\frac{t-t_i}{\tau_2}})$ describes the open probability of the AMPA-like synaptic conductances. Here, t_i denotes presynaptic spike timing, w_i denotes the synaptic weight, and g_{max} is a scaling factor.

The parameters $c_m, r_m, r_{shunt}, E_n, g_{max}, \tau_1$, and τ_2 were determined only once each for MCs and GCs using a synthetic data set (Borthakur and Cleland, 2017) and remained unchanged during the application of the algorithm to real datasets. The value of w_i at each synapse also was set to a fixed starting value based on synthetic data, but was dynamically updated according to the STDP learning rule. The spiking thresholds v_{th} of MCs and GCs were determined by assessing algorithm performance on the training and validation sets. Because we observed that using heterogeneous values of v_{th} across GCs improved performance, the values of v_{th} were randomly assigned across GCs from a uniform distribution.

Gamma Clock and Spike Precedence Code

Oscillations in the local field potential are observed throughout the brain, arising from the synchronization of activity in neuronal

ensembles. In the OB, gamma-band (30–80 Hz) oscillations are associated with the coordinated periodic inhibition of MCs by GCs (Li and Cleland, 2017; Peace et al., 2017) that constrains MC spike timing (Kashiwadani et al., 1999), thereby serving as a common clock. For this work, we modeled a single cycle gamma oscillation as a sinusoidal shunting inhibition r_{shunt} delivered onto all MCs,

$$r_{shunt} = -3.8 * \cos\left(\frac{2\pi * f * t}{1000}\right) + 5 \quad (4)$$

where f is the oscillation frequency (40 Hz) and t is the simulation time. We used a spike precedence coding scheme for MCs (Panzeri et al., 2010) where earlier MC spike phases correspond to stronger sensor input and are correspondingly more effective at growing and maintaining spike timing-dependent plastic synapses (Linster and Cleland, 2010). In the full model, the gamma clock serves as the iterative basis for the attractor; for present purposes in the *EPLff* context it served only to structure the spike times of active MCs converging onto particular GCs (precedence coding), and thereby to govern the changes in excitatory synaptic weights according to the STDP rule (see below).

Connection Topology

MC lateral dendrites support action potential propagation to GCs across the entire extent of the OB (Xiong and Chen, 2002; Peace et al., 2017), whereas inhibition of MCs by GCs is more localized. Excitatory MC-GC synapses were initialized with a uniformly distributed random probability cp of connection and a uniform weight w_0 ; synaptic weights were modified thereafter by learning. The initial connection probability cp was determined using a synthetic data set (Borthakur and Cleland, 2017), and was set to $cp = 0.4$ in the present simulations. For present purposes, as noted above, GC-MC inhibitory weights were set to zero to disable attractor dynamics.

Spike Timing-Dependent Plasticity Rule

We used a modified spike timing-dependent plasticity rule (STDP; Song et al., 2000; Dan and Poo, 2004) to regulate MC-GC excitatory synaptic weight modification. Briefly, synaptic weight changes were initiated by GC spikes and depended exponentially upon the spike timing difference between the postsynaptic GC spike and the presynaptic MC spike. When a presynaptic MC spike preceded its postsynaptic GC spike within the same gamma cycle, w for that synapse was increased; in contrast, when MC spikes followed GC spikes, or when a GC spike occurred without a presynaptic MC spike, w was decremented. Synaptic weights were limited by a maximum weight w_{max} . The pairing of STDP with MC spike precedence coding discretized by the gamma clock generated a *k winners take all* rule, in which the value of k depended substantially on the GC spike threshold v_{th} and the maximum excitatory synaptic weight w_{max} . Under this rule, activated GCs were transformed from non-specialized cells receiving weak inputs from a broad and random distribution of MCs into specialized, fully differentiated neurons that responded only to coordinated activation across a specific ensemble of k MCs. Under all training conditions, for present purposes, we set

a high learning rate such that, after one cycle of learning, each of the synapses could have one of only three values: w_0 , w_{\max} , or 0.

The STDP parameters were similar to our previous work using a synthetic data set (Borthakur and Cleland, 2017); among these, only the maximum synaptic weight w_{\max} was tuned based on validation set performance. For this feedforward implementation, online learning without the requirement of storing training data yielded its best validation set performance when $w_{\max} = w_0$, such that learning was limited to long-term synaptic depression (Borthakur and Cleland, 2017).

Classification

For the classification of test odorants in this reduced feedforward *EPLff* implementation, we calculated the Hamming distance between the binary vectors of GC odorant representations. Specifically, for every input, GCs generated a binary vector based upon whether the GC spiked (1) or did not spike (0). We matched the similarity of test set binary vectors with the training set vector(s) using the Hamming distance and classified the test sample based upon the label of the closest training sample. Alternatively, an overlap metric between GC activation patterns also was calculated (Equation 6 from Linster and Cleland, 2010); results based on this method were reliably identical to those of the Hamming distance and hence were omitted from this report. Classification was set to *none of the above* if the Hamming distance of the GC binary vectors was >0.5 , or if the overlap metric was <0.5 .

Dataset

We tested our algorithm on the publicly available UCSD gas sensor drift dataset (Vergara et al., 2012; Rodriguez-Lujan et al., 2014), slightly reorganized to better demonstrate online learning. The original dataset contains 13,910 measurements from an array of 16 polymer chemosensors exposed to six gas-phase odorants spanning a wide range of concentrations (10–1,000 *ppmv*) and distributed across 10 batches that were sampled over a period of 3 years to emphasize the challenge of sensor drift over time (Table 1). Owing to drift, the sensors' output statistics change drastically over the course of the 10 batches; between this property, the six different gas types, and the wide range of concentrations delivered, this dataset is well-suited to test the capabilities of the present algorithm without exceeding the learning capacity of its feedforward architecture (Figure 1). For the online learning scenario, we sorted each batch of data according to the odorant trained, but did not organize the data according to concentration. Hence, each training set comprised 1–10 odorant stimuli of the same type but at randomly selected concentrations. Test sets always included all six different odorants, again at randomly selected concentrations. For sensor scaling and the fine-tuning of the algorithm, we used 10% of the Batch 1 data as a validation set. The six odorants in the dataset are, in the order of training used herein: ammonia, acetaldehyde, acetone, ethylene, ethanol, and toluene. Batches 3–5 included only five different odorant stimuli, omitting toluene.

Eight features per chemosensor were recorded in the UCSD dataset, yielding a 128-dimensional feature vector. However, in contrast to previous efforts (Liu et al., 2015; Zhang and Zhang,

2015; Yan et al., 2017; Ma et al., 2018), we chose to use only one feature per sensor in our analysis (the steady state response level), for a total of 16 features. We imposed this restriction to challenge our algorithm, and because generating features from raw data requires additional processing, energy and time, all of which can impair the effectiveness of field-deployable hardware (Yin et al., 2018). Importantly, however, the sensor scaling and concentration tolerance preprocessors described above (section Data Preprocessing) would enable the *EPLff* network to utilize the full 128-dimensional dataset without specific adaptations other than expanding the number of columns accordingly.

RESULTS

Data Preprocessing

All sensory input data were preprocessed before being presented to the network. First, sensor scaling was applied to weight the 16 sensors equally in subsequent computations. The mean raw responses of the 16 sensors differed widely, with some sensors exhibiting an order of magnitude greater variance than others across the 10 odorants tested (Figure 2A). Sensor scaling (Figure 2B) mitigated this effect by scaling each sensor's gain such that the dynamic ranges of all sensors across the test battery were effectively equal. This process enabled each sensor to contribute a comparable amount of information to subsequent computations (up to a limit imposed by each sensor's signal to noise ratio), and improved network performance by maintaining consistent mean activity levels across test odorants.

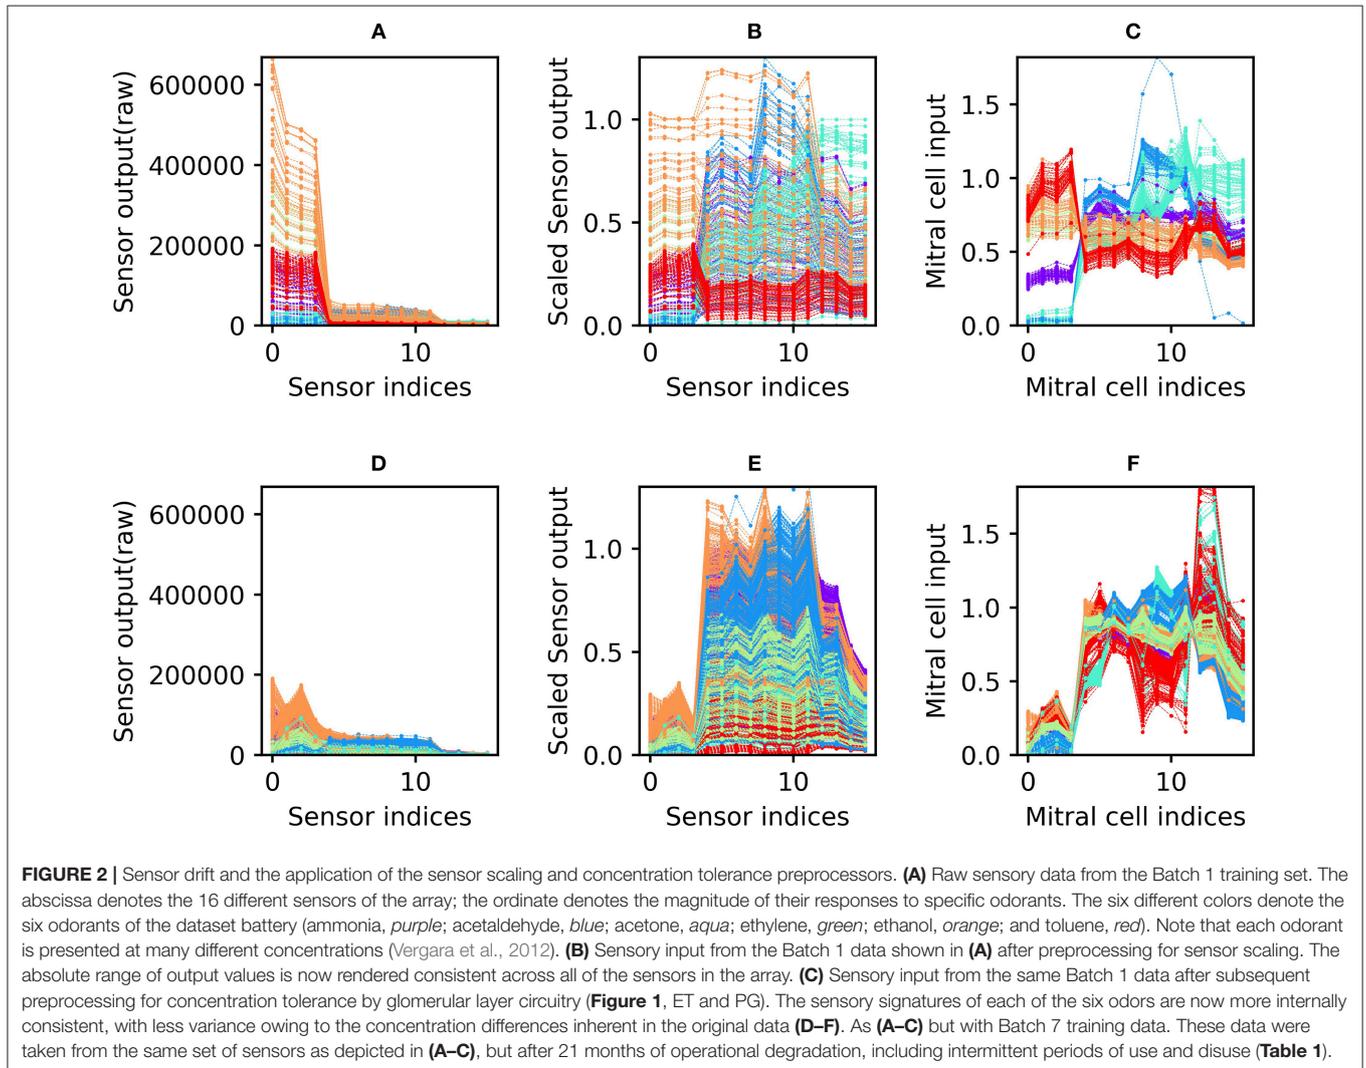
Since each odorant was presented at a wide range of randomly selected concentrations, the response of the sensor array to a given odorant varied widely across presentations (most clearly observable in Figure 2B). Application of the unsupervised concentration tolerance preprocessor sharply and selectively reduced the concentration-specific variance among responses to presented odorants (Figure 2C). These preprocessed odorant signatures then were presented to the plastic *EPLff* network for training or classification. Notably, this preprocessor step greatly facilitated cross-concentration odorant recognition, even enabling the accurate classification of samples presented at concentrations that were not included in the training set. This was particularly important for one- and few-shot learning, in which the network was trained on just one or a few exemplars (respectively), at unknown concentration(s), such that most of the odorants in the test set were presented at concentrations on which the network had never been trained.

The sensor scaling preprocessor (retaining the scaling factors determined from the 10% validation set of Batch 1), combined with the normalization effects of the subsequent concentration tolerance preprocessor, had the additional benefit of restoring the dynamic range of degraded sensors in order to better match classifier network parameters. Because of this, the network did not need to be reparameterized to effectively analyze the responses of the degraded sensors in the later batches of this dataset. Compared to the raw sensor output of Batch 1 (Figure 2A; collected from new sensors), the raw sensor output of Batch 7 (Figure 2D; collected after 21 months of sensor deterioration) was reduced to roughly a third of its original range.

TABLE 1 | Properties of the UCSD gas sensor drift dataset.

	Batch 1	Batch 2	Batch 3	Batch 4	Batch 5	Batch 6	Batch 7	Batch 8	Batch 9	Batch 10
Months	1–2	3–10	11–13	14–15	16	17–20	21	22–23	24–30	36
#Samples	445	1,244	1,586	161	197	2,300	3,613	294	470	3,600

Months denotes the age of the sensor array during the sampling of the corresponding dataset. #Samples denotes the number of samples provided by the dataset in that particular batch.



Sensor scaling (**Figure 2E**) mitigated this effect by magnifying sensor responses into the dynamic range expected by the network. Subsequent preprocessing for concentration tolerance effectively reduced concentration-specific variance, revealing a set of odorant profiles (**Figure 2F**) that, while qualitatively dissimilar to their profiles based on the same sensors 21 months prior (**Figure 2C**), appear only modestly degraded in terms of their distinctiveness from one another.

For many machine olfaction applications, it is useful to estimate the concentrations of gases in the vicinity of the sensors. We sought to use the information extracted

from the concentration tolerance preprocessor to estimate the concentrations of test samples after classification. The concentration estimation curve was a function of both odorant identity and the total sensor response profile. Using the sum of the 16 sensor responses (S), we fitted an odorant-specific quadratic curve for an implicit model of response profiles across concentrations C : $C = ax^2 + b$, where the parameters a and b were determined from the training set. **Figure 3** illustrates total sensor responses across concentrations compared to this theoretical prediction for all six odorant gases in Batches 1 and 7. The mean absolute error (MAE) of the prediction (in *ppmv*)

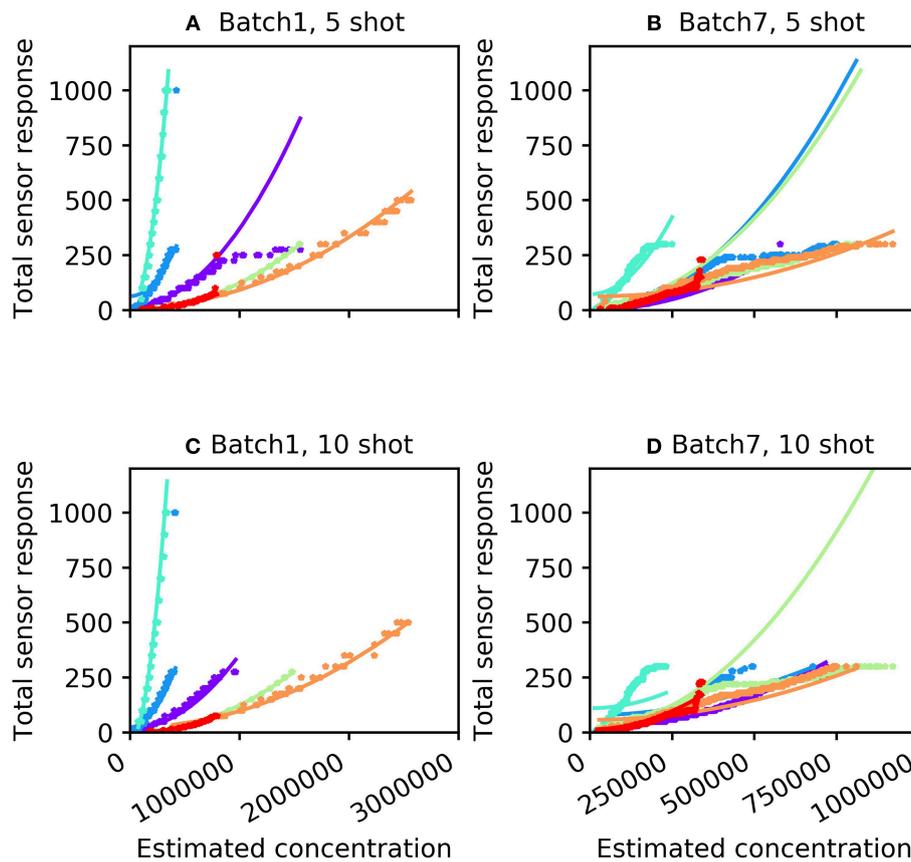


FIGURE 3 | Concentration response function predicted by the algorithm (curves) compared with measured sensor responses across multiple concentrations (stars). **(A)** Batch 1 data with five-shot training. **(B)** Batch 7 data with five-shot training. **(C)** Batch 1 data with 10-shot training. **(D)** Batch 7 data with 10-shot training. The colors denoting particular odorants are the same as in **Figure 2**.

TABLE 2 | Concentration estimation performance on test sets of all batches of UCSD gas sensor drift dataset for 5- and 10-shot learning (see **Figure 3**).

	Batch 1	Batch 2	Batch 3	Batch 4	Batch 5	Batch 6	Batch 7	Batch 8	Batch 9	Batch 10	Mean error
5 shot	35.14	51.39	32.00	44.73	66.02	37.01	76.60	17.89	0.34	71.06	43.22
10 shot	23.35	33.60	28.97	35.75	64.71	25.23	58.18	19.66	0.84	52.61	34.29

Concentrations were estimated using the predicted labels and raw sensor input. Errors represent experimental deviation from the predicted quadratic concentration curves and are in units of ppmv.

was estimated as

$$\frac{\sum_n |C_{pred} - C_{actual}|}{n} \quad (5)$$

where n denotes the total number of samples. For the five-shot training of Batch 1 (i.e., five random samples drawn from Batch 1 for each odorant), the MAE was 35.14 units (**Table 2**). This error was reduced to 23.35 for 10-shot learning (**Table 2**). Similarly, the MAE for Batch 7 decreased from 76.60 (five-shot) to 58.18 (10-shot). To the best of our knowledge, this is the first parallel network architecture to provide an estimate of concentration along with concentration tolerance.

Online Learning

Unlike biological odor learning, artificial neural networks optimized for a certain task tend to suffer from catastrophic forgetting, and the pursuit of online learning capabilities in deep networks is a subject of active study (McCloskey and Cohen, 1989; Kemker and Kanan, 2017; Kirkpatrick et al., 2017; Velez and Clune, 2017; Zenke et al., 2017; Serrà et al., 2018). In contrast, the *EPLff* learning network described herein naturally resists catastrophic forgetting, exhibiting powerful online learning using a fast spike timing-based coding metric. Moreover, we include a *none of the above* outcome which permits classification only above a threshold level of confidence (Huerta and Nowotny, 2009). Hence, after being trained on one odorant, the network

could identify a test sample as either that odorant or *none of the above*. After subsequently training the network on a second odorant, it could classify a test sample as either the first trained odorant, the second trained odorant, or *none of the above*. This online learning capacity enables *ad hoc* training of the network, with intermittent testing if desired, with no need to train on or even establish the full list of classifiable odorants in advance. It also facilitates training under missing data conditions (e.g., batches 3–5 contain samples from only five odorants, unlike the other batches which include six odorants), and could be utilized to trigger new learning in an unsupervised exploration context. Finally, once learned, the training set data need not be stored.

To analyze the 16-sensor UCSD dataset, we constructed a 16-column spiking network with 4800 GC interneurons and a uniformly random MC-GC connection probability $cp = 0.4$. This number of GCs was selected because it was the smallest network that achieved asymptotic performance on the validation dataset (Batch 1, one-shot learning; **Table 3**). We then trained this network on ammonia using 10 different few-shot training schemes: one-shot, two-shot, three-shot, up through 10-shot in order to measure the utility of additional training. Test data (across all trained odorants and all concentrations in the dataset) were classified with 100.0% accuracy in all cases (**Figure 4A**; average of three runs). We subsequently trained each of these trained networks on acetaldehyde, using the same number of training trials in each case. After one-shot learning of acetaldehyde, the network classified all trained odorants with $99.61 \pm 0.28\%$ accuracy (average of three runs). After subsequent one-shot learning of acetone, classification performance was $95.65 \pm 0.19\%$; after ethylene, $96.06 \pm 0.17\%$; after ethanol, $90.94 \pm 0.0\%$, and finally, after one-shot training on the sixth and final odorant, toluene, test set classification performance across all odorants was $90.27 \pm 0.12\%$. Multiple-shot learning generally produced correspondingly higher classification performance as the training regimen expanded (**Figure 4A**). Classification using an overlap metric (Linster and Cleland, 2010) rather than the Hamming distance yielded almost identical results (not shown). Critically, classification performance did not catastrophically decline as additional odorants were learned in series (**Figure 4**, purple to red (orange) traces in order), particularly when higher-quality sensors were used (**Figures 4A–E**) or when larger multiple-shot training sets were employed (**Figure 4**, panel abscissas). These results illustrate that the *EPLff* network, even in the absence of the full model's recurrent component, exhibits true online learning.

The availability of data in the UCSD dataset from over 3 years of sensor deterioration enabled the testing of this online learning algorithm with both fresh and degraded sensor arrays. **Figures 4B–J** presents classification results from the same procedures described above but using progressively older and more degraded sensors (Batches 2–10; **Table 1**; Vergara et al., 2012). Classification performance declined overall as the sensors deteriorated in later batches (**Figures 4F–J**), but could be substantially rescued by expanding the training regimen from one-shot to few-shot learning. Overall, multiple-shot training reliably improved classification performance, though the residual variance across different training regimes suggests that the

TABLE 3 | Effect of increased numbers of GCs in the network (GC vector length) on *EPLff* classification accuracy by the Hamming distance criterion, based on one-shot learning using the Batch 1 validation set.

#GC	1 class trained	2 classes trained	3 classes trained	4 classes trained	5 classes trained	6 classes trained
160	100.0	100.0	95.65	96.15	85.71	84.44
1,600	100.0	100.0	95.65	96.15	85.71	86.67
4,800	100.0	100.0	95.65	96.15	85.71	88.89
9,600	100.0	100.0	95.65	96.15	85.71	88.89
14,400	100.0	100.0	95.65	96.15	85.71	88.89
19,200	100.0	100.0	95.65	96.15	85.71	88.89

Connection probabilities and initial synaptic weights were consistent across all simulations.

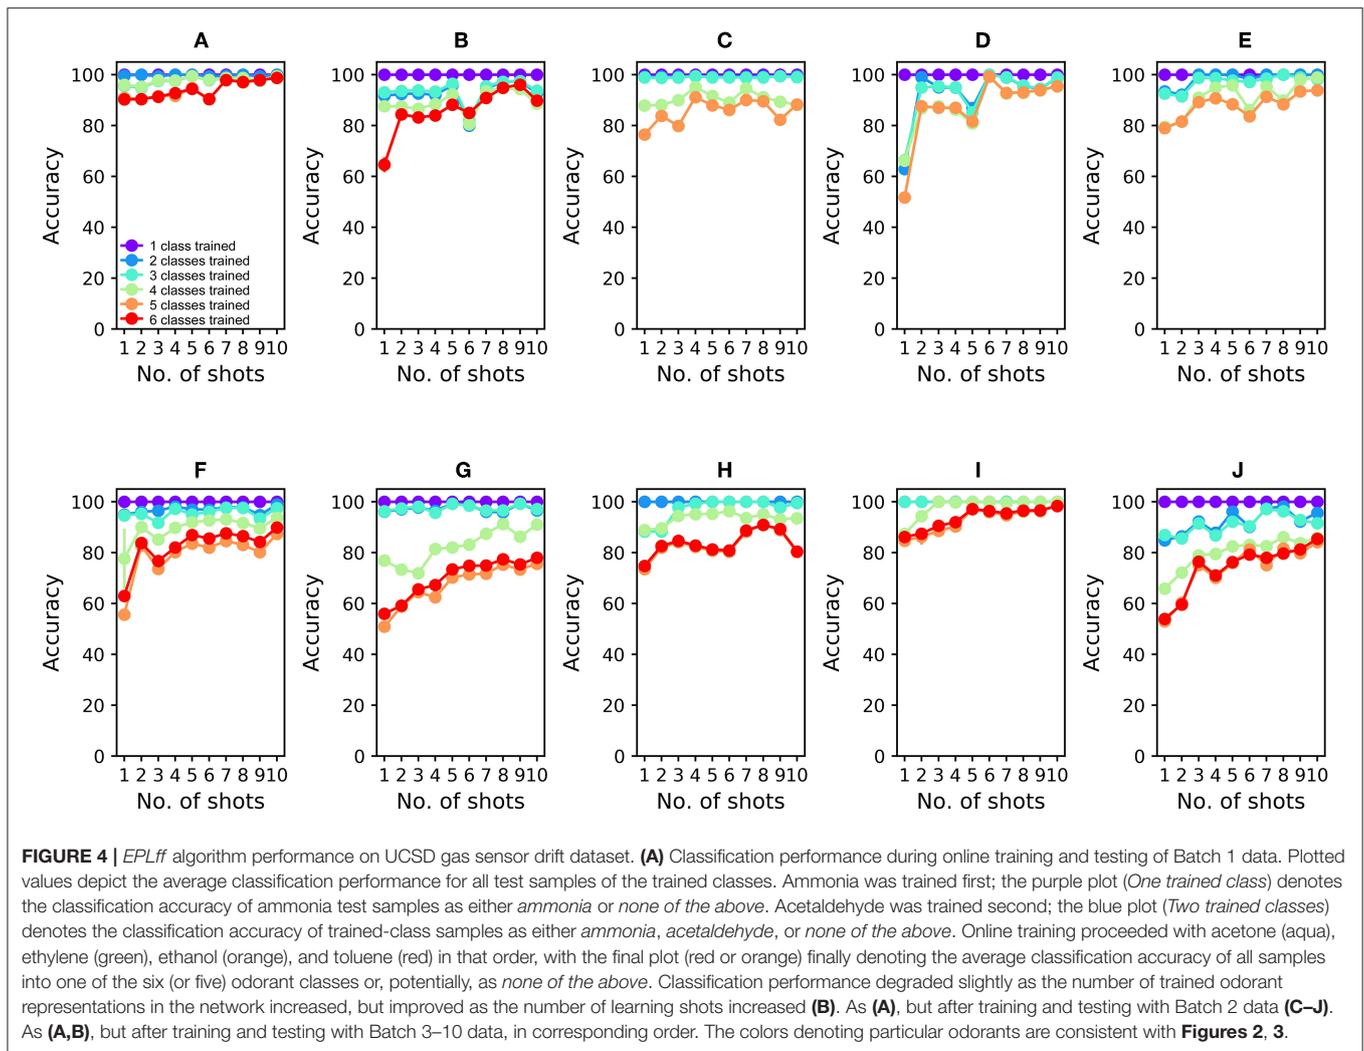
random selection of better or poorer class exemplars for training (particularly noting the uncontrolled variable of concentration) exerted a measurable effect on performance (**Figure 4**; **Table 4**).

Batch 10 of the UCSD dataset poses a relatively challenging classification problem. To produce it, the sensors were intentionally degraded and contaminated by turning off sensor heating for 5 months following the production of Batch 9 data (Vergara et al., 2012). Prior work with this dataset has achieved up to 73.28% classification performance on Batch 10, without online learning and using a highly introspective approach tailored for this specific dataset (Yan et al., 2017). In contrast, 10-shot learning on Batch 10 using the present *EPLff* algorithm achieved 85.43% classification accuracy.

To compare the *EPLff* network's resistance to catastrophic forgetting against an existing standard method, we built a 16-input multi-layer perceptron (MLP) comprising 16 input units for raw sensor input (ReLU activation), 4,800 hidden units (ReLU activation), and six output units for odorant classification. The MLP was trained using the Adam optimizer (Kingma and Ba, 2014) with a constant learning rate of 0.001. Since there was no straightforward way of implementing *none of the above* in an MLP, the MLP was only trained using two or more odorants (**Figure 5**). After initial, interspersed training on two odorants from Batch 1, the MLP classified test odorants at high accuracy ($99.41 \pm 0.0\%$; average of three runs; **Figure 5A**). However, its classification accuracy dropped sharply after the subsequent, sequential learning of odorant 3 ($30.61 \pm 0.0\%$ accuracy), odorant 4 ($16.24 \pm 9.29\%$), odorant 5 ($18.13 \pm 0.0\%$), and odorant 6 ($15.99 \pm 0.0\%$) (**Figure 5**). Catastrophic forgetting is a well-known limitation of MLPs, and is presented here simply to quantify the contrast in online learning performance between the *EPLff* implementation and a standard network of similar scale.

Online Reset Learning for Mitigating Sensor Drift

One of the most challenging problems of machine olfaction is *sensor drift*, in which the sensitivity and selectivity profiles of chemosensors gradually change over weeks to months of use or disuse. Efforts to compensate for this drift have taken many forms, from simply replacing sensors to designing highly introspective or specific corrective algorithms. For example,



one approach requires the non-random, algorithmically guided selection of relevant samples across batches and/or the utilization of test data as unlabeled data for additional training (Zhang and Zhang, 2015; Yan et al., 2017; Ma et al., 2018). Despite some partial successes in these approaches, the real-world challenge of sensor drift is a fundamentally ill-posed problem, in which the rapidity and nature of functional drift is highly dependent on the idiosyncratic chemistry of individual sensors and specific sensor-analyte pairs.

We argue that the most practical solution to this challenge is to retrain the network as needed to maintain performance, leveraging its rapid, online learning capacity. Specifically, MC-GC synaptic weights are simply reset to their untrained values and the network then is rapidly retrained using the new (degraded) sensor response profiles (*reset learning*). Retraining is not a new approach, of course, but overtly choosing a commitment to heuristic retraining as the primary method for countering sensor drift is important, as it determines additional criteria for real-world device functionality that candidate solutions must address, such as the need for rapid, ideally online

retraining in the field and potentially a tolerance for lower-fidelity training sets. Specifically, retraining a traditional classification network may require:

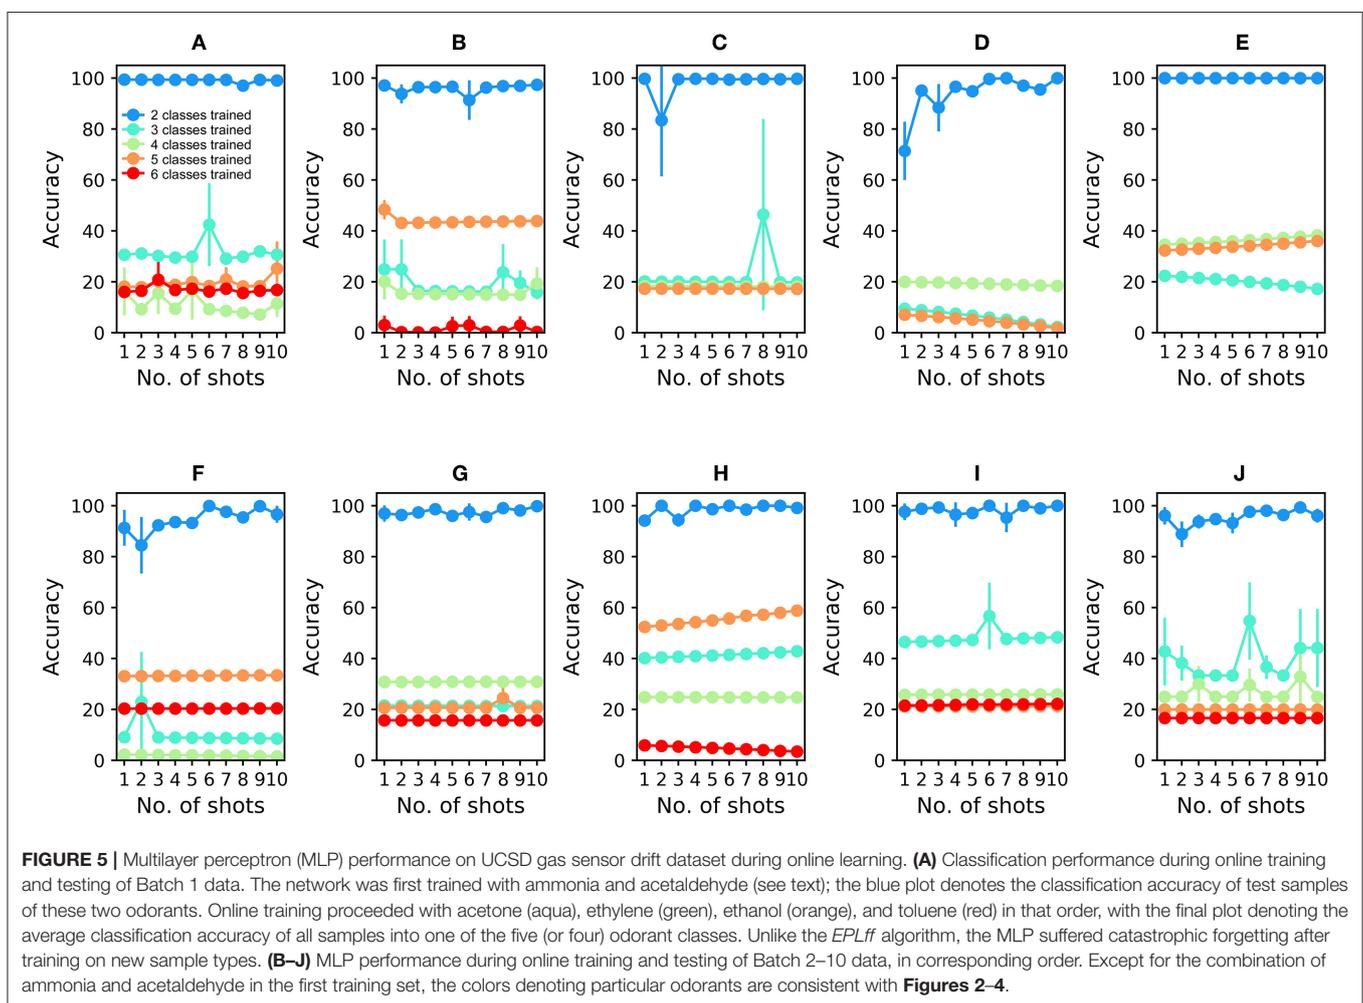
1. Prior knowledge of the number of possible odor classes to be identified,
2. A sufficiently large and representative training set incorporating each of these classes,
3. The retuning of network hyperparameters to match the altered characteristics of the degraded sensors, requiring an indeterminate number of training iterations.

The *EPL* network is not constrained by the above requirements. As demonstrated above, it can be rapidly retrained using small samples of whatever training sets are available and then be updated thereafter—including the subsequent introduction of new classes. The storage of training data for retraining purposes is unnecessary as the network does not suffer from catastrophic forgetting. Finally, the present network does not require hyperparameter retuning. Here, only the MC-GC weights were updated during retraining

TABLE 4 | Mean *EPLff* classification accuracies across all test odorants on the UCSD drift data set by the Hamming distance criterion.

	Batch 1	Batch 2	Batch 3	Batch 4	Batch 5	Batch 6	Batch 7	Batch 8	Batch 9	Batch 10	Average
1 shot	95.42	83.62	92.48	69.43	88.89	80.97	79.37	87.59	93.04	74.11	84.49
2 shot	95.15	90.38	93.97	93.68	89.35	91.35	80.98	90.41	94.68	77.36	89.73
3 shot	96.28	89.91	93.60	93.02	95.78	87.30	82.98	93.52	96.51	85.72	91.46
4 shot	96.62	90.29	97.02	92.64	96.87	91.25	83.89	93.26	97.02	82.58	92.14
5 shot	97.99	93.47	95.54	86.94	96.40	92.42	87.35	92.89	99.05	87.19	92.92
6 shot	96.00	85.44	94.70	99.69	93.01	92.22	87.77	92.93	98.73	87.36	92.78
7 shot	98.80	94.22	96.54	96.66	97.03	93.47	87.79	95.11	98.30	88.39	94.63
8 shot	98.59	96.45	95.79	95.48	95.64	92.78	89.43	96.16	98.81	90.34	94.95
9 shot	98.39	96.92	94.11	95.35	98.06	90.37	88.92	94.84	98.82	88.32	94.41
10 shot	99.39	92.44	94.95	97.73	98.22	94.55	89.74	92.30	99.48	90.46	94.93

Odorant-specific classification accuracies are depicted in **Figure 4**.



(using the same STDP rule); sensor scaling factors and all other parameters were ascertained once, using the 10% validation set of Batch 1, and held constant thereafter. Moreover, the *none of the above* classifier confidence feature facilitates awareness of when the network may require retraining; an increase in *none of the above* classifications

provides an initial cue that then can be evaluated using known samples.

To assess the efficacy of this approach, we tested the *EPLff* algorithm on the UCSD dataset framed as a sensor drift problem. The procedure for this approach, and consequently the results, are identical to those of section Online Learning above (**Figure 4**;

Table 4). Importantly, the sensor scaling factors and network parameters were tuned only once, using the validation set from Batch 1, on the theory that the concept of rapid reset was incompatible with a strategy of re-optimizing multiple network hyperparameters. Hence, no parameter changes were permitted, other than the MC-GC excitatory synaptic weights that were updated normally during training according to the STDP rule (In order to avoid duplication of figures, this constraint was observed in the simulations of section Online Learning as well). As described above (**Figure 4**), Batch 1 training samples from all six odorants again were presented to the network in an online learning configuration, and classification performance then was assessed by Batch 1 test data. MC-GC synaptic weights then were reset to the default values (the *reset*), after which Batch 2 training samples were presented to the network in the same manner, followed by testing with Batch 2 test data including all odorants and concentrations. We repeated this process for batches 3–10. We also assessed post-reset classification performance across all batches based on a maximally rapid reset (i.e., one-shot learning) and compared this to performance after expanded training protocols up through 10-shot learning. All classification performance results (averaged across three full repeats each) are depicted in **Figure 4** and **Table 4**. In general, while modest increases in classification accuracy were observed when the training set size was larger, these results demonstrate scalability, showing that the *EPLff* algorithm classifies large sets of test data with reasonable accuracy even based on small training sets and lacking control over the concentrations of presented odorants.

DISCUSSION

We present a neural network algorithm that achieves superior classification performance in an online learning setting while not being specifically tuned to the statistics of any particular dataset. This property, coupled with its few-shot learning capacity and SNN architecture, renders it particularly appropriate for field-deployable devices based on learning-capable SNN hardware (Davies et al., 2018; Imam and Cleland, 2019), recognizing that the interim use of the Hamming distance for nearest-neighbor classification in the present *EPLff* framework will not be part of such a deployable system. This algorithm is inspired by the architecture of the mammalian olfactory bulb, but is comparably applicable to any high-dimensional dataset that lacks internal low-dimensional structure.

The present *EPLff* incarnation of the network utilizes one or more preprocessor algorithms to prepare data for effective learning and classification by the core network. Among these is an unsupervised concentration tolerance algorithm derived from feedback normalization models of the biological system (Cleland et al., 2007, 2012; Banerjee et al., 2015), a version of which has been previously instantiated in SNN hardware (Imam et al., 2012). Inclusion of this preprocessor enables our algorithm to quickly learn reliable representations based on few-shot learning from odorant samples presented at different and unknown concentrations. Moreover, the network then can generalize

across concentrations, correctly classifying unknown test odorants presented at concentrations on which the network was never trained, and even estimating the concentrations of these unknowns.

The subsequent, plastic EPL layer of the network is based on a high-dimensional projection of sensory input data onto a network of interneurons known as granule cells (GCs). In the present feed-forward implementation, our emphasis is on the roles and capacities of two sequential preprocessor steps followed by the STDP-driven plasticity of the excitatory MC-GC synapses. Subsequent extensions of this work will restore the feedback architecture of the original model (Imam and Cleland, 2019) while enabling a more sophisticated development of learned classes within the high-dimensional projection field. Even in its present feedforward form, however, the *EPLff* algorithm exhibits (1) rapid, online learning of arbitrary sensory representations presented in arbitrary sequences, (2) generalization across concentrations, (3) robustness to substantial changes in the diversity and responsivity of sensor array input without requiring network reparameterization, and, by virtue of these properties, is capable of (4) effective adaptation to ongoing sensor drift via a rapid reset-and-retraining process termed reset learning. This capacity for fast reset learning represents a practical strategy for field-deployable devices, in which a training sample kit could be quickly employed in the field to retune and restore functionality to a device in which the sensors may have degraded. Importantly for such purposes, the *EPLff* algorithm was not, and need not be, crafted to the statistics of any particular data set, nor was the network pre-exposed to testing set data as has been done in some approaches (Zhang and Zhang, 2015; Yan et al., 2017).

Because field-deployable devices require a level of generic readiness for undetermined or underdetermined problems, and these *EPLff* properties favor such readiness, we have emphasized the portability of these algorithms to neuromorphic hardware platforms that may come to drive such devices. Interestingly, many of the features of the biological olfactory system that have inspired this design are appropriate for such devices. Spike timing and event-based algorithms are attractive candidates for compact, energy-efficient hardware implementation (Imam et al., 2012; Merolla et al., 2014; Qiao et al., 2015; Diehl et al., 2016; Esser et al., 2016; Davies et al., 2018). Spike timing metrics can compute similar transformations as analog and rate-based representations; indeed, it has been proposed that spike based computations could in principle exhibit all of the computational power of a universal Turing machine (Maass, 1996, 2015). STDP is a localized learning algorithm that is highly compatible with the colocalization of memory and compute principle of neuromorphic design, and its theoretical capacities have been thoroughly explored in diverse relevant contexts (Nessler et al., 2009; Linster and Cleland, 2010; Schmiedt et al., 2010; Bengio et al., 2015; O'Connor et al., 2018). Our biologically constrained approach to algorithm design also provides a unified and empirically verified framework to investigate the interactions of these various algorithms and information metrics, to better interpret and apply them to artificial network design.

Other groups have previously proposed networks for gas sensor data analysis inspired by biological olfactory systems. Models of olfactory bulb and piriform cortical activity have been applied to analyze chemosensor array data (Raman and Gutierrez-Osuna, 2005; Raman et al., 2006). Algorithms based on the insect olfactory system have been employed to learn and identify odor-like inputs (Diamond et al., 2016; Delahunt et al., 2018) as well as to identify handwritten digits—visual inputs incorporating additional low-dimensional structure (Huerta and Nowotny, 2009; Delahunt and Kutz, 2018; Diamond et al., 2019). More broadly, insect mushroom bodies in particular have been deeply studied in terms of both their pattern separation and associative learning capacities (Hige, 2018; Cayco-Gajic and Silver, 2019). These capacities potentiate one another in service to odor learning and the classification of learned odor-like signals, though they also have been applied to more complex tasks (Ardin et al., 2016; Peng and Chittka, 2017). In the present work, we sought to design artificial learning networks to replicate some of the most powerful capabilities of the biological olfactory system, in particular its capacity for rapid online learning and the fast and effective classification of learned odorants despite ongoing changes in sensor properties and the unpredictability of odor concentrations. Future work will extend this framework to incorporate the

feedback dynamics of the biological system, increase the dimensionality of sensor arrays, and develop more sophisticated biomimetic classifiers.

AUTHOR CONTRIBUTIONS

TC originally conceived the algorithm, which was vetted and modified for present purposes by AB and TC. AB designed, programmed, and performed the simulations. AB and TC designed the figures and wrote the paper.

FUNDING

This work was supported by a Cornell University Sage fellowship to AB and an Intel Neuromorphic Research Community faculty award and NIH/NIDCD awards DC014367 and DC014701 to TC.

ACKNOWLEDGMENTS

The authors acknowledge Dr. Nabil Imam for interesting discussions regarding the *EPLff* algorithm, and Dr. Ramon Huerta and Dr. Jordi Fonollosa for discussions regarding the UCSD Gas Sensor Array Drift Dataset.

REFERENCES

- Ardin, P., Peng, F., Mangan, M., Lagogiannis, K., and Webb, B. (2016). Using an insect mushroom body circuit to encode route memory in complex natural environments. *PLoS Comput. Biol.* 12:e1004683. doi: 10.1371/journal.pcbi.1004683
- Banerjee, A., Marbach, F., Anselmi, F., Koh, M. S., Davis, M. B., Garcia da Silva, P., et al. (2015). An interglomerular circuit gates glomerular output and implements gain control in the mouse olfactory bulb. *Neuron* 87, 193–207. doi: 10.1016/j.neuron.2015.06.019
- Becherelli, R., Zampetti, E., Pantalei, S., Bernabei, M., and Persaud, K. C. (2010). Design of a very large chemical sensor system for mimicking biological olfaction. *Sens. Actuators B Chem.* 146, 446–452. doi: 10.1016/j.snb.2009.11.031
- Bengio, Y., Lee, D.-H., Bornschein, J., Mesnard, T., and Lin, Z. (2015). Towards biologically plausible deep learning. *arXiv:1502.04156* [cs]. Available online at: <http://arxiv.org/abs/1502.04156> (accessed May 13, 2018).
- Borthakur, A., and Cleland, T. A. (2017). “A neuromorphic transfer learning algorithm for orthogonalizing highly overlapping sensor array responses,” in *2017 ISOS/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)* (Montreal, QC), 1–3.
- Cayco-Gajic, N. A., and Silver, R. A. (2019). Re-evaluating circuit mechanisms underlying pattern separation. *Neuron* 101, 584–602. doi: 10.1016/j.neuron.2019.01.044
- Cleland, T. A. (2014). Construction of odor representations by olfactory bulb microcircuits. *Prog. Brain Res.* 208, 177–203. doi: 10.1016/B978-0-444-63350-7.00007-3
- Cleland, T. A., Chen, S.-Y. T., Hozer, K. W., Ukatu, H. N., Wong, K. J., and Zheng, F. (2012). Sequential mechanisms underlying concentration invariance in biological olfaction. *Front. Neuroeng.* 4:21.
- Cleland, T. A., Johnson, B. A., Leon, M., and Linster, C. (2007). Relational representation in the olfactory system. *Proc. Natl. Acad. Sci. U.S.A.* 104, 1953–1958. doi: 10.1073/pnas.0608564104
- Cleland, T. A., and Sethupathy, P. (2006). Non-topographical contrast enhancement in the olfactory bulb. *BMC Neurosci.* 7:7. doi: 10.1186/1471-2202-7-7
- Dan, Y., and Poo, M.-M. (2004). Spike timing-dependent plasticity of neural circuits. *Neuron* 44, 23–30. doi: 10.1016/j.neuron.2004.09.007
- Davies, M., Srinivasa, N., Lin, T. H., Chinya, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Delahunt, C. B., and Kutz, J. N. (2018). *A Moth Brain Learns to Read MNIST*. Available online at: <https://openreview.net/forum?id=HyYuqoCUz> (accessed May 13, 2018).
- Delahunt, C. B., Riffell, J. A., and Kutz, J. N. (2018). Biological mechanisms for learning: a computational model of olfactory learning in the *Manduca sexta* moth, with applications to neural nets. *Front. Comput. Neurosci.* 12:102. doi: 10.3389/fncom.2018.00102
- Diamond, A., Schmuker, M., Berna, A. Z., Trowell, S., and Nowotny, T. (2016). Classifying continuous, real-time e-nose sensor data using a bio-inspired spiking network modelled on the insect olfactory system. *Bioinspir. Biomim.* 11:026002. doi: 10.1088/1748-3190/11/2/026002
- Diamond, A., Schmuker, M., and Nowotny, T. (2019). An unsupervised neuromorphic clustering algorithm. *Biol. Cybern.* doi: 10.1007/s00422-019-00797-7
- Diehl, P. U., Pedroni, B. U., Cassidy, A., Merolla, P., Neftci, E., and Zarella, G. (2016). TrueHappiness: neuromorphic emotion recognition on truenorth. *arXiv:1601.04183* [q-bio, cs]. Available online at: <http://arxiv.org/abs/1601.04183> (accessed May 13, 2018).
- Esser, S. K., Merolla, P. A., Arthur, J. V., Cassidy, A. S., Appuswamy, R., Andreopoulos, A., et al. (2016). Convolutional networks for fast, energy-efficient neuromorphic computing. *Proc. Natl. Acad. Sci. U.S.A.* 113, 11441–11446. doi: 10.1073/pnas.1604850113
- Gire, D. H., Franks, K. M., Zak, J. D., Tanaka, K. F., Whitesell, J. D., Mulligan, A. A., et al. (2012). Mitral cells in the olfactory bulb are mainly excited through a multistep signaling path. *J. Neurosci.* 32, 2964–2975. doi: 10.1523/JNEUROSCI.5580-11.2012
- Gonzalez, J., Monroy, J. G., Garcia, F., and Blanco, J. L. (2011). “The multi-chamber electronic nose (MCE-nose),” in *2011 IEEE International Conference on Mechatronics (Istanbul)*, 636–641.

- Hige, T. (2018). What can tiny mushrooms in fruit flies tell us about learning and memory? *Neurosci. Res.* 129, 8–16. doi: 10.1016/j.neures.2017.05.002
- Huerta, R., and Nowotny, T. (2009). Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural Comput.* 21, 2123–2151. doi: 10.1162/neco.2009.03-08-733
- Imam, N., and Cleland, T. A. (2019). Rapid online learning and robust recall in a neuromorphic olfactory circuit. *arXiv:1906.07067 Cs Q-Bio*. Available online at: <http://arxiv.org/abs/1906.07067> (accessed June 17, 2019).
- Imam, N., Cleland, T. A., Manohar, R., Merolla, P. A., Arthur, J. V., Akopyan, F., et al. (2012). Implementation of olfactory bulb glomerular-layer computations in a digital neurosynaptic core. *Front. Neurosci.* 6:83. doi: 10.3389/fnins.2012.00083
- Iskierko, Z., Sharma, P. S., Bartold, K., Pietrzyk-Le, A., Noworyta, K., and Kutner, W. (2016). Molecularly imprinted polymers for separating and sensing of macromolecular compounds and microorganisms. *Biotechnol. Adv.* 34, 30–46. doi: 10.1016/j.biotechadv.2015.12.002
- Kashiwadani, H., Sasaki, Y. F., Uchida, N., and Mori, K. (1999). Synchronized oscillatory discharges of mitral/tufted cells with different molecular receptive ranges in the rabbit olfactory bulb. *J. Neurophysiol.* 82, 1786–1792. doi: 10.1152/jn.1999.82.4.1786
- Kemker, R., and Kanan, C. (2017). FearNet: brain-inspired model for incremental learning. *arXiv:1711.10563 [cs]*. Available online at: <http://arxiv.org/abs/1711.10563> (accessed January 24, 2019).
- Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. *arXiv:1412.6980 [cs]*. Available online at: <http://arxiv.org/abs/1412.6980> (accessed January 30, 2019).
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proc. Natl. Acad. Sci.* 114, 3521–3526. doi: 10.1073/pnas.1611835114
- Länge, K., Rapp, B. E., and Rapp, M. (2008). Surface acoustic wave biosensors: a review. *Anal. Bioanal. Chem.* 391, 1509–1519. doi: 10.1007/s00216-008-1911-5
- Lepousez, G., and Lledo, P.-M. (2013). Odor discrimination requires proper olfactory fast oscillations in awake mice. *Neuron* 80, 1010–1024. doi: 10.1016/j.neuron.2013.07.025
- Lepousez, G., Nissant, A., Bryant, A. K., Gheusi, G., Greer, C. A., and Lledo, P.-M. (2014). Olfactory learning promotes input-specific synaptic plasticity in adult-born neurons. *Proc. Natl. Acad. Sci. U.S.A.* 111, 13984–13989. doi: 10.1073/pnas.1404991111
- Li, G., and Cleland, T. A. (2017). A coupled-oscillator model of olfactory bulb gamma oscillations. *PLoS Comput. Biol.* 13:e1005760. doi: 10.1371/journal.pcbi.1005760
- Linster, C., and Cleland, T. A. (2010). Decorrelation of odor representations via spike timing-dependent plasticity. *Front. Comput. Neurosci.* 4:157. doi: 10.3389/fncom.2010.00157
- Liu, Q., Cai, H., Xu, Y., Li, Y., Li, R., and Wang, P. (2006). Olfactory cell-based biosensor: a first step towards a neurochip of bioelectronic nose. *Biosens. Bioelectron.* 22, 318–322. doi: 10.1016/j.bios.2006.01.016
- Liu, Q., Hu, X., Ye, M., Cheng, X., and Li, F. (2015). Gas recognition under sensor drift by using deep learning. *Int. J. Intell. Syst.* 30, 907–922. doi: 10.1002/int.21731
- Ma, Z., Luo, G., Qin, K., Wang, N., and Niu, W. (2018). Online sensor drift compensation for E-nose systems using domain adaptation and extreme learning machine. *Sensors* 18:E742. doi: 10.3390/s18030742
- Maass, W. (1996). Lower bounds for the computational power of networks of spiking neurons. *Neural Comput.* 8, 1–40. doi: 10.1162/neco.1996.8.1.1
- Maass, W. (2015). To spike or not to spike: that is the question. *Proc. IEEE* 103, 2219–2224. doi: 10.1109/JPROC.2015.2496679
- Mandairon, N., Kuczewski, N., Kermen, F., Forest, J., Midroit, M., Richard, M., et al. (2018). Opposite regulation of inhibition by adult-born granule cells during implicit versus explicit olfactory learning. *eLife* 7:e34976. doi: 10.7554/eLife.34976
- McCloskey, M., and Cohen, N. J. (1989). Catastrophic interference in connectionist networks: the sequential learning problem. *Psychol. Learn. Motiv.* 24, 109–165. doi: 10.1016/S0079-7421(08)60536-8
- Mehta, D., Altan, E., Chandak, R., Raman, B., and Chakrabarty, S. (2017). “Behaving cyborg locusts for standoff chemical sensing,” in *2017 IEEE International Symposium on Circuits and Systems (ISCAS)* (Baltimore, MD), 1–4.
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Nagayama, S., Homma, R., and Imamura, F. (2014). Neuronal organization of olfactory bulb circuits. *Front. Neural Circuits* 8:98. doi: 10.3389/fncir.2014.00098
- Nessler, B., Pfeiffer, M., and Maass, W. (2009). “STDP enables spiking neurons to detect hidden causes of their inputs,” in *Advances in Neural Information Processing Systems 22*, eds Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta (Curran Associates, Inc.), 1357–1365. Available online at: <http://papers.nips.cc/paper/3744-stdp-enables-spiking-neurons-to-detect-hidden-causes-of-their-inputs.pdf> (accessed May 16, 2018).
- O’Connor, P., Gavves, E., Reisser, M., and Welling, M. (2018). *Temporally Efficient Deep Learning with Spikes*. Available online at: <https://openreview.net/forum?id=HyYuoCUz> (accessed May 18, 2018).
- Panzeri, S., Brunel, N., Logothetis, N. K., and Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends Neurosci.* 33, 111–120. doi: 10.1016/j.tins.2009.12.001
- Peace, S. T., Johnson, B. C., Li, G., Kaiser, M. E., Fukunaga, I., Schaefer, A. T., et al. (2017). Coherent olfactory bulb gamma oscillations arise from coupling independent columnar oscillators. *bioRxiv* 213827. doi: 10.1101/213827
- Peng, F., and Chittka, L. (2017). A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory. *Curr. Biol.* 27, 224–230. doi: 10.1016/j.cub.2016.10.054
- Persaud, K., and Dodd, G. (1982). Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature* 299, 352–355. doi: 10.1038/299352a0
- Persaud, K. C., Marco, S., and Gutiérrez-Gálvez, A. (eds.). (2013). *Neuromorphic Olfaction*. Boca Raton: CRC Press/Taylor & Francis.
- Qiao, N., Mostafa, H., Corradi, F., Osswald, M., Stefanini, F., Sumislawska, D., et al. (2015). A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses. *Front. Neurosci.* 9:141. doi: 10.3389/fnins.2015.00141
- Raman, B., and Gutierrez-Osuna, R. (2005). “Mixture segmentation and background suppression in chemosensor arrays with a model of olfactory bulb-cortex interaction,” in *Proceedings 2005 IEEE International Joint Conference on Neural Networks* (Montreal, QC), Vol. 1, 131–136.
- Raman, B., Sun, P. A., Gutierrez-Galvez, A., and Gutierrez-Osuna, R. (2006). Processing of chemical sensor arrays with a biologically inspired model of olfactory coding. *IEEE Trans. Neural Netw.* 17, 1015–1024. doi: 10.1109/TNN.2006.875975
- Rodríguez-Lujan, I., Fonollosa, J., Vergara, A., Homer, M., and Huerta, R. (2014). On the calibration of sensor arrays for pattern recognition using the minimal number of experiments. *Chemom. Intell. Lab. Syst.* 130, 123–134. doi: 10.1016/j.chemolab.2013.10.012
- Schmiedt, J., Albers, C., and Pawelzik, K. (2010). “Spike timing-dependent plasticity as dynamic filter,” in *Advances in Neural Information Processing Systems 23*, eds J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta (Curran Associates, Inc.), 2110–2118. Available online at: <http://papers.nips.cc/paper/3917-spike-timing-dependent-plasticity-as-dynamic-filter.pdf> (accessed May 18, 2018).
- Schmuker, M., Nawrot, M., and Chicca, E. (2015). “Neuromorphic sensors, olfaction,” in *Encyclopedia of Computational Neuroscience*, eds D. Jaeger and R. Jung (New York, NY: Springer), 1991–1997.
- 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
- Serrà, J., Suris, D., Miron, M., and Karatzoglou, A. (2018). Overcoming catastrophic forgetting with hard attention to the task. *arXiv:1801.01423 [cs, stat]*. Available online at: <http://arxiv.org/abs/1801.01423> (accessed January 24, 2019).
- Serrano, E., Nowotny, T., Levi, R., Smith, B. H., and Huerta, R. (2013). Gain control network conditions in early sensory coding. *PLoS Comput. Biol.* 9:e1003133. doi: 10.1371/journal.pcbi.1003133
- Shi, H., Tsai, W.-B., Garrison, M. D., Ferrari, S., and Ratner, B. D. (1999). Template-imprinted nanostructured surfaces for protein recognition. *Nature* 398, 593–597. doi: 10.1038/19267

- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Velez, R., and Clune, J. (2017). Diffusion-based neuromodulation can eliminate catastrophic forgetting in simple neural networks. *PLoS ONE* 12:e0187736. doi: 10.1371/journal.pone.0187736
- Vergara, A., Vembu, S., Ayhan, T., Ryan, M. A., Homer, M. L., and Huerta, R. (2012). Chemical gas sensor drift compensation using classifier ensembles. *Sens. Actuators B Chem.* 166–167, 320–329. doi: 10.1016/j.snb.2012.01.074
- Xiong, W., and Chen, W. R. (2002). Dynamic gating of spike propagation in the mitral cell lateral dendrites. *Neuron* 34, 115–126. doi: 10.1016/S0896-6273(02)00628-1
- Yan, K., Zhang, D., and Xu, Y. (2017). Correcting instrumental variation and time-varying drift using parallel and serial multitask learning. *IEEE Trans. Instrum. Meas.* 66, 2306–2316. doi: 10.1109/TIM.2017.2707898
- Yin, H., Wang, Z., and Jha, N. (2018). A hierarchical inference model for internet-of-things. *IEEE Trans. Multi-Scale Comput. Syst.* 4, 260–271. doi: 10.1109/TMSCS.2018.2821154
- Zaidi, Q., Victor, J., McDermott, J., Geffen, M., Bensmaia, S., and Cleland, T. A. (2013). Perceptual spaces: mathematical structures to neural mechanisms. *J. Neurosci.* 33, 17597–17602. doi: 10.1523/JNEUROSCI.3343-13.2013
- Zenke, F., Poole, B., and Ganguli, S. (2017). Continual learning through synaptic intelligence. *arXiv:1703.04200* [cs, q-bio, stat]. Available online at: <http://arxiv.org/abs/1703.04200> (accessed January 24, 2019).
- Zhang, L., and Zhang, D. (2015). Domain adaptation extreme learning machines for drift compensation in E-nose systems. *IEEE Trans. Instrum. Meas.* 64, 1790–1801. doi: 10.1109/TIM.2014.2367775
- Zhang, Y., Zhao, J., Du, T., Zhu, Z., Zhang, J., and Liu, Q. (2017). A gas sensor array for the simultaneous detection of multiple VOCs. *Sci. Rep.* 7:1960. doi: 10.1038/s41598-017-02150-z

Conflict of Interest Statement: Both authors are listed as inventors on a Cornell University provisional patent (8631-01-US) covering other aspects of this algorithm.

Copyright © 2019 Borthakur and Cleland. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.