

# A Neuromorphic Transfer Learning Algorithm for Orthogonalizing Highly Overlapping Sensor Array Responses

Ayon Borthakur

Cornell University, Ithaca, USA 14853

Email: ab2535@cornell.edu

Thomas A. Cleland

Cornell University, Ithaca, USA 14853

Email: tac29@cornell.edu

**Abstract**—In biological and artificial noses, similar odorants activate substantially overlapping populations of sensors. Increasing the discriminability of such odorants through learning is challenging, as learning-related plasticity in neurons activated by one such odorant then may inappropriately alter the representations of similar, overlapping odorants. We here describe a transfer learning algorithm, based on mammalian olfactory bulb architecture and asymmetric STDP, in which interneurons construct experience-dependent higher-order receptive fields (HORFs) that become diagnostic of more complex odorant signatures (covariance patterns). Consequently, even highly similar odorants can evoke nonoverlapping ensembles of interneurons, and thereby independently regulate principal neuron spike timing patterns and prevent the cross-contamination of plasticity. HORF construction is regulated by parameters such as maximum synaptic weight, initial connection probabilities, and the interneuron membrane time constant and spike threshold. This core algorithm can in principle allocate inhibition so as to regulate generalization and discrimination among specific odorant representations.

## I. INTRODUCTION

Chemosensory identification in the real world is fundamentally a problem of variance management. The chemosignals emitted by sources of interest are variable in quality and in concentration, and their sensory representations overlap substantially, occluding one another in unpredictable ways. The observation that biological chemosensory (olfactory) systems enable robust source segregation in complex environments is the basis for current interest in neuromorphic chemosensory systems. A useful algorithm must be able to discriminate between very similar chemosignals when warranted while also generalizing across the variability of emissions from the same source. Correctly resolving this challenge requires an adaptive categorization process based upon experience (transfer learning [1]). Moreover, when multiple chemosignals are mixed, their sources often must be categorically identified based on highly occluded sensor responses in which none of the source signals are clearly recognizable.

The mammalian olfactory system employs multiple coordinated algorithms to help resolve difficult problems such as the segregation of concentration-dependent variance from genuine quality-dependent variance, a necessary prerequisite for concentration tolerance in olfaction [2]. Glomerular-layer algorithms in the olfactory bulb sharpen high-dimensional olfactory sensory representations [3], but are not well suited to iden-

tify or regulate higher-order correlations among chemosensory representations, as would be required for adaptive categorization. However, the external plexiform layer (EPL) network within olfactory bulb, which mediates lateral inhibition among principal neurons (mitral cells; MCs) via inhibitory granule cell interneurons (GCs), is well suited for this task. Notably, the topology of lateral synaptic weights is not correlated with physical proximity. We propose that the topology of EPL lateral inhibition is based on learning, and here present a tractable, data-constrained transfer learning algorithm for functional plasticity in the EPL that enables experience-dependent source identification. The allocation of inhibition based on higher-order correlations also can reduce the cross-contamination of plasticity (interference), in which learning about an odorant improperly alters the representations of similar odorants owing to their shared resources.

## II. MODEL DESCRIPTION

In the EPL, extensive MC lateral dendrites support action potential propagation and deliver synaptic excitation onto a large population of GC interneurons distributed across the entire extent of the OB. These GCs, in turn, deliver divisive (shunting) synaptic inhibition onto MCs. Whereas regenerative excitation propagates losslessly along these lateral dendrites, the propagation of inhibition is limited by cable properties [4]. Accordingly, in the model, all MCs and GCs are associated with a specific column (these columns in turn correspond to specific odorant receptors and their associated olfactory bulb glomeruli); MCs excite an initially-random population of GCs drawn from the entire model EPL, whereas GCs inhibit MCs only within their own column (presently with nonplastic synapses). The distribution of MC-GC excitatory synaptic weights is altered by unsupervised learning as described below.

In the model, information in MC activity patterns is represented by a spike timing-based metric – specifically, a spike phase precedence code [5] based on a common underlying  $\gamma$ - $\beta$  periodicity that can be observed in field potential oscillations. The reasons for proposing this framework and the biophysical principles on which it is predicated are described elsewhere [4], [6]; briefly, its advantages in the present context include fast readout and a means to avoid recursively altering sensory input levels. In our spike timing model, earlier MC




Fig. 1. Generation and illustration of sequentially similar representations of odorants A, B, C, D, and E. *Control* is a non-overlapping (dissimilar) odorant generated by randomly reassigned receptor activities. See text for details.




Fig. 2. STDP learning causes both potentiation and depression of MC-to-GC synaptic weights. Note that many weights are increased prior to being reduced as the GC spike progressively phase-advances within the  $\gamma$ - $\beta$  cycle.

spike phases within  $\gamma$ - $\beta$  oscillation cycles correspond to stronger excitation, and EPL lateral inhibition delays MC spike times rather than preventing spike initiation or blocking spike propagation. [4]. Using timing-dependent learning rules (e.g. STDP), this metric can mediate with spike delays what rate-coding metrics achieve by silencing neuronal activity [6].

The evolution of neuronal membrane voltage over time is described by a first order differential equation typical of a LIF neuron:  $\tau \left( \frac{dV}{dt} \right) = -V + I \left( \frac{R_m}{R_{shunt}} \right)$ , where  $\tau = R_m C_m$ ,  $R_m$  is the neuronal input resistance,  $C_m$  is the membrane capacitance,  $\tau$  is the membrane time constant,  $V$  is the neuronal membrane potential, and  $R_{shunt}$  for MCs is the oscillatory shunting inhibition at  $\gamma$ - $\beta$  frequency (here  $40\text{Hz}$ ). For GCs, which do not receive inhibition,  $R_{shunt}$  is unity. We model the total synaptic input current to a GC as  $I = g_w (E_n - V)$ , where  $E_n$  is the Nernst potential and  $g_w = \sum_{i=1}^n w_i g_{max} \left( \frac{\tau_1 \tau_2}{\tau_1 - \tau_2} \right) \left( \exp \left( \frac{-(t-t_i)}{\tau_1} \right) - \exp \left( \frac{-(t-t_i)}{\tau_2} \right) \right)$  resembling the dynamics of AMPA synapses ( $\tau_1$ ,  $\tau_2$  are the rise and decay time constants respectively). We computed the overlap of neuronal activities between two odorants as the *normalized L<sub>2</sub> norm*  $= \frac{\sum_{i=1}^n s_{1i} s_{2i}}{\|s_1\| \|s_2\|}$ , where  $s_i$  is a vector whose length equals the total number of GCs, and each element corresponds to the total number of spikes evoked during the period of odorant presentation.

### III. THE GENERATION OF SEQUENTIALLY SIMILAR ODORANT REPRESENTATIONS

Odorant representations have a dimensionality equal to the number of receptor types (100 in the present model), and are defined by activity levels across this population of receptors. Accordingly, proximity in this 100-dimensional space corresponds to odorant similarity. Owing to the intrinsic variability of odorant sources, odorant source representations constitute manifolds (‘clouds’) within this high-dimensional

space. To generate sequentially similar odorant series, we added a small random variation to each of the odorant-receptor binding affinities (Figure 1a); specifically, *odorant B* = *odorant A* +  $X_1$ ; *odorant C* = *odorant B* +  $X_2$ , etc., where  $X_i$  is a vector whose elements are independently drawn from a skewed normal distribution (the skew served to counteract the accumulation of variance so that all odor representations exhibited similar activity distributions; Figure 1b). Because random walks in high-dimensional spaces reliably move away from the point of origin (Pólya’s theorem), this procedure generated odorant series A–E ordered in terms of decreasing similarity to A; non-overlapping control odorants also were generated (Figure 1c). The distances between neighboring odors A–E in these series were roughly constant (Figure 1d).

### IV. LEARNING-DEPENDENT CONSTRUCTION OF HIGHER-ORDER RECEPTIVE FIELDS IN GRANULE CELLS

We implemented a spike timing-based asymmetric STDP rule at all excitatory (MC to GC) synapses [6], [7]. Inhibitory plasticity was omitted. In this STDP paradigm, synapses in which presynaptic spikes precede postsynaptic spikes within a common  $\gamma$ - $\beta$  cycle are potentiated (LTP), whereas synapses in which presynaptic spikes follow postsynaptic spikes are depressed (LTD). Hence, MCs whose spikes contribute to evoking a spike in a common postsynaptic GC have their synapses potentiated, whereas those convergent MCs that fire too late to contribute to postsynaptic spike generation have their synapses weakened. Importantly, as the synaptic weights of contributing MCs increase, the phase of the postsynaptic GC spike advances, until it is able to be generated by a smaller number of stronger presynaptic inputs. This process results in the progressive refinement of GC response selectivity, culminating in a fixed sensitivity to a particular ensemble of  $k$  activated MCs (i.e., a higher-order receptive field, or HORF, of order  $k$ ; Figure 2). Via this “ $k$  winners take all” rule, GCs become selective for feature sets that are arbitrarily diagnostic of particular odorants. We found that the asymptotic order of GC HORFs depended upon the maximum permitted potentiation of individual synapses (amongst other parameters); lower potentiation caps yielded higher-order receptive field distributions (Figure 3) and reduced the overlap among GC population responses to similar odorants after learning. Indeed, an LTD-only learning rule produced the highest-order GC HORFs (Figure 3a). Initial network connection probabilities and the GC spiking threshold and membrane time constant also affected the asymptotic order of GC HORFs (not shown).




Fig. 3. Orders of GC receptive fields after asymptotic learning are reduced from  $\sim 20$  to  $\sim 8$  as the potentiation cap ( $w_{max}$ ) of MC-to-GC synapses is increased from panel *a* (no LTP; LTD only) to panels *b*, *c*, and *d*.  $w_{init}$  = initial synaptic weight. The ordinate depicts the number of maximized synaptic weights (where  $w = w_{max}$ ) connecting to a given GC after learning. The peaks around order  $\sim 40$  (*a*) and order = 0 (*b,c,d*) both depict GCs that did not spike, hence generating no synaptic plasticity (note that, in *a*,  $w_{max} = w_{init}$ , so these GCs have a spurious nonzero order).

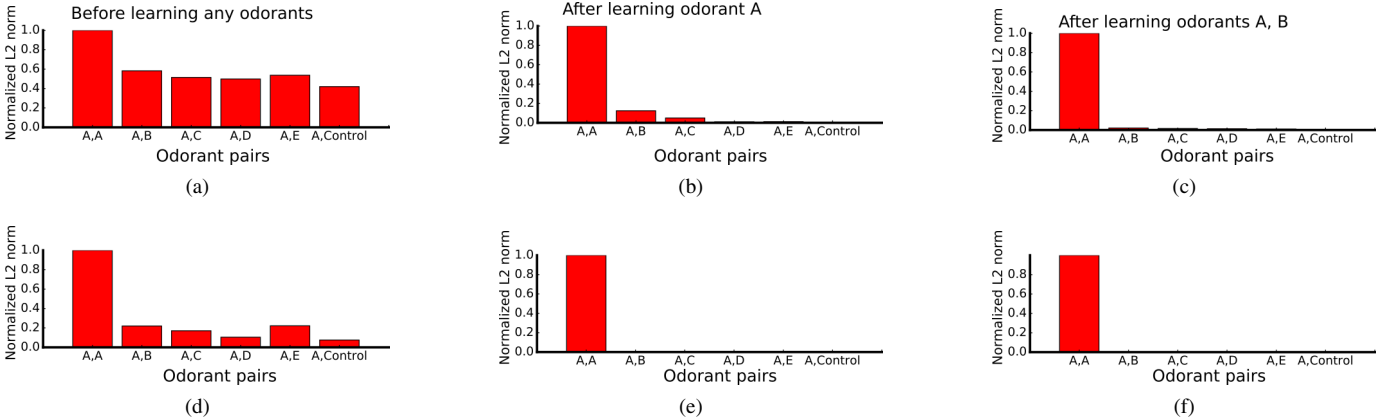


Fig. 4. Overlap of GC spiking activity representations between a given odorant *A* and a series of sequentially similar odorants *B-E* or a dissimilar control odorant. *Top row*: GC spike threshold of 0.35,  $\tau = 50 \mu s$ . A given amount of training on odorant *A* orthogonalizes its representation against moderately similar odorants *D* & *E*; additional training with odorant *B* further orthogonalizes *A* against *B* and less-similar odorants. *Bottom row*: GC spike threshold of 0.42,  $\tau = 64 \mu s$ . The same amount of learning under these more stringent parameters enables the GC representation of odorant *A* to be orthogonalized against all odorants tested; further training with odorant *B* then is superfluous.

The orthogonalizing capacity of the algorithm depended on both network properties and learning parameters (Figure 4). Training the network on odorant *A* progressively orthogonalized its representation with respect to similar odorants, up to asymptotic limits determined by network and synaptic properties (Figure 4b). Additional training on a similar odorant (*B*) could further orthogonalize the representations of *A* with respect to *B* and related odors (Figure 4c). In contrast, with different network and synaptic parameters, training on *A* alone could attain full orthogonality in its GC ensemble compared to the same set of similar odorants (Figure 4e). In sum, this learning algorithm progressively differentiates trained odorant representations from other, similar representations (existing or future), up to an asymptotic level determined by network and synaptic state variables. These properties determine the level of inductive transfer, or the balance between discrimination and generalization, to be associated with that odor representation.

## V. CONCLUSION

The vast coding space provided by the high dimensionality of olfactory systems potentially supports an enormous capacity for learning distinct odorant representations with their attendant variance profiles. The unsupervised learning at the MC-to-GC synapse described here enables similar odorants that activate highly overlapping ensembles of principal neurons to

develop non-overlapping activated populations of GC interneurons, thereby enabling the construction of independent, odor source-specific patterns of feedback inhibition.

## REFERENCES

- [1] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, ‘How transferable are features in deep neural networks?’, *ArXiv 14111792 Cs*, Nov. 2014.
- [2] T. A. Cleland, S.-Y. T. Chen, K. W. Hozer, H. N. Ukatu, K. J. Wong, and F. Zheng, ‘Sequential mechanisms underlying concentration invariance in biological olfaction’, *Front. Neuroengineering*, vol. 4, Jan. 2012.
- [3] T. A. Cleland and P. Sethupathy, ‘Non-topographical contrast enhancement in the olfactory bulb’, *BMC Neurosci.*, vol. 7, p. 7, 2006.
- [4] A. B. R. McIntyre and T. A. Cleland, ‘Biophysical constraints on lateral inhibition in the olfactory bulb’, *J. Neurophysiol.*, vol. 115, no. 6, pp. 2937-2949, Jun. 2016.
- [5] S. Panzeri, N. Brunel, N. K. Logothetis, and C. Kayser, ‘Sensory neural codes using multiplexed temporal scales’, *Trends Neurosci.*, vol. 33, no. 3, pp. 111-120, Mar. 2010.
- [6] C. Linster and T. A. Cleland, ‘Decorrelation of Odor Representations via Spike Timing-Dependent Plasticity’, *Front. Comput. Neurosci.*, vol. 4, Dec. 2010.
- [7] P. J. Sjöström, G. G. Turrigiano, and S. B. Nelson, ‘Rate, timing, and cooperativity jointly determine cortical synaptic plasticity’, *Neuron*, vol. 32, no. 6, pp. 1149-1164, Dec. 2001.