

Improving odor classification through self-organized lateral inhibition in a spiking olfaction-inspired network

Bahadır Kasap^{1,§} and Michael Schmuker^{1,2,*}

Abstract—In this study, we propose unsupervised learning of the lateral inhibition structure through inhibitory spike-timing dependent plasticity (iSTDP) in a computational model for multivariate data processing inspired by the honeybee antennal lobe. After exposing the network to a sufficient number of input samples, the inhibitory connectivity self-organizes to reflect the correlation between input channels. We show that this biologically realistic, local learning rule produces an inhibitory connectivity that effectively reduces channel correlation and yields superior network performance in a multivariate scent recognition scenario. The proposed network is suited as a preprocessing stage for spiking data processing systems, like for example neuromorphic hardware or neuronal interfaces.

I. INTRODUCTION

The insect olfactory system is capable of classifying odorants by encoding and processing the neural representations of chemical stimuli. The insect olfactory system is well described at a structural level [1]. In the antennae of insects, odors are transformed into a neuronal representation by a number of receptor classes, each of which encodes a certain combination of chemical features. Axons from each receptor class converge into separate compartments (so-called glomeruli) in the antennal lobe (AL). The activity pattern of those glomeruli resembles a multivariate representation of the stimulus space [2]. Olfactory receptors are broadly tuned and hence the response spectra of glomeruli overlap, that is, they exhibit channel correlation. It has been proposed that the AL reduces this channel correlation through lateral inhibition, expanding coding space and using it more efficiently for distributed odor representations [3], [4], [5].

The insect olfactory system thus provides an efficient basis for bio-inspired computational methods to process and classify multivariate data with channel correlation. The AL operates as a decorrelation filter on neural representations of odors before they are delivered to higher brain areas. In previous work, we demonstrated how lateral inhibition in an olfaction inspired network reduces correlation between channels and facilitates separation of multivariate patterns [6], [7]. In this network, the strength of lateral inhibition between any two glomeruli was set according to the correlation in their odor response spectra, as previously suggested by modeling studies [8].

Spike-timing dependent plasticity (STDP) is a well established mechanism for synaptic regulation. STDP adapts synaptic strength according to the temporal relation of pre- and postsynaptic spikes [9]. STDP is a local, unsupervised learning mechanism that depends only on information available at the synapse and doesn't rely on a teacher signal. STDP is ubiquitously observed in various sensory systems *in vitro* and *in vivo* [10]. Furthermore, STDP is experimentally observed in inhibitory synapses as well [11], [12]. Recently, it has been shown that STDP in inhibitory synapses facilitates establishing a state of irregular, asynchronous activity in a balanced spiking network [13].

Here, we propose an unsupervised learning of the lateral inhibition structure in an olfaction-inspired neuronal network via inhibitory spike-timing dependent plasticity (iSTDP). To this end, we implemented an olfaction-inspired spiking network model with lateral inhibitory connections that support iSTDP. We show how the inhibitory connectivity self-organizes to effectively reduce channel correlation. Furthermore, channel decorrelation ensures odor pattern decorrelation which enhances the performance of a Naïve Bayes Classifier in a scent recognition scenario.

II. METHODS & RESULTS

A. Input data & Response patterns

The input data for the AL network model contained 836 odorants from Sigma-Aldrich Flavors and Fragrances Catalog [14]. Odorants in this data set are labeled according to their scent ('fruity', 'balsamic', 'green', 'nutty' etc.) and one odorant may bear more than one label. Odorants in the data set were represented by 184 physico-chemical molecular descriptors. These 184-dimensional vectors were transformed into a ten-dimensional multivariate firing-rate representation using ten virtual receptors (VRs [6]). In brief, a VR is defined as a point in n -dimensional data space and is used to encode n -dimensional real-valued multivariate data into a k -dimensional firing rate representation using k virtual receptors. The activity r_k of the k -th VR depends on the distance $d(\mathbf{s}, \mathbf{p}_k)$ between a data point \mathbf{s} and a VR \mathbf{p}_k according to

$$r_k = \frac{d(\mathbf{s}, \mathbf{p}_k) - d_{\min}}{d_{\max} - d_{\min}} \quad (1)$$

with d_{\max} and d_{\min} , the largest and the smallest distance encountered in the data set. In other words, a VR responds strongly if the stimulus data point is close, and weakly if the distance is large. Thus, VRs encode the real-valued input data set into a bounded, positive representation that can be converted into population firing rates.

¹Theoretical Neuroscience, Institute of Biology, Freie Universität Berlin, Berlin, 14195, Germany

²Bernstein Center for Computational Neuroscience Berlin, Berlin, 10115, Germany

[§]present address: Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, 6525 AJ, the Netherlands

*m.schmuker@fu-berlin.de

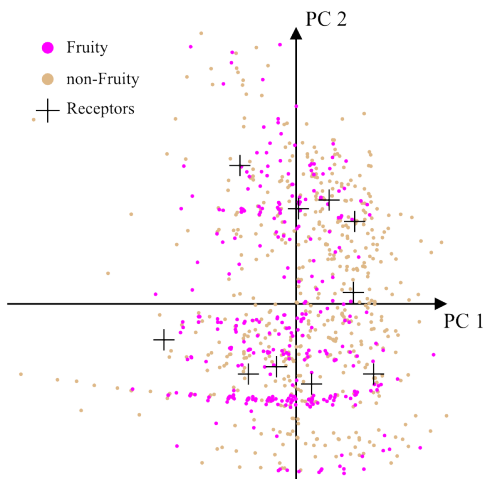


Fig. 1. Distribution of “Fruity” vs. “non-Fruity” odorants in The Flavors and Fragrances data [14]. 2D PCA projection of 184-dimensional space, 48% total variance explained. Crosses mark positions of virtual receptors.

Virtual receptors are placed in data space using a self-organizing map (SOM) [15]. Consequently, VRs cover all relevant parts of chemical space and preserves the local topology in their low-dimensional projections. Fig. 1 depicts a two-dimensional projection of the original 184-dimensional data set and the locations of the VR placed by the SOM.

We used the SOMMER application [16] to train SOMs with a 2×5 toroidal architecture, yielding a ten-dimensional VR representation. In the scope of our model, the evoked response patterns correspond to activation of receptors in the AL.

The VR encoding mimicks two important properties of the olfactory code: Each odorant activates multiple receptors (redundant coding) and each receptor is activated by several odorants (broad tuning curves). Furthermore, odors that are chemically similar according to their location in the 184-dimensional chemical property space evoke similar VR response patterns.

B. Network Layout

Our network model comprises two stages and is directly inspired by the insect olfactory system (Fig. 2). In the first stage, the response of the VRs is translated into spiking activity of k populations of olfactory receptor neurons (ORNs) using a stochastic point process model. To this end, we used Poisson neurons as ORNs to generate independent spiking events based on the population firing rate. The VR activity pattern was mapped into the firing rate interval $[0, 120]$ spikes/s, such that the lowest VR response was transformed into a poisson spike train with 0 spikes/s, and the highest response elicited 120 spikes/s in the ORNs.

The second stage implements a decorrelation network in analogy to the AL network in insects. Each ORN population feeds on a separate population of excitatory projection neurons (PNs). Local inhibitory neurons (LNs) are excited by one PN population and project their inhibitory synapses on all other PN populations. Those inhibitory synapses

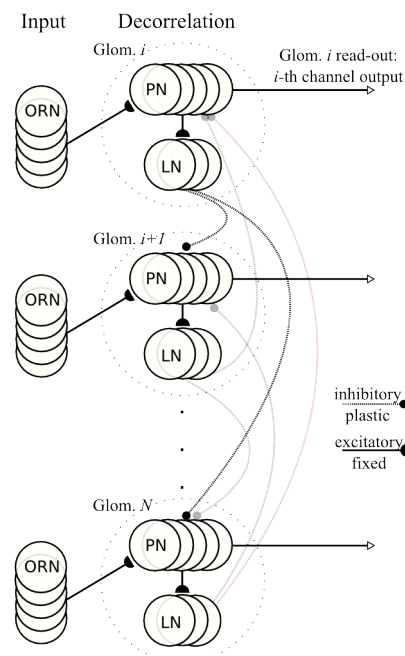


Fig. 2. The network, schematic. ORN: Olfactory receptor neuron, PN: projection neuron, LN: local inhibitory neuron. Each ORN projects to PNs in one glomerulus, which project on LNs in the same glomerulus, which project to PNs in all other glomeruli. Weights from LNs to PNs support STDP.

constitute lateral inhibition and are subject to iSTDP. In the computational model, we used leaky integrate and fire neurons as PN and LN population neurons with conductance-based synapses. Connections from ORNs to PNs, from PNs to LNs and LNs to PNs were made with a probability of 0.4. We used 30 ORNs, 40 PNs and 10 LNs in each glomerulus.

The network scheme with lateral inhibition connections applies a ‘winner-takes-most’ condition (a soft winner-takes-all), where the populations with stronger response profiles inhibit the ones with weaker responses to an odor pattern. Consequently, lateral inhibition processes the input data such that the contrast between channels is enhanced. The connectivity of lateral inhibition was initialized uniformly (that is, all weights equal) with a low initial value.

C. Self-organizing lateral inhibition with spike-timing dependent plasticity

We exposed the network to stimulus patterns. Data points were presented consecutively in random order. We presented each of the 836 odor patterns for one second of biological time. During training, synapses from inhibitory LNs to PNs were subject to iSTDP as described in [13]. According to this learning rule, pairs of pre- and post-synaptic spikes caused potentiation of the inhibitory synapse by an amount proportional to L (eq. 2),

$$L = \frac{e^{-\frac{|\Delta t|}{\tau_{\text{STDP}}}}}{2\tau_{\text{STDP}}}, \quad (2)$$

with Δt the time difference between pre- and postsynaptic spikes, and τ_{STDP} defining the width of the iSTDP window

(fixed at 20ms in this study). Since the STDP time window L is symmetric, the temporal order of pre- and postsynaptic spikes is not important. In addition, each presynaptic spike leads to depression of the synapse. The total change in synaptic weight Δw can be summarized as in eq. 3,

$$\Delta w = \eta [\text{pre} \times \text{post} - \rho_0 \times \text{pre}], \quad (3)$$

where pre and post are the pre- and postsynaptic activity (their interaction defined by L), η is the learning rate and ρ_0 is a constant that acts as a target rate for the postsynaptic neuron. A detailed mathematical description and analysis of the learning rule is provided in [13].

As a consequence of the learning rule, those synapses between presynaptic LNs and postsynaptic PNs which respond strongly to similar odors were potentiated, leading to strong lateral inhibition between the neuron populations with similar response profiles. Note that this relationship imposes a self-regulating mechanism on the strength of inhibitory synapses: Strong lateral inhibition between neuronal populations increases contrast in their firing rates, reducing the similarity in their response profiles. Therefore, correlation decreases as lateral inhibition increases. In consequence, the strength of iSTDP-enabled LN-PN synapses will converge to a value that reflects the similarity of response spectra of the pre- and postsynaptic glomeruli. In this manner, the lateral inhibition structure created through iSTDP reflects the correlation structure in the input data (Fig. 3 A). As a result, channel correlation at PN output after training is significantly reduced compared to VR correlation, and the available coding space is used more efficiently (Fig. 3 B).

D. Decorrelation and benchmarking network performance

In order to determine the benefit of self-organized lateral inhibition for stimulus separation, we trained a Naïve Bayes classifier to predict the scent labeling of odorants based on the activity pattern they evoke in PNs. We measured classifier performance according to the recognition of “Fruity” and “Non-Fruity” odors (5-fold stratified cross-validation, 1000 repetitions with random train and test data splits). Decorrelation by self-organized lateral inhibition performed significantly better than unprocessed data ($p < 10^{-174}$, Wilcoxon rank sum test; Fig. 3 C). It also outperformed decorrelation by uniform lateral inhibition (all inhibitory weights identical, $p < 10^{-77}$).

Furthermore, we analyzed the final inhibitory weights and resulting classification scores depending on different learning rates η and target rates ρ_0 . The learning rate η strongly influenced how well the network reflected the correlation structure in the data (Fig. 4 A). This structure is well reproduced in the inhibitory connectivity for low learning rates. For high learning rates, subsequent odor patterns overwhelm the structure built during the representation of previous odor patterns, counteracting a convergence to a value that matches the input correlation.

Classification scores increased when the inhibitory connectivity matched the VR correlation structure, but only for high target rates ρ_0 (Fig. 4 B). Setting ρ_0 to a low value

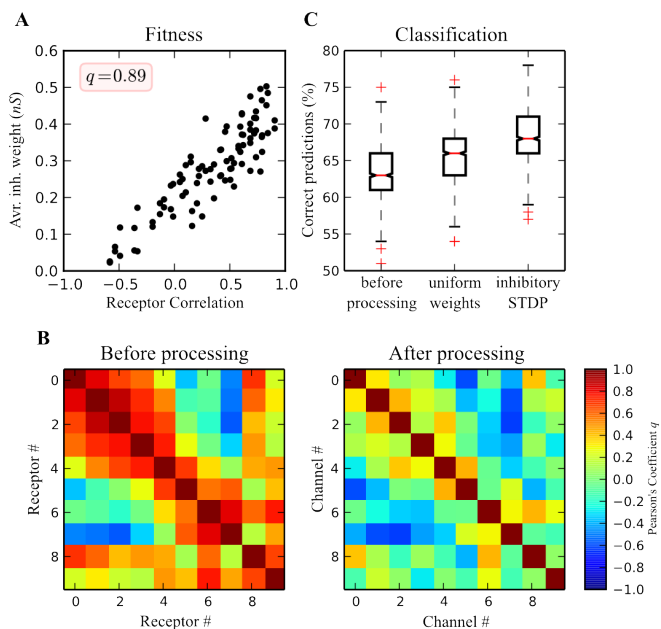


Fig. 3. Impact of iSTDP on inhibitory connectivity and channel correlation. ($\eta = 10^{-3}$, $\rho_0 = 10$ Hz) A) Relationship between channel correlation and average inhibitory weight after stimuli exposure. B) Correlation matrix of the virtual receptors and output channels. C) Classifier performance before and after exposure, and compared to uniform inhibitory weights.

resulted in stronger inhibition that reduced the postsynaptic firing rate to a level close to ρ_0 . However, the spike count observed from a poisson process in a fixed time interval (1 s in our case) exhibits higher variance when the average firing rate is low. Thus, the spike count variance at low rates acts like adding noise to the stimulus representation and consequently has negative impact on the classification performance.

In order to compare these results to a non-spiking approach to decorrelation, we performed principle component analysis (PCA) on the data and trained the Naïve Bayes classifier on the 10 components explaining the most variance in the data set (total 80% explained). This approach achieved slightly better results than the spiking iSTDP approach presented above (avg. 72.0% correct, $P_{25}/P_{75} = 69.9/74.4$ for 1000 repetitions of fivefold cross-validation). Similarly, the performance on the original untreated 184-dimensional dataset yielded slightly better performance values (avg. 73.8% correct, $P_{25}/P_{75} = 72.0/76.2$). However, a direct comparison of absolute classification scores between spiking and non-spiking approaches is difficult, since the conversion of continuous values to spike trains adds noise and potentially degrades the representation of the input data.

III. CONCLUSIONS

We presented the unsupervised learning of input data correlation structure in a scent recognition scenario. We tested how channel correlation affects the weights of STDP-enabled inhibitory synapses in a network inspired by the insect antennal lobe. We showed that upon exposure to artificial stimulus patterns, the inhibitory connectivity of this network

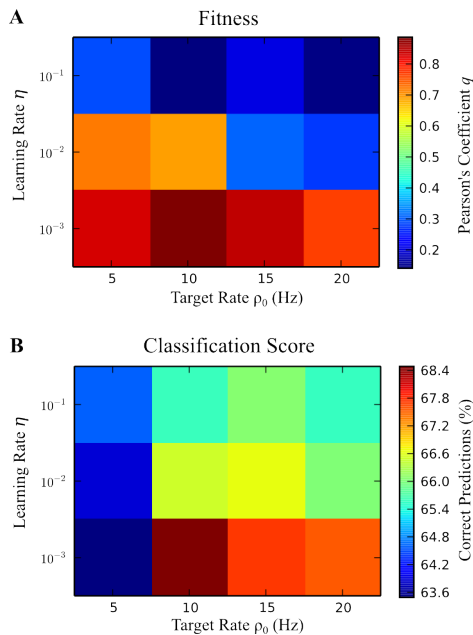


Fig. 4. Fitness and classification scores of different learning rates and target rates. A) How well the final average inhibitory weights between populations reflect the correlation between VR responses. B) Classification scores after the data is processed by iSTDP.

adapts to the correlation structure of the input. The trained network effectively reduces channel correlation in the output indicating that iSTDP is a suitable unsupervised mechanism to assert decorrelation of the input data without a priori information about the input. Reducing channel correlation with the network improved stimulus separability when assessed with a Naïve Bayes classifier in an odor classification task. Our network model incorporates a completely spike-based method to reduce channel correlation, and is therefore suited as a building block for bio-inspired data analysis frameworks, for example on neuromorphic hardware systems supporting spiking neural network models [17], [18].

The classification scores obtained using the transformed data from the spiking network model were slightly lower than obtained by a purely numerical approach. However, it is also clear that the conversion of input data from numerical representation to spike trains introduces noise, with likely negative impact on the classification scores. Hence, it is difficult to compare the absolute performance scores obtained with spiking and non-spiking approaches. We have performed such an analysis in an earlier study using a rate-code model that was not affected by such stochastic noise [6]. The result was that both PCA and correlation-based lateral inhibition achieve similar performance. PCA had maximum performance at about 5-10 PCs, which declined as more PCs were taken into account. In contrast, correlation-based lateral inhibition reached maximum performance at higher dimensionalities, which didn't decline as more dimensions were used. Hence, correlation-based lateral inhibition is

better suited for brain-like, massively parallel computing where data dimensionality potentially has less impact on computational efficiency, but rather the distributedness of the code can be exploited without penalty on computational efficiency. This circumstance argues for a neuromorphic implementation of the present network as a preprocessing step in neuromorphic data processing schemes.

ACKNOWLEDGMENT

This work was funded by a grant from DFG (SCHM2474/1-2 to MS), BMBF (01GQ1001D to MS) and EC Marie Curie-ITN-NETT (Project No. 28914).

REFERENCES

- [1] Wilson, R. I., & Mainen, Z. F. (2006). Early events in olfactory processing. *Annu. Rev. Neurosci.*, 29, 163-201.
- [2] Huerta, R. & Nowotny, T. (2009). Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural computation*, 21(8), 2123-51.
- [3] Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nature Reviews Neuroscience*, 3(11), 884-895.
- [4] Wiechert, M. T., Judkewitz, B., Riecke, H., & Friedrich, R. W. (2010). Mechanisms of pattern decorrelation by recurrent neuronal circuits. *Nature neuroscience*, 13(8), 1003-1010.
- [5] Bernacchia, A., & Wang, X. (2011). Decorrelation by recurrent inhibition in heterogeneous neural circuits. *arXiv preprint arXiv:1107.3111*, 1-31.
- [6] Schmuker, M., & Schneider, G. (2007). Processing and classification of chemical data inspired by insect olfaction. *Proceedings of the National Academy of Sciences*, 104(51), 20285-20289.
- [7] Schmuker, M., Yamagata, N., Nawrot, M. P., & Menzel, R. (2011). Parallel representation of stimulus identity and intensity in a dual pathway model inspired by the olfactory system of the honeybee. *Frontiers in Neuroengineering*, 4, 17.
- [8] Linster, C., Sachse, S., & Galizia, C. G. (2005). Computational modeling suggests that response properties rather than spatial position determine connectivity between olfactory glomeruli. *Journal of neurophysiology*, 93(6), 3410-3417.
- [9] Bi, G., & Poo, M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual review of neuroscience*, 24, 139-66.
- [10] Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nature neuroscience*, 3 Suppl(November), 1178-83.
- [11] Haas, J. S., Nowotny, T., & Abarbanel, H. D. I. (2006). Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *Journal of neurophysiology*, 96(6), 3305-13.
- [12] Lamsa, K. P., Kullmann, D. M., & Woodin, M. a. (2010). Spike-timing dependent plasticity in inhibitory circuits. *Frontiers in synaptic neuroscience*, 2(June), 8.
- [13] Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., & Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science (New York, N.Y.)*, 334(6062), 1569-73.
- [14] Sigma-Aldrich (2004). *Flavors and Fragrances* (Sigma-Aldrich, Milwaukee, WI).
- [15] Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59-69.
- [16] Schmuker, M., Schwarte, F., Brück, A., Proschak, E., Tanrikulu, Y., Givehchi, A., Scheffele, K., et al. (2007). SOMMER: self-organising maps for education and research. *Journal of molecular modeling*, 13(1), 225-8.
- [17] Hausler, C., Nawrot, M. P., & Schmuker, M. (2011). A spiking neuron classifier network with a deep architecture inspired by the olfactory system of the honeybee. (2011) 5th International IEEE/EMBS Conference on Neural Engineering, 198-202.
- [18] Pfeil, T., Grübl, A., Jeltsch, S., Müller, E., Müller, P., Petrovici, M. a., Schmuker, M., et al. (2013). Six networks on a universal neuromorphic computing substrate. *Frontiers in neuroscience*, 7(February), 11.