Proceedings of the 5th International
IEEE EMBS Conference on Neural Engineering
Cancun, Mexico, April 27 - May 1, 2011

A spiking neuron classifier network with a deep architecture inspired by the olfactory system of the honeybee

Chris Häusler*†, Martin Paul Nawrot*† and Michael Schmuker*†
*Neuroinformatics & Theoretical Neuroscience, Institute of Biology
Freie Universität Berlin, 14195 Berlin, Germany
† Bernstein Center for Computational Neuroscience Berlin, 10115 Berlin, Germany
Email: chris.hausler@bccn-berlin.de, {martin.nawrot, m.schmuker}@fu-berlin.de

***Abstract*—We decompose the honeybee's olfactory pathway into local circuits that represent successive processing stages and resemble a deep learning architecture. Using spiking neuronal network models, we infer the specific functional role of these microcircuits in odor discrimination, and measure their contribution to the performance of a spiking implementation of a probabilistic classifier, trained in a supervised manner. The entire network is based on a network of spiking neurons, suited for implementation on neuromorphic hardware.**

I. INTRODUCTION

Honeybees can identify odorant stimuli in a quick and reliable fashion [1], [2]. The neuronal circuits involved in odor discrimination are well described at the structural level. Primary olfactory receptor neurons (ORNs) project to the antennal lobe (AL, circuit 1 in Fig. 1). Each ORN is believed to express only one of about 160 olfactory receptor genes [3]. ORNs expressing the same receptor protein converge in the AL in the same glomerulus, a spherical compartment where ORNs synapse onto their downstream targets. There are about 160 glomeruli in the honeybee AL, matching the number of receptor genes. ORNs from one class (i.e., expressing the same receptor protein) convey information about certain chemical properties of the olfactory stimulus [4]. Each odorant activates many different ORN classes, inducing a spatial response pattern across glomeruli in the AL.

Strong lateral inhibitory interactions between glomeruli in the AL make an impact on information processing [5], [6]. Lateral inhibition has also been shown to significantly increase the performance of a classifier on high-dimensional data in a machine-learning context [7].

From the glomeruli in the AL, uniglomerular projection neurons (PNs) send their axons to Kenyon cells (KCs) in the mushroom body, a central brain structure where different sensory modalities converge and stimulus associations may be formed [8]. At this stage, about 950 PNs project onto a sizeable fraction of the total 160.000 KCs, providing a massive "fan-out" of connections. It has been shown that this fan-out enables efficient classification of odors [9], [10]. Connections between PNs and KCs are realized within small local microcircuits, where PNs and KCs interact with an inhibitory cell population [11] (circuit 2 in Fig. 1). The functional role of those microcircuits is still unclear.

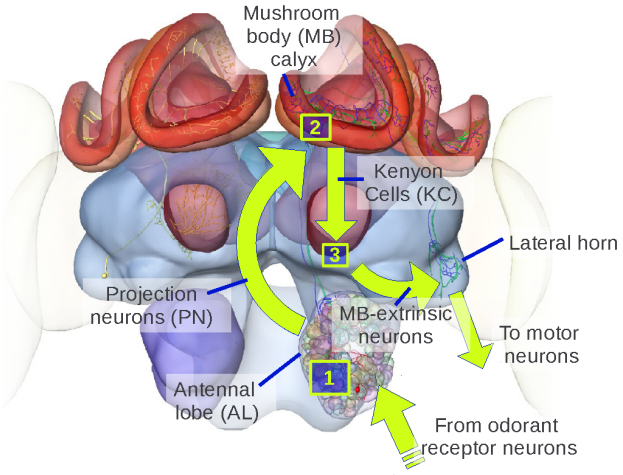

Fig. 1. Processing hierarchy in the honeybee brain. The stages relevant to this study are numbered for reference. Brain image taken from [12].

The arrangement of circuits along the processing hierarchy resembles a deep architecture in which each stage plays a distinct functional role in transforming the input data. In the present work, we illustrate how lateral inhibition enhances linear separability of stimulus patterns by increasing contrast between input dimensions. We assess pattern separability using a classifier network based on spiking neurons that operates with a biologically plausible learning rule. In addition, we demonstrate how microglomerular microcircuits create nonlinear transformations of the input patterns in order to overcome fundamental limitations of the neuronal classifier.

II. SIMULATION TECHNIQUES

A. Simulations of networks of spiking neurons

We used PyNN [13] with the NEURON backend [14] to implement and simulate networks of spiking neurons, using the `IF_facets_hardware1` neuron model (a standard integrate-and-fire type neuron model with conductance-based synapses). Simulations were distributed on a computing cluster using PyDi (http://pydi-parallel.sourceforge.net). Data

Authors listed in alphabetical order. Author contributions: M. Schmuker designed and performed research and wrote the paper, C. Häusler performed research, M.P. Nawrot critically revised and approved the paper.

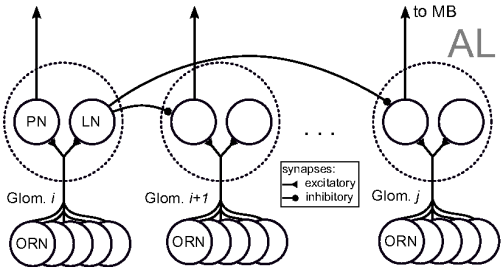

Fig. 2. Network schematic of the antennal lobe where lateral inhibition takes place.

from parallel simulations was collected and analyzed using a MySQL database.

B. Data generation

We sampled data points from two univarate, two-dimensional gaussian distributions. The centroids were located in the first quadrant of the coordinate system, with their coordinates in the interval $[0, 1]$. We controlled correlation between data classes by adjusting the separation angle between the centroids (smaller angle = higher correlation). The standard deviation of the underlying distributions were adjusted to control the "noisiness" of the data classes.

III. NETWORK MODULES

A. Antennal lobe: Input and lateral inhbition

Input from receptor neurons arrives in the AL, where lateral inhibition takes place (Fig. 2). ORNs receiving input from one channel (i.e., one data dimension) converge onto PNs and LNs in the same glomerulus. LNs from each glomerulus project to PNs in all other glomeruli. PNs project to the mushroom body (and to the classifier circuit via KCs, not shown). Data points are presented to the network by setting the ORN's spike rate according to the numeric value of the respective input dimension.

B. Mushroom body: Transforming input space

In the honeybee, olfactory information is processed by circuits within microglomeruli in the lip region of the MB calyx (Fig. 3A). Within microglomeruli, presynaptic boutons from PNs synapse onto KC dendrites, and also onto dendrites of neurons running in the protocerebral-calycal tract (PCT neurons) [11]. There is evidence that those PCT neurons are GABAergic and thus mediate inhibition [15].

We broke down the population connectivity within microglomeruli into two individual circuit realizations (Fig. 3B). The trivial "pass-through" circuit relays information from PNs directly onto KCs without any additional processing. In the more interesting "inverter" circuit a PN projects onto a PCT neuron, which in turn forms an inhibitory projection onto a KC. The KC receives converging direct input from PNs from many different glomeruli. Hence, the KC will always fire if a stimulus is presented, except when the PN is active which will inhibit KC firing indirectlyv via the PCT neuron. When

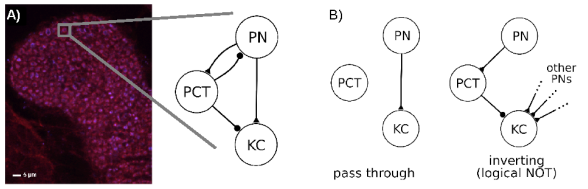

Fig. 3. A) Microglomeruli in the MB calyx in the honeybee brain (blue: PN presynapses, red: KC denrites, from [16] with permission), and the population connectivity within microglomeruli as suggested by [11]. B) Putative realizations of individual circuits within microglomeruli (not exhaustive).

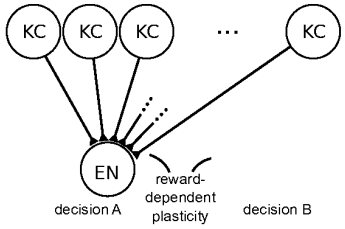

Fig. 4. Schematic of the neuronal classifier. KC: Kenyon Cells, EN: Mushroom-body extrinsic neurons. Excitatory synapses between KCs and ENs are subject to reward-dependent plasticity. Two ENs encode for two decision options.

tuned accordingly, this circuit inverts any input coming from the PN that drives the PCT neuron.

There are many more possible arrangements of these circuits which produce interesting transformations of the input, even more so when not only two input dimensions but a higher-dimensional input space is taken into account, which presents interesting questions for future studies.

C. Mushroom body: Classifier stage

Recently it has been shown that neurons at the MB output (MB-extrinsic neurons, circuit 3 in Fig. 1) exhibit reward-dependent plasticity and reliably encode odor-reward associations [17]. In this light, the output of the MB corresponds can be seen as the stage where stimuli get assigned a meaning, or *value*, similar to data points getting assigned class labels in a classifier. We modeled the classifier stage using a simple decision circuit suitable for supervised learning.

Neurons in the classifier stage receive input from Kenyon cells (KC). In the classifier circuit, input neurons converged on two output neurons, i.e., *decision* neurons (Fig. 4). Upon stimulus presentation, the output neuron with the higher firing rate determined the decision of the classifier (A or B in the figure).

We trained the classifier in a supervised fashion by presenting stimuli (i.e., *labeled data points*) to the network and updating the synaptic weights between input and decision neurons after each stimulus presentation according to the learning rule by Fusi et al. [18]: If the decision was correct (i.e., if the presented data point was from class A and the corresponding neuron "A" fired with a higher rate than the "B" neuron), the weights of synapses on the winner neuron which were active during stimulus presentation were potentiated. If

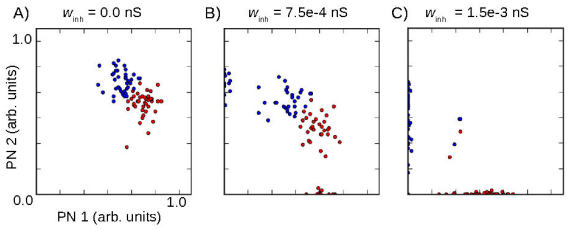

Fig. 5. Impact of lateral inhibition on data representation. Data points are colored according to class adherence. A) No, B) weak, C) strong lateral inhibition. Weight indicates the conductance of inhibitory synapses.

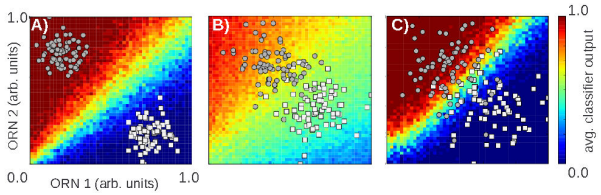

Fig. 6. Impact of lateral inhibition on classifier performance. Disks and boxes indicate data points from two classes (input space). Color represents classifier output sampled from a 50×50 grid averaged over 20 training runs (0.0=A, 1.0=B). A) weakly correlated classes, no lateral inhibition. B) Strongly correlated classes, increased variance, no lateral inhibition. C) High correlation and variance, with lateral inhibition.

the decision was incorrect, active synapses on the winner neuron were depressed.

Synapses were tagged as "active" if the corresponding input neuron fired at 50% or higher of is maximum designed firing rate. Weights were updated according to eq. 1

$$\begin{aligned} \Delta w^{+} &= q \cdot (w_{\max} - w) \quad \text{for potentiation,} \\ \Delta w^{-} &= q \cdot (w_{\min} - w) \quad \text{for depression,} \end{aligned} \tag{1}$$

where $w_{\max} = 0.0033$ nS and $w_{\min} = 0.0$ nS are maximal and minimal bounds to the synaptic weight w, and q is the stepwidth (set to 0.5).

Note that although we did not model reward explicitly, in the scope of this study "reward" corresponded to "having made the right choice". Likewise, synaptic depression was induced by the wrong choice being made, which corresponds to punishment.

IV. CLASSIFIER PERFORMANCE

Our goal was to assess the contribution of two characteristic neuronal microcircuits in the honeybee brain on processing and classification of olfactory stimuli. To quantify the impact of these circuits on stimulus discrimination, we implemented a neuronal classifier based on spiking neurons, which can be trained in a supervised manner (see III-C). Our goal was to explore the capabilities of the neuronal classifier, and to assess how the AL and MB circuits improve its performance. We first analyze the impact of processing in the AL, then we demonstrate a fundamental limitation of the learning rule, and finally we show how this limitation can be overcome by extending the network with microglomerular circuits in the MB.

A. Decorrelation through lateral inhibition in the AL

First we analyzed the impact of lateral inhibition in the AL on the representation of stimuli. Throughout all experiments, we used two-dimensional toy data sets (sec. II-B). Fig. 5 shows the impact of lateral inhibition on stimulus representation, illustrated by a toy data set of two moderately overlapping classes (leftmost panel). When lateral inhibition was increased, data points were pushed towards either axis (*winner-take-most*). In other words, inter-dimensional contrast increased with increasing lateral inhibition.

B. Impact of lateral inhibition on classifier performance

We trained the neuronal classifier to discriminate between classes from the toy data set. After training, we probed the classifier's decision on a regular grid of points in input space. We repeated the training/testing procedure 50 times, each time with a different, random order of the data points, and calculated the average decision of the classifier at each grid point (Fig. 6). When trained with data generated from distributions with well-separated centroids (i.e., low correlation between classes) and low intra-class variance, the classifier robustly succeeded in separating the input space into two partitions corresponding to the two classes (Fig. 6A). When variance and correlation were increased, the location of the classifier's separating hyperplane varied stronger across repetitions, leading to the average decision the be more "smeared out" across the input space (Fig. 6B). Strengthening lateral inhibition in the AL recovered the sharp hyperplane when intra-class variance and inter-class correlation were high (Fig. 6C). Lateral inhibition thus improved the performance of the neuronal classifier on correlated data.

C. Classifier performance on rotated data sets

When rotating the data sets by 45 degrees, 90 degrees and 135 degrees (counterclockwise), we found that the classifier failed to learn any of those arrangements which are not separable along the identity axis (Fig. 7). This behaviour can be explained by the characteristic of the learning rule which only changes synapses which were *active* during stimulus presentation (section III-C). Hence, the classifier was not able to learn *absence* of input in a dimension, which would have been required to learn to separate the data in Fig. 7B, and, to a lesser degree, 7A and C.

D. Benefit of microglomerular circuits in the MB

The above limitation can be overcome by extending the network with microglomerular circuits (see III-B). With lateral inhibiton disabled, we added inverter circuits and pass-through circuits to the architecture. The classifier was then also able to learn the rotated versions of the data (Fig. 8A, B and C).

We then increased the strength of lateral inhibition in the AL network (presynaptic to the MB circuits), and increased data variance and correlation as in Fig. 6. The classifier

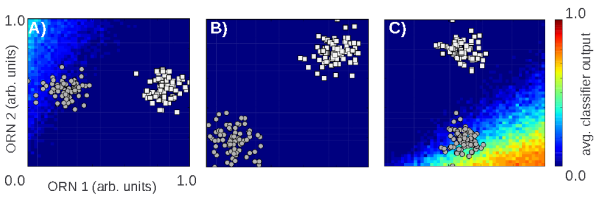

Fig. 7. Performance of the neuronal classifier for data classes which are not separable along the identity diagonal, but rotated by A) 45 degrees, B) 90 degrees and C) 135 degrees counterclockwise.

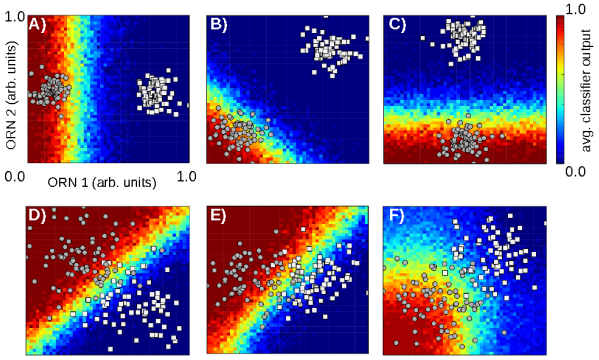

Fig. 8. Classifier performance with microglomerular circuits. A) Without lateral inhibition, data set from Fig. 7, rotated by 45 degrees, B) 90 degrees, C) 135 degrees counterclockwise. D) With lateral inhibition and microglomerular circuits, unrotated, E) rotated by 45 degrees, F) rotated by 135 degrees.

delivered the same performance for non-rotated data as without microglomerular circuits, meaning that inverter circuits did not impact negatively on classifier performance when used together with lateral inhibition (Fig. 8D). The classifier also learned to separate the data rotated by 45 degrees, but the hyperplane was aligned along the identity diagonal (Fig. 8E). This observation indicates that there may be an attractor for the hyperplane to be oriented along the identity diagonal under these conditions. Moreover, for the data set that was rotated by 90 degrees, the hyperplane appeared to be curved in input space (Fig. 8F).

E. Quantification of classifier performance

We quantified the performance of the classifier by training it with datasets of varying difficulty rating and rotation, and calculating classification accuracy using Matthew's correlation coefficient (MCC, [19]). MCC is defined as in eq. 2:

$$MCC = \frac{\mathrm{TP}\cdot\mathrm{TN} + \mathrm{FP}\cdot\mathrm{FN}}{\sqrt{(\mathrm{TN}+\mathrm{FN})\cdot(\mathrm{TN}+\mathrm{FP})\cdot(\mathrm{TP}+\mathrm{FN})\cdot(\mathrm{TP}+\mathrm{FP})}} \quad (2)$$

where TP is the number of true positives (member of class 1 and predicted class 1), TN the number of true negatives, FP the number of false positives (member of class 0 but predicted class 1), and FN the number of false negatives.

To allow for better visualization of the classifier's performance against a single metric of the data, we defined a *Difficulty Rating*, whereby datasets with large separation

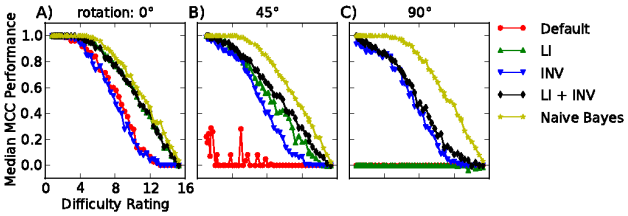

Fig. 9. Comparison of classifier performance of the various network flavors and a Naive Bayes classifier. LI: Lateral inhibition; INV: inverter circuits. A) Performance on unrotated data, B) on data rotated by 45°, B) on data rotated by 90°.

angles and small standard deviations were considered 'easy' and had a low Difficulty Rating whilst those with small separation angles and larger standard deviations were assigned higher Difficulty Ratings. Difficulty rating is defined by eq. 3:

$$\text{Difficulty Rating} = (90 - \Phi) * \sigma \quad (3)$$

where Φ is the angle of separation between the centroids and σ is the standard deviation of the class.

We compared the performance of the neuronal classifier against a Naive Bayes classifier, a probabilistic classifier that utilises Bayes theorem with an assumption of independence. The Naive Bayes classifier is particularly well suited as a reference classifier because it delivers decent perfomance on many real-world data sets [20], [21], and it has no tuneable parameters which may skew any comparison when not adjusted optimally (like, for example, in a support vector machine).

All networks succeed in learning to separate the data which is separable along the identity diagonal ("unrotated data", Fig. 9A). Naturally, performance degrades as the separation becomes more difficult. Lateral inhibition delivers a clear benefit in performance – the networks with lateral inhibition perform almost as well as the Naive Bayes classifier.

For the data which was rotated for 45°, the "Default" network (without lateral inhibition and without inverter circuits) fails completely (Fig. 9B). Inverter circuits rescue some performance. But best performance is achieved when combining inverter circuits with lateral inhibition, pointing out the benefit of inverter circuits in this situation.

Inverter circuits are absolutely required to learn separation of the 90° rotated data (Fig. 9C). Pairing lateral inhibition with inverter circuits increases performance. The Naive Bayes classifier not affected by data set rotation, as expected.

Taken together, these results show that by equipping a neuronal classifier with a lateral inhibition stage and inverter circuits, we can robustly work around the limitation of the biologically plausible learning rule. However, it is also obvious that there is still potential for optimization in the circuits.

V. CONCLUSION AND FUTURE WORKS

A. Conclusion

We presented a neuronal classifier based entirely on a network of spiking neurons. Inspired by the olfactory system of the honeybee, we created a deep network architecture, in

which each processing stage performs a specific task. Lateral inhibition in the AL supports separation of correlated data and increases performance on overlapping classes. Microglomerular circuits in the MB provided transformations of the data that enabled the classifier stage to learn separation of classes which are not separable along the identity diagonal. Depending on the location of data classes in input space, our network performs comparably to a Naive Bayes classifier.

B. Future works

1) Real-world data sets: While the two-dimensional toy data set we used here was very useful to probe the capabilities and limitations of the neuronal classifier and the effect of the circuits that we add, its usefulness for practical applications is limited. We are currently testing the classifier on the surrogate ORN responses from [7] and other high-dimensional benchmark data sets.

2) Neuromorphic Hardware: We are porting our models on neuromorphic hardware [22] as a first step towards implementations of fast and powerful neuromorphic classification devices, applicable to a wide range of parallel sensor data. At the time of writing this manuscript, we have a basic implementation of the classifier network running on the hardware. Preliminary results suggest that the performance on the hardware is similar to that obtained in the simulated network.

VI. ACKNOWLEDGMENTS

This work was supported by the German Ministry for Education and Research through grant 01GQ1001D (Bernstein Center for Computational Neuroscience Berlin), grant 01GQ0941 (Bernstein Focus Learning and Memory in Decision Making), grant 01GQ0772 (Bernstein Partner Project "Olfactory coding: the Mushroom body of the Honeybee"), by DFG grant SCHM2474/1-1, and by an internal grant from Freie Universität Berlin.

REFERENCES

[1] M. Giurfa, "Behavioral and neural analysis of associative learning in the honeybee: a taste from the magic well." *Journal of comparative physiology. A, Neuroethology, sensory, neural, and behavioral physiology*, vol. 193, no. 8, pp. 801–24, Aug. 2007. [Online]. Available: http://www.springerlink.com/content/e6408p3gwt228217

[2] G. A. Wright, M. Carlton, and B. H. Smith, "A honeybee's ability to learn, recognize, and discriminate odors depends upon odor sampling time and concentration." *Behavioral neuroscience*, vol. 123, no. 1, pp. 36–43, Mar. 2009. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/2632763

[3] The Honeybee Genome Sequencing Consortium, "Insights into social insects from the genome of the honeybee Apis mellifera." *Nature*, vol. 443, no. 7114, pp. 931–49, Oct. 2006. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/2048586

[4] M. Schmuker, M. de Bruyne, M. Hähnel, and G. Schneider, "Predicting olfactory receptor neuron responses from odorant structure." *Chemistry Central journal*, vol. 1, p. 11, 2007. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/17880742

[5] R. I. Wilson and Z. Mainen, "Early events in olfactory processing," *Annu. Rev. Neurosci.*, vol. 29, pp. 163–201, 2006. [Online]. Available: http://arjournals.annualreviews.org/doi/full/10.1146/annurev.neuro.29.051605.112950

[6] S. Krofczik, R. Menzel, and M. P. Nawrot, "Rapid odor processing in the honeybee antennal lobe network." *Frontiers in computational neuroscience*, vol. 2, no. January, p. 9, 2008. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19221584

[7] M. Schmuker and G. Schneider, "Processing and classification of chemical data inspired by insect olfaction," *Proceedings of the National Academy of Sciences*, vol. 104, no. 51, pp. 20 285–20 289, 2007. [Online]. Available: http://www.pnas.org/content/104/51/20285.long

[8] M. Heisenberg, "Mushroom body memoir: from maps to models." *Nature reviews. Neuroscience*, vol. 4, no. 4, pp. 266–75, Apr. 2003. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/12671643

[9] R. Huerta, T. Nowotny, M. Garcia-Sanchez, H. D. I. Abarbanel, and M. Rabinovich, "Learning classification in the olfactory system of insects," *Neural computation*, vol. 16, no. 8, pp. 1601–1640, 2004. [Online]. Available: http://www.mitpressjournals.org/doi/abs/10.1162/089976604774201613

[10] T. Nowotny, R. Huerta, H. D. I. Abarbanel, and M. I. Rabinovich, "Self-organization in the olfactory system: one shot odor recognition in insects." *Biological Cybernetics*, vol. 93, pp. 436–446, 2005. [Online]. Available: http://dx.doi.org/10.1007/s00422-005-0019-7

[11] O. Ganeshina and R. Menzel, "GABA-immunoreactive neurons in the mushroom bodies of the honeybee: an electron microscopic study," *Journal of Comparative Neurology*, vol. 437, pp. 335–349, 2001.

[12] J. Rybak, A. Kuß, H. Lamecker, S. Zachow, H.-C. Hege, M. Lienhard, J. Singer, K. Neubert, and R. Menzel, "The digital bee brain: integrating and managing neurons in a common 3D reference system," *Frontiers in Systems Neuroscience*, vol. 4, no. July, p. 30, 2010. [Online]. Available: http://www.frontiersin.org/neuroscience/systemsneuroscience/paper/10.3389/fnsys.2010.00030/

[13] A. P. Davison, D. Brüderle, J. Eppler, J. Kremkow, E. Muller, D. Pecevski, L. Perrinet, and P. Yger, "PyNN: A Common Interface for Neuronal Network Simulators." *Frontiers in neuroinformatics*, vol. 2, p. 11, Jan. 2008. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19194529

[14] M. L. Hines, A. P. Davison, and E. Muller, "NEURON and Python." *Frontiers in neuroinformatics*, vol. 3, no. January, p. 1, Jan. 2009. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19198661

[15] B. Grünewald, "Morphology of feedback neurons in the mushroom body of the honeybee, Apis mellifera." *The Journal of comparative neurology*, vol. 404, no. 1, pp. 114–26, Feb. 1999. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/9886029

[16] S. Krofczik, "The honeybee olfactory system," Ph.D. dissertation, Freie Universität Berlin, 2007. [Online]. Available: http://www.diss.fu-berlin.de/diss/receive/FUDISS_thesis_000000002676

[17] M. Strube-Bloss, M. P. Nawrot, and R. Menzel, "Mushroom Body Output Neurons Encode Odor-Reward Associations," *The Journal of neuroscience*, vol. 31, no. 8, pp. 3129–3140, feb 2011.

[18] S. Fusi, W. F. Asaad, E. K. Miller, and X.-J. Wang, "A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales." *Neuron*, vol. 54, no. 2, pp. 319–33, Apr. 2007. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/17442251

[19] B. W. Matthews, "Comparison of the predicted and observed secondary structure of T4 phage lysozyme," *Biochimica Biophysica Acta*, vol. 405, pp. 442–451, 1975.

[20] M. Sahami, S. Dumais, D. Heckerman, and E. Horvitz, "A bayesian approach to filtering junk e-mail," in *AAAI Workshop on Learning for Text Categorization*, 1998.

[21] A. Bender, H. Y. Mussa, R. C. Glen, and S. Reiling, "Similarity searching of chemical databases using atom environment descriptors (molprint 2d): evaluation of performance." *J Chem Inf Comput Sci*, vol. 44, no. 5, pp. 1708–1718, 2004. [Online]. Available: http://dx.doi.org/10.1021/ci0498719

[22] J. Schemmel, D. Brüderle, K. Meier, and B. Ostendorf, "Modeling synaptic plasticity within networks of highly accelerated I&F neurons," in *IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2007, pp. 3367–3370. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4253401