Odor-Driven Attractor Dynamics in the Antennal Lobe Allow for Simple and Rapid Olfactory Pattern Classification

Roberto Fdez. Galán
r.fdez@biologie.hu-berlin.de
Institute for Theoretical Biology, Humboldt–Universität zu Berlin,
D-10115 Berlin, Germany

Silke Sachse
sachses@mail.rockefeller.edu
Institute for Neurobiology, Freie Universität Berlin, D-14195 Berlin, Germany,
and Laboratory of Neurogenetics and Behavior, Rockefeller University,
New York, New York 10021, U.S.A.

C. Giovanni Galizia
galizia@ucr.edu
Institute for Neurobiology, Freie Universität Berlin, D-14195 Berlin, Germany, and
Department of Entomology, University of California Riverside, CA 92521, USA

Andreas V. M. Herz
a.herz@biologie.hu-berlinde
Institute for Theoretical Biology, Humboldt–Universität zu Berlin,
D-10115 Berlin, Germany

The antennal lobe plays a central role for odor processing in insects, as demonstrated by electrophysiological and imaging experiments. Here we analyze the detailed temporal evolution of glomerular activity patterns in the antennal lobe of honeybees. We represent these spatiotemporal patterns as trajectories in a multidimensional space, where each dimension accounts for the activity of one glomerulus. Our data show that the trajectories reach odor-specific steady states (attractors) that correspond to stable activity patterns at about 1 second after stimulus onset. As revealed by a detailed mathematical investigation, the trajectories are characterized by different phases: response onset, steady-state plateau, response offset, and periods of spontaneous activity. An analysis based on support-vector machines quantifies the odor specificity of the attractors and the optimal time needed for odor discrimination. The results support the hypothesis of a spatial olfactory code in the antennal lobe and suggest a perceptron-like readout mechanism that is biologically implemented in a downstream network, such as the mushroom body.

1 Introduction

The antennal lobe (AL) is the primary brain structure of insects involved in odor coding (Galizia & Menzel, 2001) and short-term memory (Menzel, 1999). This neural structure is functionally homologous to the olfactory bulb in mammals, fish, and amphibia. The neural activity in the AL is highly structured in both space and time, which makes this system especially interesting for studies about neural dynamics and coding on the system's level. The spatial order arises from the precise arrangement of neural clusters called glomeruli. Honeybees possess around 160 glomeruli (Flanagan & Mercer, 1989), and the three to five projection neurons (PN) of each glomerulus relay olfactory information to higher processing areas: the mushroom bodies and the lateral protocerebrum. The temporal response of the PNs shows different features on various timescales: spike activity of single neurons with timing precision of a few milliseconds, membrane-potential oscillations of around 20 Hz, and slow modulations of the membrane potential at around 2 Hz. Membrane oscillations reflect network oscillations that emerge from the transient synchronization of PNs during odor presentation (Wehr, 1999). The slow membrane response is odor specific and matches the overall PN firing rate (Wehr, 1999; Laurent et al., 2001).

The choice of a timescale to analyze the neural dynamics in the antennal lobe has profound effects on the interpretation of the olfactory code. Two main hypotheses have been proposed to explain how odors are encoded. According to the first hypothesis, odors are encoded through a stimulus-induced transient synchronization of PN assemblies during stimulation (Laurent, 1996). The second hypothesis emphasizes the role of spatial patterns of neural activity for odor coding (Hildebrand & Shepherd, 1997; Korsching, 2002), as already suggested by the modular structure of the AL. This hypothesis is supported by calcium-imaging studies that show that a given set of glomeruli is specifically activated by a given odor (Galizia & Menzel, 2001). The activity of the PNs in these glomeruli, however, is not static but modulated by inhibitory mechanisms (Sachse & Galizia, 2002).

In this work, we focus on the detailed temporal development of glomerular activity patterns in the antennal lobe recorded with calcium imaging. A novel experimental technique to stain exclusively projection neurons (Sachse & Galizian, 2002) permits us to specifically study the evolution of those odor-evoked responses that are transmitted to higher processing areas. We represent the AL dynamics in a multidimensional space in which each dimension corresponds to the activity of one glomerulus. Our analysis reveals that during stimulation, different odors trigger different odor-specific trajectories that separate rapidly after stimulus onset. For odor stimuli with a steplike time course, the trajectories converge to stable spatial activity patterns within odor-specific regions after about 1 second and remain unchanged up to small fluctuations until the end of stimulation. Using methods from dynamical systems theory and support-vector machines, we quantify the fine structure of the trajectories and their discriminability. Our find-

ings support the hypothesis that odors are encoded through odor-specific attractors that may easily be read out by a downstream network.

2 Methods

2.1 Animal Preparation and Data Recording. The preparation and calcium imaging of PNs was performed as described in Sachse and Galizia (2002). Briefly, adult worker honeybees were caught and fixed in a plexiglas stage. Projection neurons were backfilled from the protocerebrum with the calcium-sensitive dye fura-dextran (potassium salt, 3000 MW, Molecular Probes, Eugene, OR, USA). Imaging was done using a T.I.L.L. Photonics imaging system (Gräfelfing, Germany). For each measurement, a series of 60 double frames was taken at a sampling frequency of 6 Hz. The interstimulus interval was 40 s. Odors with a duration of 2 s were delivered to the antennae using a custom-made and computer-controlled olfactometer. For each of the tested odors (isoamylacetate, 1-hexanol, 1-octanol, and 1-nonanol; Sigma-Aldrich, Deisenhofen, Germany), 6 μl of the odorant dissolved in mineral oil were applied on a filter paper (1 cm^2) in a plastic syringe. Neural responses were calculated as absolute changes of the fluorescence ratio between 340 nm and 380 nm excitation light. Signals were attributed to identified glomeruli by reconstructing the glomerular structure in the fura ratio images. Glomeruli were identified on the basis of their morphological borderlines using a digital atlas of the AL (Galizia, McIlwrath, & Menzel, 1999). In total, data from seven bees were analyzed.

2.2 Data Analysis. At each point in time t, the activity pattern of the AL network is represented by a point in a multidimensional space in which each dimension corresponds to the calcium signal from one glomerulus. Depending on the bee, between 18 and 23 glomeruli could be identified. For simplicity, we will refer to this glomerular subspace as the AL space. During an odor presentation, the AL activity draws an open curve in this multidimensional space. We first study the system's evolution along these trajectories. Since the data were sampled at a frequency of 6 Hz, the temporal resolution is $\Delta t = 0.167$ s. We denote as $\vec{r}_A(t)$ the trajectory triggered by odor A. To quantify the rate of activity changes, the velocity $v_A(t)$ of the activity state along the trajectories is used;

$$v_A(t) = \left| \frac{\Delta \vec{r}(t)}{\Delta t} \right| = \frac{|\vec{r}_A(t + \Delta t) - \vec{r}_A(t)|}{\Delta t}.$$

As a measure of the force causing changes of the neural activity patterns, we introduce the magnitude $a_A(t)$ of the acceleration vector,

$$a_A(t) = \frac{|\vec{v}_A(t + \Delta t) - \vec{v}_A(t)|}{\Delta t}.$$

The acceleration vector can be decomposed into a tangential component $a_{tA}(t)$ and a normal component $a_{nA}(t)$ relative to the trajectory. These two components are given by

$$a_{tA}(t) = \frac{|\vec{v}_A(t + \Delta t)| - |\vec{v}_A(t)|}{\Delta t}$$

and

$$a_{nA}(t) = \sqrt{|a_A(t)|^2 - |a_{tA}(t)|^2}.$$

The first component provides information about the amount of force parallel to the current trajectory, preserving the movement direction but changing the movement velocity. The second component measures the amount of force that leads to changes in the direction of the trajectory.

In addition to their detailed dynamical properties, we also study the reproducibility and specificity of the odor-driven trajectories. To this end we use support-vector machines (SVMs). An SVM (Vapnik, 1998; Burges, 1998; Boser, Guyon, & Vapnik, 1992) is an algorithm to calculate a hyperplane that separates a multidimensional space into two regions: the first region contains all points of a specified data set within the given data, and the other contains the remaining points. Let us call these regions I and II, respectively. A hyperplane is completely determined by a normalized orthogonal vector $\vec{w}$ and the distance $|b|$ to the origin. In fact, every point $\vec{x}_{hp}$ on the hyperplane satisfies

$$\vec{w} \cdot \vec{x}_{hp} = b.$$

The optimal separating hyperplane between the two data sets is calculated by maximizing the margin with respect to the points of both regions. This mathematical problem can be turned into the minimization of a quadratic form in the positive quadrant. A separating hyperplane exists only if both data sets do not overlap. However, the minimization problem can be reformulated to tolerate some overlap, and the calculation of an optimal separating hyperplane is also possible.

According to SVM theory (Vapnik, 1998; Burges, 1998; Boser et al., 1992), the vector $\vec{w}$ that characterizes the separating hyperplane can always be expressed as a linear combination of the points that lie on the margin—the so-called support vectors. Once $\vec{w}$ and b of the optimal hyperplane are known, it is straightforward to cast new data $\vec{x}$ into I or II according to the following criterion:

if $\vec{w} \cdot \vec{x} > b$, $\vec{x}$ belongs to I

   otherwise $\vec{x}$ belongs to II.   (2.1)

The SVM method can be extended to separating manifolds different from hyperplanes by applying appropriate coordinate transformations (ex-

pressed as kernel functions) equivalent to a change in the metric of the space under consideration—in our case, the AL space. The scalar product, whose associated separating manifold is the hyperplane, was the one that provided the best separation of odors among the tested kernels (scalar product, second- and third-order polynomial, gaussian, and sigmoidal).

To quantify the ability of the SVM to classify an odor correctly, we use a measure called classification performance. This measure is defined as the fraction of points cast into the correct region expressed as a percentage. It yields 100% when the data set of a given odor does not overlap with the rest and yields less when it does.

When two data sets overlap, the optimal separating hyperplane lies at the intersection, and its orientation is very sensitive to the particular distribution of points there. In this case, it is interesting to quantify the stability of the plane through the generalization performance, defined as the relative number of points that can be singly removed without affecting the classification performance. A lower-bound estimator of the generalization performance is obviously the fraction of points that do not belong to the intersection (unambiguous patterns). The exact value of the generalization performance is calculated based on its definition, which represents a bootstrap or leave-one-out method (Efron & Tibshirani, 1993) commonly used to test SVMs (Boser et al., 1992). Let A be the set of n points we want to separate from the rest. Let us now remove one point P of the set A and compute the separating hyperplane. We then check whether P falls into the correct region, where the other $n-1$ points of A lie. We repeat this procedure for all n points of A. The fraction of points successfully classified is expressed as a percentage and measures the generalization performance of a hyperplane. The average of the generalization performance for all odors (all separating hyperplanes) quantifies the reliability of the partition of the AL space into odor-specific regions. We call this measure the separability of odors.

Separating any set of points from the rest of a given ensemble becomes a trivial task in a space of many dimensions. As the number of dimensions increases, it becomes easier to separate the data points. Equivalently, any subset of a small ensemble of points can typically be separated from the rest, even in a low-dimensional space. Therefore, before carrying out all SVM analyses (classification performance, generalization performance, and separability), we pool the data of the seven bees investigated. By doing this, we study the robustness of the olfactory code across trials and also across individuals. In total, 70 points were available, corresponding to four different odors. Up to eight common glomeruli could be identified in all bees. Thus, the AL space used for the SVM analysis has eight dimensions.

3 Results

We represent the activity of the AL network by considering each glomerulus as a component of a vector in a multidimensional space (see section 2.2). At
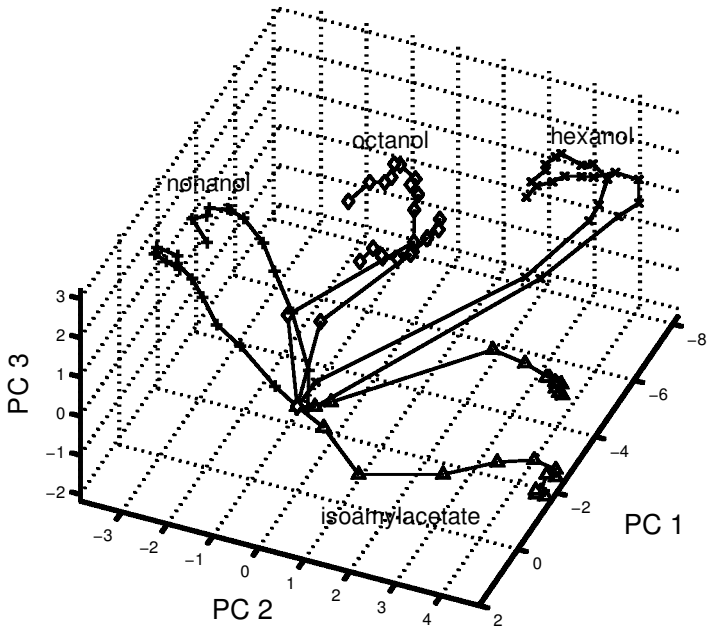
Figure 1: Odor-specific trajectories of the antennal lobe activity. Several odors were presented twice to the same bee, and data were recorded at fixed time intervals (167 ms). Accordingly, the distance between successive data points represents the speed of activity changes. The trajectories depart rapidly from the origin and slow down when they approach odor-specific regions. To generate this figure, the original 21-dimensional AL space has been projected down onto the first three principal components (PC). These three components account for more than 58% of the whole variance of the trajectories.

each instant, the state of the network thus corresponds to one point in the AL space; the temporal evolution of the network is represented by a trajectory. Figure 1 shows such trajectories, caused by several odors presented to a single bee. These data demonstrate that repeating the same stimulus yields similar trajectories. The trajectories do not evolve at a constant speed but start to slow down after approximately 0.5 s and reach odor-specific stable activity pattern after about 1 s. These patterns can be interpreted as stimulus-induced attractors of the network dynamics. After stimulation, the network returns slowly to the resting state (see Figure 2). This relaxation follows a returning path that is different from the trajectory during stimulation. Although trajectories corresponding to the same odor are similar, they exhibit some variability over several repetitions and also across individuals (data not shown).
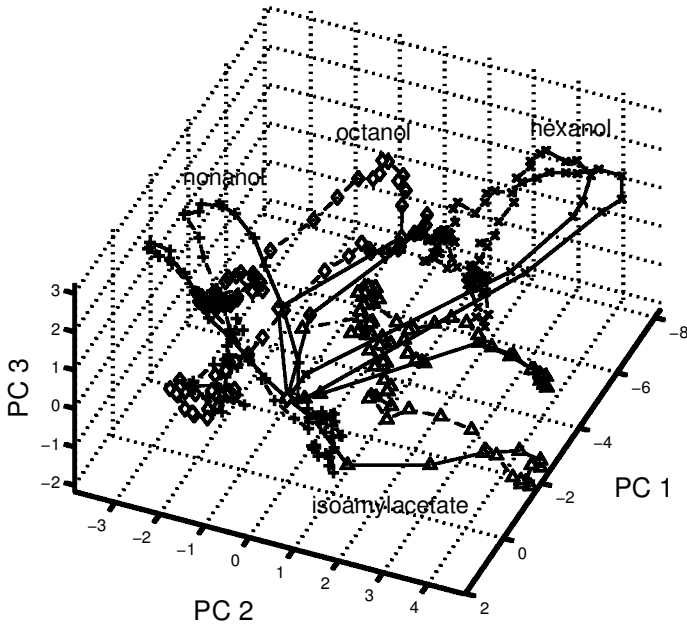
Figure 2: Trajectories of the antennal lobe activity during post-stimulus relaxation in one bee. The solid lines represent the trajectories during stimulation, as in Figure 1. The dashed lines represent the evolution of the trajectories 5 seconds after stimulation. The trajectories during stimulation and relaxation do not coincide; the latter are more irregular.

The observations based on Figures 1 and 2 can be quantified by computing the velocity and the acceleration (see section 2) of the trajectories (see Figure 3). The velocity increases rapidly in the first 400 ms of stimulation, reaches a maximum, and then slows down. After approximately 1 s, the trajectories do not evolve further, and only small fluctuations remain, which are of the same order as those during ongoing activity (see Figure 3A). When the stimulation ends, there is a second peak in the velocity, which is smaller than the first one. This finding demonstrates that the relaxation to the resting state is slower than the excitation during stimulation.

The decomposition of the acceleration vector into a tangential and a normal component relative to the trajectory provides further insight into the origin of the fluctuations (see Figures 3B and 3C). Only during odor presentation is the tangential component substantially different from zero. At the beginning, it is positive, indicating an increase of velocity. After approximately 400 ms, it becomes negative, which means that the velocity slows down. During spontaneous activity, the tangential acceleration fluctuates
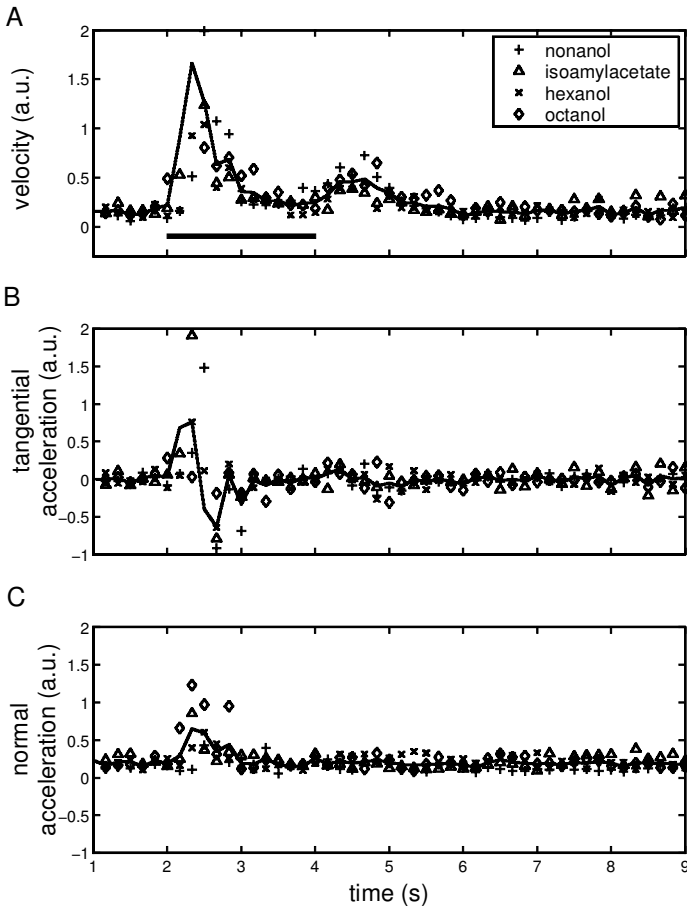
around zero. In contrast, the normal component of the acceleration vector, responsible for changes in the direction of the trajectory, has a certain bias during spontaneous activity and increases during stimulation. However, the larger contribution to the total acceleration during stimulation comes from the tangential component. Thus, in the AL space, odors represent deterministic forces that pull and smoothly bend trajectories of neural activity, whereas the background activity represents a stochastic force that primarily causes random changes in the direction of the trajectories.

How odor specific are the stable activity patterns to which the trajectories converge? The separability measure (see section 2) allows one to answer this question as it quantifies the stability of the hyperplanes that divide the AL space into odor-specific regions. The separability at each point in time (see Figure 4, dotted line) shows that the trajectories become more and

more odor specific during stimulation. The separability reaches a maximum after approximately 1.5 s from stimulus onset. At this point in time, an optimal separating hyperplane for each odor is calculated. The classification performance (see section 2) yields for our data 100% for octanol and nonanol, 93% for isoamylacetate, and 79% for hexanol. These values are substantially larger than chance level (50%) and contrast with the fact that all odors used evoke similar overlapping patterns. The result supports the hypothesis that odors are encoded in the AL as stable spatial patterns of neural activity.

4 Discussion

Based on a multidimensional representation of the calcium dynamics, our analysis shows that trajectories associated with different odors rapidly evolve toward different regions in AL space. As a result of the network dynamics, the trajectories converge to odor-specific attractors. This finding is incompatible with the "winnerless competition" network proposed as a model for the AL (Rabinovich et al., 2001; Laurent et al., 2001), since such a network encodes stimuli in complex trajectories on heteroclinic orbits. However, our results are in agreement with the findings on the dynamic representation of odors in the olfactory bulb of the zebrafish (Friedrich & Laurent, 2001), where the activity patterns also become increasingly odor specific during stimulation.

Figure 3: *Facing page.* Kinematics of the antennal lobe activity. The symbols represent the mean values of the velocity (A) and acceleration (B, C) of activity changes at each point in time for a given odor. The continuous line denotes the mean computed from all odors. The black bar in *A* illustrates the duration of stimulus presentation. The whole data set, pooled over trials with seven bees, was taken into account for the calculation. (A) Before stimulation, the velocity has a roughly constant value, which deviates from zero due to spontaneous activity. At the beginning of the stimulation, the sudden change in velocity indicates a transition leading to a stimulus-dependent state, where activity changes are mainly driven by noise as for the spontaneous activity but with slightly higher values. When the stimulation stops, there appears another peak, though not as pronounced as the first one. This indicates that relaxation to the resting state is slower. (B, C) The acceleration can be decomposed into a tangential and a normal component. During the stimulation period, the tangential component is larger than the normal component. This indicates that a given odor drives the network into a new stable, odor-specific attractor state along an open, smooth curve. The negative peak of the tangential component corresponds to the arrival at the attractor. The normal component has a certain bias in the absence of stimulation. This means that the effect of the background activity on the AL dynamics consists of random changes of the direction of the trajectory.
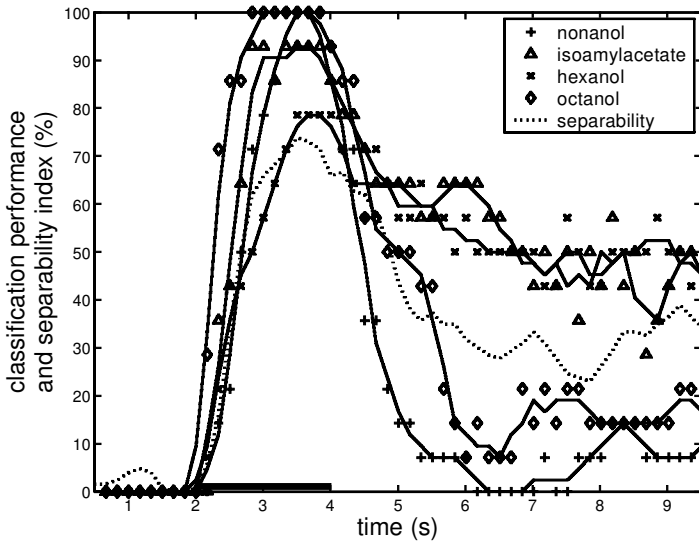
Figure 4: Classification performance and separability as a function of time. For the calculations, all data collected across trials with seven bees were taken into account. During stimulation (black bar), the partition of the antennal lobe space into odor-specific regions becomes possible, as the increase of the separability (dotted lines) shows. At the time of highest separability, one SVM is trained for each odor. Then, the SVMs are tested over time. The classification performance of a given odor (solid lines) is the fraction of points that are recognized as that odor by the respective SVM. Note that except for hexanol, the odors can be very well discriminated even before the trajectories reach the attractors.

Our experimental technique permits us to accurately resolve the large-scale activity patterns in the antennal lobe. These dynamical patterns reproduce reliably odor-specific firing rates of the PNs involved (Wehr, 1999). However, we cannot resolve fast events, such as 20 Hz network oscillation or transient odor-specific synchronizations between the spikes of PNs (Laurent & Davidowitz, 1994; Wehr & Laurent, 1996; Laurent, 1996; Laurent, Wehr, & Davidowitz, 1996). This implies that we cannot answer the question whether ultimately, the olfactory code is based on sequences of oscillating PNs or whether it is based on the attractors of the slow temporal patterns reported here. As part of the ongoing controversy between spike-timing and firing-rate codes, this question remains open for further investigation.

Recent work by Lei, Christensen, and Hildebrand (2000) on the moth's antennal lobe helps us understand the relation between the phenomena observed in electrophysiological and imaging techniques. These authors have studied with intracellular recordings the activity of PNs enclosed in

glomeruli that show an odor-specific response in imaging experiments. They found that PNs within the same glomerulus fire coherently during odor presentation and that the degree of synchrony is modulated by lateral inhibition between active neighboring glomeruli. Hence, the modulation of the odor-induced synchrony is translated into variations of the calcium concentration within the glomerulus, as observed with imaging techniques.

According to our results, the regions of stability, or correspondingly the associated stable calcium patterns, are a reliable mapping of the input onto the AL, and they potentially provide a reliable output to the next neural structures: the mushroom body and the lateral protocerebrum. Within this framework, neurons of a network downstream of the AL could carry out the same task as we have done in order to interpret the AL code: explore the AL space and "look" at which region the evoked response patterns converge to. Support-vector machines provide a good starting point to realize such a readout process. Criterion 1 can be mapped onto a perceptron network (see Figure 5, Rumelhart & McClelland, 1986), where some (or all) units x_n of a bottom layer (the antennal lobe) synapse with strength w_{An} to a given unit in an upper layer (e.g., a Kenyon cell of the mushroom body). If the activation threshold of this unit is set to b, it fires only in response to odor A:

$$\text{if } \sum_n^N x_n \cdot w_{An} = \vec{x} \cdot \vec{w}_A > b \text{ then } \vec{x} \text{ belongs to A}$$

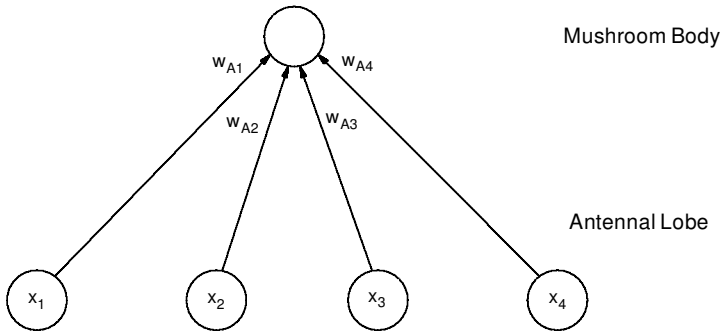$$\text{otherwise } \vec{x} \text{ does not belong to A.}$$

Figure 5: Sketch of the readout mechanism for the network next to the AL. The algorithm used to analyze the imaging data can be mapped onto a perceptron-like neural network, whose architecture is compatible with the anatomy of the bee brain. The units in the lower layer represent individual glomeruli in the antennal lobe. The unit in the upper layer represents a neuron of the mushroom body (Kenyon cell). This unit responds to a given odor A only if the whole activity of the lower units weighted by their synaptic strength $\sum_n^N x_n \cdot w_{An}$ exceeds a threshold b.

Within the theory of artificial neural networks (see, e.g., Hertz, Krogh, & Palmer, 1991) it has been shown how synaptic weights can adapt through Hebbian learning to the vector $\vec{w}_A$, which together with the threshold b determines the optimal separating plane in the AL space for odor A. Notice, however, that no learning process is necessary to identify odors with such a network. Suppose that the threshold b is set constant for all units in the upper layer. By connecting each upper unit with random synaptic weights $\vec{w}$ to the units of the bottom layer, each upper unit will look at a random direction of the AL space. If the number of units in the upper layer is large enough, sufficiently many directions of the AL space will be covered, and, hence, for any given odor-specific attractor, there will always be at least one upper unit that reacts to it. This perceptron-based readout mechanism is consistent with the anatomy of the olfactory system (Strausfeld, 1976). The very high number of Kenyon cells in the mushroom body suggests that there might even be multiple upper units reacting to one odor. The resulting population code would lead to higher signal-to-noise ratios and thus be beneficial for further downstream information processing.

Although the trajectories need about 1 second to reach the attractors, when they can be optimally separated, odors are already well discriminated after 300 ms. This result leads to a testable prediction: If it could be shown that bees can differentiate odors in behavioral tests with minimal reaction times that are significantly shorter than 1 second (the time to reach the steady plateau state), this would provide direct evidence for the coding scheme proposed here. In addition, it would be interesting to investigate whether and to what extent the behavioral discriminability matches that obtained from the neurobiological model.

Acknowledgments

References

Boser, B. E., Guyon, I., & Vapnik, V. (1992). A training algorithm for optimal margin classifiers. In *Annual Workshop on Computational Theorey Archive, Proceedings of the Fifth Annual Workshop on Computational Learning Theorey*, (pp. 144–152). New York: ACM Press. Available on-line: http://citeseer.nj.nec.com/boser92training.html.

Burges, C. J. C. (1998). A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery, 2*, 121–167, Available on-line: http://citeseer.nj.nec.com/burges98tutorial.html.

Efron, B., & Tibshirani, R., (1993). *An introduction to the bootstrap*. London: Chapman and Hall.

Flanagan, D., & Mercer, A. R., (1989). An atlas and 3-dimensional reconstruction of the antennal lobes in the worker honeybee. *International Journal of Insect Morphology and Embriology, 18,* 145–159.

Friedrich, W. R., & Laurent, G. (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science, 291,* 889–894.

Galizia, C. G., McIlwrath, S. L., & Menzel, R. (1999). A digital three-dimensional atlas of the honeybee antennal lobe glomeruli based on optical sections acquired using confocal microscopy. *Cell and Tissue Research, 295,* 383–394.

Galizia, C. G., & Menzel, R. (2001). The role of glomeruli in the neural representation of odours: Results from optical recording studies. *Journal of Insect Physiology, 47,* 115–130.

Hertz, J., Krogh, A., & and Palmer, R. G. (2001). *Introduction to the theory of neural computation.* New York: Perseus Books.

Hildebrand, J. G., & Shepherd, G. M. (1997). Mechanisms of olfactory discrimination: Converging evidence for common principles across phyla. *Annu. Rev. Neurosci., 20, 595–631.*

Korsching, S. (2002). Olfactory maps and odor images. *Curr. Opin. Neurobiol., 12,* 387–392.

Laurent, G. (1996). Dynamical representation of odors by oscillating and evolving neural assemblies. *Trends Neurosci., 19,* 489–496.

Laurent, G., & Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science, 265,* 1872–1875.

Laurent, G., Stopfer, M., Friedriech, R. W., Rabinovich, M. I., Volkovskii, A., & Abarbanel, H. D. (2001). Odor encoding as an active, dynamical process: Experiments, computation and theory. *Annu. Rev. Neurosci., 24,* 263–297.

Laurent, G., Wehr, M., & Davidowitz, H. (1996). Temporal representations of odors in an olfactory network. *J. Neurosci., 16,* 3837–3847.

Lei, H., Christensen, T. A., & Hildebrand, J. G. (2002). Local inhibition modulates odor evoked synchronization of glomerulus specific output neurons. *Nature Neuroscience, 5,* 557–565.

Menzel, R. (1999). Memory dynamics in the honeybee. *J. Comp. Physiol. A, 185,* 323–340.

Rabinovich, M., Volkovskii, A., Lecanda, P., Huerta, R., Abarbanel, H., & Laurent, G. (2001). Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lett., 87,* 68102.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition, Vol. 1: Foundations.* Cambridge, MA: MIT Press.

Sachse, S., & Galizia, C. G. (2002). Role of inhibition for temporal and spatial odor representation in olfactory output neurons: A calcium imaging study. *Journal of Neurophysiology, 87* 1106–1117.

Strausfeld, N. (1976). *Atlas of an insect brain.* Berlin: Springer-Verlag.

Vapnik, V. (1998). *Statistical learning theory.* New York: Wiley.

Wehr, M. (1999). *Oscillatory sequences of firing in the locust olfactory system: mechanisms and functional significance*. Unpublished doctoral dissertation, California Institute of Technology. Available on-line: http://www.cns.caltech.edu/~mike/thesis.html.

Wehr, M., & Laurent, G. (1996). Odor encoding by temporal sequences of firing in oscillating neural assemblies. *Nature, 384*, 162–166.