

Dominique Martinez

Detailed and abstract phase-locked attractor network models of early olfactory systems

Received: 12 May 2005 / Accepted: 18 July 2005 / Published online: 8 September 2005
© Springer-Verlag 2005

Abstract Across species, primary olfactory centers show similarities both in their cellular organization and their types of olfactory information coding. In this article, we consider an excitatory-inhibitory spiking neural network as a model of early olfactory systems (antennal lobe for insects, olfactory bulb for vertebrates). In line with experimental results, we show that, in our network, odor-like stimuli evoke synchronization of excitatory cells, phase-locked to the oscillations of the local field potential. As revealed by a mathematical analysis, the phase-locking probability of excitatory cells is given by an inverted-U function and the firing probability of inhibitory cells is well described by a sigmoid function. These neural response functions are used to reduce the spiking model to a more abstract model with discrete-time dynamics (oscillatory cycles) and binary-state neurons (phase-locked or not). An iterative map, built for explaining the dynamics of the binary model, reveals that it converges to fixed point attractors similar to those obtained with the spiking model. This result is consistent with odor-specific attractors found in recent experimental studies. It also provides insights for designing bio-inspired olfactory associative memories applicable for data analysis in electronic noses.

1 Introduction

Primary olfactory centers in insects and vertebrates show striking similarities both in their cellular organization and their types of olfactory information coding (Strausfeld and Hildebrand, 1999). Odorant molecules are first captured by olfactory sensory neurons distributed on the insect antennae or in the vertebrate epithelium. It is now well established that the large number of sensory neurons expressing the same odorant receptor gene converge onto one or few specific

glomeruli in the insect antennal lobe (AL) or the vertebrate olfactory bulb (OB), where they connect to other cells.

Both the AL and the OB are networks of excitatory cells interconnected via local inhibitory neurons. Olfactory stimuli produce fast oscillations of the local field potential (LFP), resulting from the interplay between inhibitory and excitatory cells in the AL/OB. Indeed, the oscillations persist after ablation of higher brain structures (Laurent and Davidowitz, 1994) but are disrupted when the GABAergic synapses from inhibitory to excitatory cells in the AL/OB are pharmacologically blocked (Stopfer et al., 1997; Hosler et al., 2000; Lagier et al., 2004). An oscillatory activity in the LFP indicates the presence of an olfactory stimulus but the LFP oscillations are independent on the applied stimulus, as any odor leads to the same mean frequency, e.g. 20 Hz in the locust (Stopfer et al., 1997). Such a robust field oscillation is believed to play an important role in the coding of olfactory information by providing a ‘clock’ or temporal frame of reference for the encoding neurons (Laurent, 1996).

It has been reported that the LFP oscillations emerge from the synchronization of the underlying neuronal population. In the presence of an olfactory stimulus, excitatory cells show a synchronized activity, phase-locked to the LFP oscillations (Kashiwadani et al., 1999; Buonviso et al., 2003; Lagier et al., 2004; Friedrich et al., 2004). The LFP oscillatory power in the beta frequency band increases during learning (Ravel et al., 2003) or repeated exposure to an odor (Stopfer and Laurent, 1999), indicating either a better synchronization, a larger number of synchronized cells or both. At each oscillatory cycle, some neurons exhibit phase-locked activity while others do not. The phase of the activity of individual neurons relative to the LFP oscillation does not convey any information about the odor identity (Laurent, 1996; Friedrich et al., 2004), nor about its intensity (Stopfer et al., 2003). Although response latencies with respect to the slower respiratory cycles depends on the odor intensity in the mammal OB (Margrie and Schaefer, 2003), there is so far no experimental evidence for a phase code with respect to the LFP. In contrast, with excitatory cells synchronize at each oscillatory cycle depends on the odor stimulus. More precisely,

D. Martinez
LORIA, Campus Scientifique BP 239,
54506 Vandoeuvre-Les-Nancy, France
E-mail: dmartine@loria.fr
Tel.: +33 383-59-30-72
Fax: +33 383-41-30-79

the identities of the phase-locked neurons evolve in time in an odor-specific manner (Laurent and Davidowitz, 1994; Laurent et al., 1996; Wehr and Laurent, 1996; Laurent, 1999; Friedrich et al., 2004). Recent experimental data indicate that activity patterns become more specific over time (Friedrich and Laurent, 2001; Friedrich et al., 2004) and eventually converge onto odor-specific attractors (Gálan et al., 2004; Sachse and Galizia, 2003).

In summary, electrophysiological studies in the AL/OB demonstrated that olfactory stimuli generate odor-specific synchronization of excitatory cells, phase-locked to the oscillations of the LFP. In order to explain the observed oscillatory synchronization, several computational models have been developed, ranging from networks with detailed conductance-based spiking neurons (Bazhenov et al., 2001; Davison et al., 2003) to networks with binary McCulloch-Pitts neurons (Holub et al., 2002; Quenet and Horn, 2003). On the one hand however, the intrinsic complexity of conductance-based models does not permit mathematical analysis and large scale simulations. On the other hand, oversimplified models derived from scratch lack biological plausibility. In this study, we present a procedure to reduce a spiking neural network model of the insect AL to a more abstract model using binary neurons. The motivation behind this work is twofold. First, the reduced complexity of the binary model will allow a deeper understanding of the mechanisms underlying olfactory coding. Second, the binary model is likely to provide a direct input for designing bio-inspired data analysis methods for artificial olfaction in electronic noses.

In the next section, we describe the spiking neural network exhibiting oscillatory synchronization as a model of early olfactory systems and present simulation results. In section 3, the spiking model is analyzed. The mathematical analysis consists in determining single neuron response probability, conditional to field potential oscillations. As revealed by the analysis, the phase-locking probability for excitatory cells is an inverted-U function while the firing probability for inhibitory cells is well described by a sigmoid-shaped function. In section 4, we present a binary model based on such neuronal activation functions and show that it converges to fixed point attractors similar to those obtained with the spiking model. In section 5 we summarize the results obtained and discuss the hypothesis that early olfactory systems work as phase-locked associative memories.

2 Spiking model : nonlinear integrate and fire neurons and continuous-time dynamics

2.1 Model description

We simulated the conductance based models of excitatory and inhibitory cells from (Bazhenov et al., 2001) and found that their firing frequency in response to a constant input current can be arbitrarily low (type I excitability). It has been shown that the quadratic integrate and fire model or the equivalent theta model are very good approximations of any type I neuron around the threshold (Ermentrout, 1996). Therefore, in

our model we use theta neurons (parameters have been fitted such as to obtain the same frequency-current response than their equivalent conductance based neurons). The state variable θ of a theta neuron obeys the following equation

$$\frac{d\theta}{dt} = (1 - \cos \theta) + (1 + \cos \theta)\alpha J \quad (1)$$

where J is the total input current and $\alpha > 0$ is a constant characterizing the neuron current-frequency response curve. Such a theta neuron can be considered as a point $(\cos \theta, \sin \theta)$ moving on the unit circle (Ermentrout, 1996; Hoppensteadt and Izhikevich, 2002) for which a spike occurs whenever θ crosses π (see Fig. 1). When $J < 0$ and constant, there are two fixed points at $\pm \arccos [(1 + \alpha J)/(1 - \alpha J)]$. The negative one is stable and corresponds to the neuron's resting state. When $J > 0$ and constant, the fixed points do not exist anymore and the neuron fires regularly at a frequency given by $\sqrt{\alpha J}/\pi$. By using the transformation $v = \tan(\theta/2)$ the theta model becomes the quadratic integrate and fire model (Gerstner and Kistler, 2002). The total input current in Eq. (1) is $J = I + I^{syn}$, where I^{syn} is the synaptic current and $I = I^{ext} - I^{th}$. I^{ext} is the external current and I^{th} denotes the threshold, i.e. the minimal current required for repetitive firing. As shown in Fig. 2, our model is a network of N_E excitatory cells (E-cells) and N_I inhibitory cells (I-cells) coupled via simple exponential synapses (model parameters are specified in figure caption).

2.2 Model simulation

Although the AL model used here is an idealized version of the one proposed in (Bazhenov et al., 2001), simulation results remain qualitatively the same. In the presence of a stimulus, the network presents cycles of activity, in line with experimental observations (see Fig. 3 for a typical spike raster plot). Each oscillatory cycle is defined by a volley of quasi-synchronized E-cell spikes followed a few milliseconds later by a similar volley of I-cell spikes. Consecutive cycles of activity are separated by periods of silence lasting for about 50 ms and giving rise to 20 Hz LFP oscillations, similar to that of recorded in the locust (Stopfer et al., 1997). Blocking

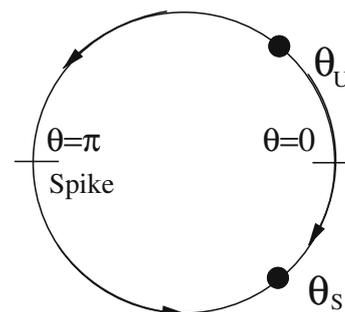


Fig. 1 Neuron model. Phase circle for the theta neuron model (Eq. (1) with $J < 0$). There are two fixed points $\pm \arccos [(1 + \alpha J)/(1 - \alpha J)]$. The positive one is unstable, the negative one is stable and corresponds to the neuron's resting state

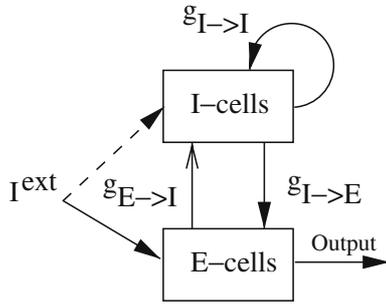


Fig. 2 Network model. E-I network model : $N_E = 90$, $N_I = 30$. The synaptic strengths are $g_{E \rightarrow I} = 0.05$, $g_{I \rightarrow E} = 0.5$, $g_{I \rightarrow I} = 0.1$ with probability of connection $p = 0.4$ and 0 with probability $1 - p$. The decay rates are $\tau_E = 5$ ms and $\tau_I = 6$ ms for the excitatory and inhibitory synapse. E-cells : $\alpha = 0.05$, $I^{th} = 0.5$. I-cells : $\alpha = 0.1$, $I^{th} = 0.8$. Because the number of activated neurons increases with odor concentration (Ng et al., 2002; Sachse and Galizia, 2003), the percentage of neurons in our network receiving the external input mimic odor concentration. Unless otherwise specified, the input is applied to 33% of the neurons chosen at random. For each one of these activated neurons, the input consists of a constant current $I^{ext} = 0.75$ with added gaussian noise (0.1 standard deviation) of 600 ms duration (stimulus onset randomly chosen between 0 and 30 ms). Note however that, for I-cells, this external current is below threshold ($I^{th} = 0.8$) and additional synaptic current coming from the E-cells is needed to induce firing. The output of the network is provided only by the E-cells which project to higher stages

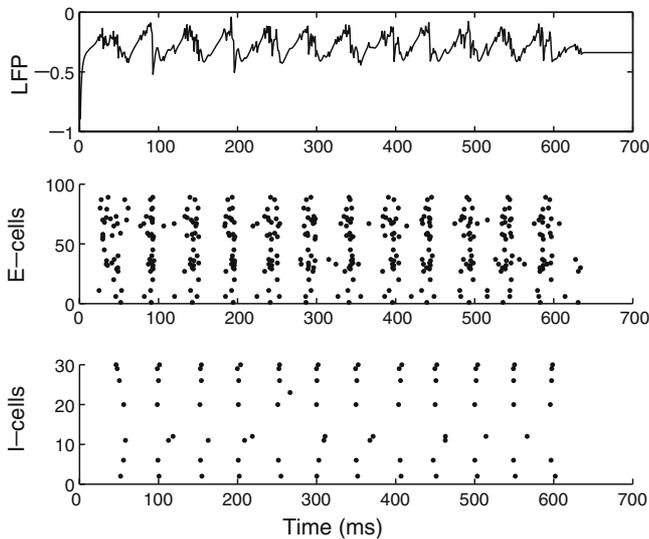


Fig. 3 Simulation of the E-I network. Raster plot of network activity (dots denote individual spikes). Network parameters are given in caption of Fig. 2. The LFP is computed as the instantaneous average of the E-cells' θ variables so that there is no phase lag between the peak of the LFP and the mean firing time \bar{T} of the E-cell population

inhibitory $I \rightarrow E$ synapses in our model disrupted E-cell synchronization and LFP oscillations. This is in agreement with previous experimental and modeling studies which support the functional relevance of fast GABAergic inhibition from I-cells in the synchronization of the E-cells (Stopfer et al., 1997; Hosler et al., 2000; Bazhenov et al., 2001). To see if the mean LFP frequency could convey some informa-

tion about the odor identity and its intensity, we performed repeated simulations with random stimuli of different intensities. Note that the number of activated neurons in our model reflects odor concentration (see model assumptions in caption of Fig. 2). As shown in Fig. 4, the LFP frequency is independent on the odor concentration. Moreover, the small standard deviations, obtained from different runs with random input and network, indicate that the LFP frequency does not depend on particular input configurations and network connectivities. In contrast, it only depends on the time constant of the inhibitory synapse, which is in agreement with previous studies, e.g. (Chow et al., 1998). Because the LFP frequency is independent on the applied stimulus (identity and intensity), it defines a 'clock' or a temporal frame of reference for the encoding neurons.

If one looks more carefully at particular oscillatory cycles displayed in Fig. 5, we see that some E-cells exhibit phase-locked activity while others do not. An E-cell is considered to be phase-locked at a given cycle if it fires within a temporal windows of $\pm \epsilon$ ms around the peak of the LFP (typically, $\epsilon = 5$ ms (Laurent, 1999; Laurent et al., 2001)). In order to see if this phase-locked coding is robust, we performed 20 repeated simulations with different initial conditions and random noise added to the input. We adapted the analysis reported by Laurent et al., (1996) and Bazhenov et al. (2001), by assigning to each E-cell spike a phase ($-\pi$, $+\pi$) relative to its closest LFP peak (zero phase = positive peak of the LFP). Note that the LFP is artificially generated in our simulations as the average of the E-cells' θ variables so that there is no phase lag between the peak of the LFP and the mean firing time \bar{T} of the E-cell population.

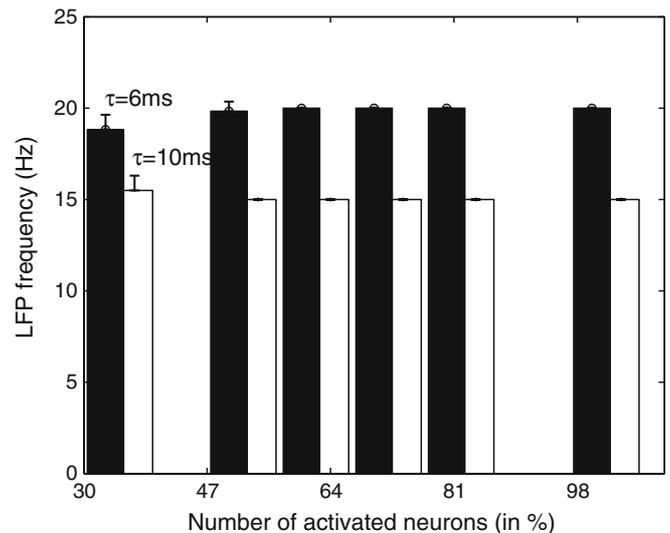


Fig. 4 LFP frequency. LFP frequency vs. odor concentration for two different decay rates of the inhibitory $I \rightarrow E$ synapse. The percentage of neurons activated by the stimulus reflects the odor concentration. This is taken as a model assumption (see caption of Fig. 2). The LFP frequency is given by the frequency of the maximum Fourier component in the power spectrum computed on the unfiltered signal. For each odor concentration, mean LFP frequency and standard deviation are obtained over 10 runs. For each run, network connectivity pattern and input pattern (set of neurons receiving the stimulus) were random

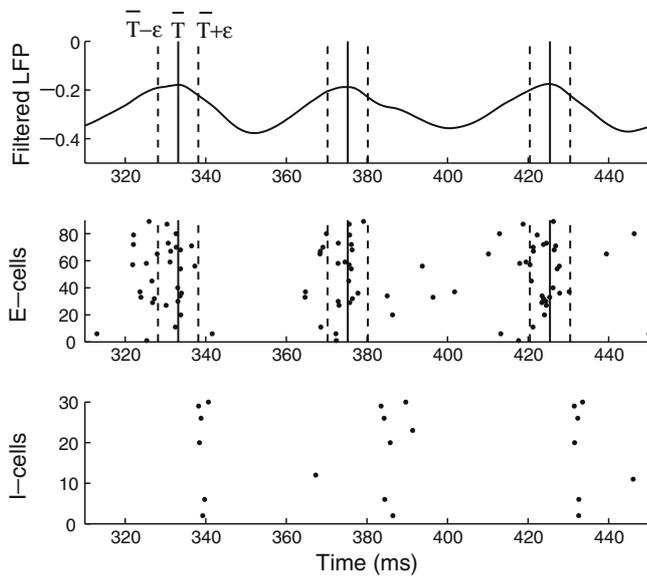


Fig. 5 E-cell Phase-locking. The LFP oscillations define a ‘clock’ or temporal frame of reference of period \bar{T} . The filtered LFP is obtained with a low pass filter (2nd order Butterworth with 30Hz cutoff frequency). At each oscillatory cycle, phase-locked E-cells are those which fire within a temporal windows of $\pm\epsilon$ ms around the mean firing time \bar{T} of the E-cell population. Typically, $\epsilon = 5$ ms (Laurent, 1999; Laurent et al., 2001)

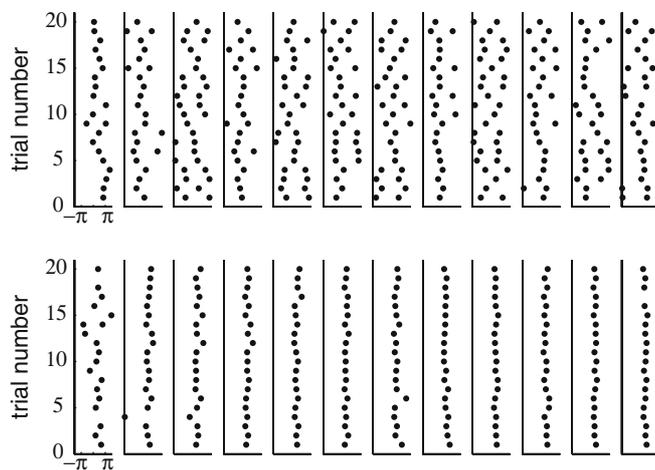


Fig. 6 Phase-locking pattern analysis over 20 simulations of the spiking model with different initial conditions $\theta(t=0)$ and different input currents due to the noise. For repeated trials, the stimulus was applied to the same 33% “random set” of neurons. The top row is for a desynchronized E-cell, the bottom row is for a phase-locked E-cell. Each box corresponds to an oscillatory cycle (from 1 to 12). The phase of each spike is represented as a dot across repeated trials (zero phase in the i th box = no jitter between the spike and the i th peak of the LFP). Note that the twelve oscillatory cycles of 50 ms each (LFP frequency is 20 Hz) have a total duration of 600 ms, which corresponds to the stimulus duration (see legend of Fig. 2)

Figure 6 shows the results of this analysis for 20 different runs of the network and two particular E-cells (one is phase-locked and the other is not). As shown in the figure, the pattern is time invariant, an E-cell being either synchronized

or desynchronized at all the oscillatory cycles. This stable state is reached very quickly, just after one oscillatory cycle. In addition, a phase-locked E-cell produces a single spike per oscillatory cycle and the phase locking pattern is reproducible across repeated trials, in agreement with previous observations (Laurent, 1996; Wehr and Laurent, 1996). As done in (Wehr and Laurent, 1996), the output of an E-cell at each oscillatory cycle is represented as a binary state 1 or 0 depending on whether its firing is phase-locked or not. At each oscillatory cycle, the stimulus is thus characterized by a binary vector lying in a multidimensional space, where each dimension corresponds to the binary state of a given E-cell. Fig. 7 shows a 2D projection of these patterns for three different stimuli. Note that logistic principal component analysis (PCA) has been used for this analysis because it is better suited to modeling binary data than conventional PCA (Schein et al., 2003). We see that different simulations of the network for each stimulus converge to the same representation despite noise and different initial conditions, and that different representations are obtained for different stimuli.

A stimulus is thus robustly encoded by the binary code formed by the identity of the phase-locked E-cells, at each oscillatory cycle. Can this binary code be generated by an abstract model using discrete-time dynamics and binary-state neurons? In order to answer this question, we will build an abstract model based on single response probabilities, conditional to field potential oscillations. In the next section, probabilities of E-cell phase-locking and I-cell firing will be determined mathematically.

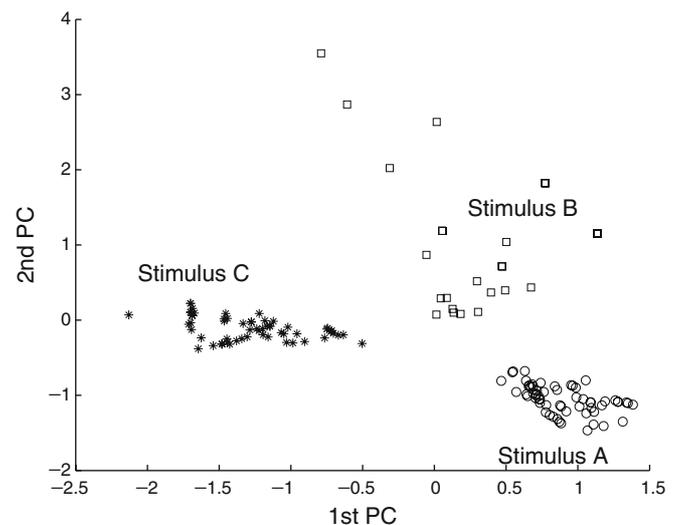


Fig. 7 Stimuli are encoded by the phase-locked binary code. At each oscillatory cycle, the output of the spiking model is provided by the E-cell only. It is represented as a binary vector in a multidimensional space ($N_E = 90$), where each dimension corresponds to the binary state of a given E-cell (1 if phase-locked and 0 otherwise). We pooled the binary data obtained at the different oscillatory cycles (12 cycles), for different odors (3 stimuli A, B and C) and from repeated trials (3 runs for each stimuli). In total, 99 binary vectors were available and projected, using logistic PCA (Schein et al., 2003), onto the first two principal components (PC). Circles, squares and stars are the projected points for stimuli A, B and C, respectively

3 Mathematical analysis of single neuron response probability

3.1 E-cell phase-locking probability

We consider here the conditional probability $P(E|k^-)$ that an E-cell is phase-locked at the current cycle n , given it received k^- inhibitory synaptic events at the previous cycle $n-1$. The conditional probability $P(E|k^-)$ is derived in appendix A1 and is given by Eq. (5) where $\langle k^- \rangle$ is the mean inhibitory drive, i.e. the mean number of unitary inhibitory post-synaptic currents (IPSCs) received by the E-cells on average at the previous cycle, and σ_I is the IPSC temporal jitter, i.e. the standard deviation in the occurring times of the inhibitory events. In the insect AL, we can assume that the inhibitory synaptic events are synchronous (σ_I small). This is justified both with our simulations, from which we obtained $\sigma_I \approx 3.5$ ms and with electrophysiological recordings in the locust AL showing a very synchronous I-cell activity ($\sigma_I \approx 3.8$ ms, see also Fig. 2B in (Laurent and Davidowitz, 1994)). Although in the mammal OB, the I-cell activity is less synchronous ($\sigma_I \approx 22$ ms (Margrie and Schaefer, 2003)), the number k^- of IPSCs received by a particular E-cell at each oscillatory cycle is expected to be large because of the number of I-cells connecting a given E-cell, estimated to be on the order of 10^4 (Davison et al., 2003). Therefore, both for the AL and the OB, the term σ_I^2/k^- is expected to be small and can be neglected in Eq. (5). Moreover, $\epsilon = 5$ ms is in the order of the decay time of the inhibitory synapses $\tau_I = 6$ ms so that the ratio $\tau_I/\epsilon \approx 1$. Considering these approximations, the phase-locking probability can be simplified to

$$P(E|k^-) \approx 1 - \left(\ln \frac{k^-}{\langle k^- \rangle} \right)^2 \quad (2)$$

Note that the simplifications made for deriving Eq. (2) necessarily decrease the negative term in Eq. (5). Thus, Eq. (2) is not a lower bound anymore. We have checked numerically that Eq. (2) is a good candidate for an approximation of the phase-locking probability. Fig. 8 compares the approximation given by (2) and the lower bound given by (5) to estimated data obtained from the simulations of the spiking model. As shown in the figure, the phase-locking probability is an asymmetric inverted-U function centered on the mean inhibitory drive $\langle k^- \rangle$. We see that the values of the phase-locking probability given by (2) are higher than those obtained with the lower bound (5), as explained above. Because of this, the approximation (2) gives a better fit to the simulated data than the original lower bound (5).

3.2 I-cell firing probability

Inhibitory synapses between I-cells introduce a complex competition mechanism in the E-I network so that only a fraction of I-cells are active at each oscillatory cycle. Blocking inhibitory $I \rightarrow I$ synapses leads to an increased number of active I-cells, in agreement with previous experimental study

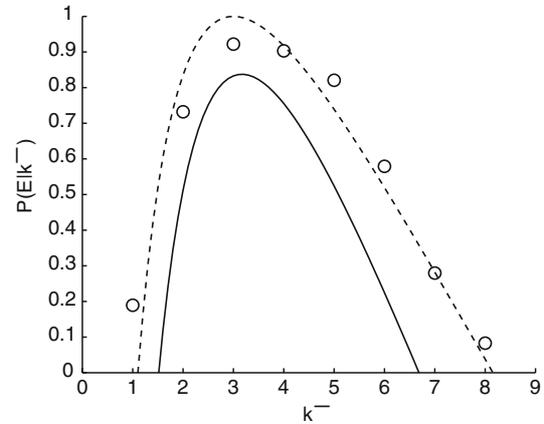


Fig. 8 The phase-locking probability of an E-cell is an inverted-U function. Circles represent the E-cell phase-locking probability ($\epsilon = 5$ ms) vs. number k^- of received IPSCs, estimated over repeated simulations as in Fig. 3 with all the I-cells receiving the external input (20 runs in total). From these simulations, the estimated mean inhibitory drive and jitter were $\langle k^- \rangle \approx 3$ and $\sigma_I \approx 3.5$ ms. The plain curve represents the lower bound given by (5) and the dashed curve is the approximation given by (2)

Nusser et al., (2001). We consider here the firing probability $P(I|k^+)$ of the I-cell receiving k^+ excitatory synaptic events at the current cycle. As shown in appendix A2, $P(I|k^+)$ is given by Eq. (7). Let us denote by $\langle k^+ \rangle$ the mean number of phase-locked unitary excitatory post-synaptic currents (EPSCs) received by the I-cells on average at the current cycle. The I-cells receiving the strongest excitatory drive will cross firing threshold first and inhibit other I-cells. When $k^+ \gg \langle k^+ \rangle$, the probability to be inhibited by another I-cell vanishes, $p_{k^+} \approx 0$ in Eq. (7) and, thus, $P(I|k^+) \approx 1$. Vice versa, when $k^+ \ll \langle k^+ \rangle$, we have $P(I|k^+) \approx 0$. The firing probability $P(I|k^+)$ given by Eq. (7) is actually a sigmoid-shaped function whose threshold depends on the mean excitatory drive $\langle k^+ \rangle$. This is in agreement with experimental data from which the relationship between neuronal firing probability and the LFP has been described by a sigmoid function (Eeckman and Freeman, 1991). Therefore, the firing probability can be simplified to

$$P(I|k^+) \approx \sigma(k^+ - \langle k^+ \rangle - \Theta') \quad (3)$$

where σ is the sigmoid function and $\langle k^+ \rangle + \Theta'$ is a global threshold. Fig. 9 compares the approximations given by (3) and (7) to estimated values obtained from simulations.

4 Binary model : Binary neurons and synapses and discrete-time dynamics

We consider here a network of stochastic binary neurons and discrete-time dynamics. Derivation of the binary model is described in appendix A3. All units in the network are (0,1) binary valued. I-cells may be active (1) or inactive (0) according to $P(I|k^+)$ given by Eq. (3) and E-cells may be phase-locked (1) or not (0) according to $P(E|k^-)$ given by Eq. (2).

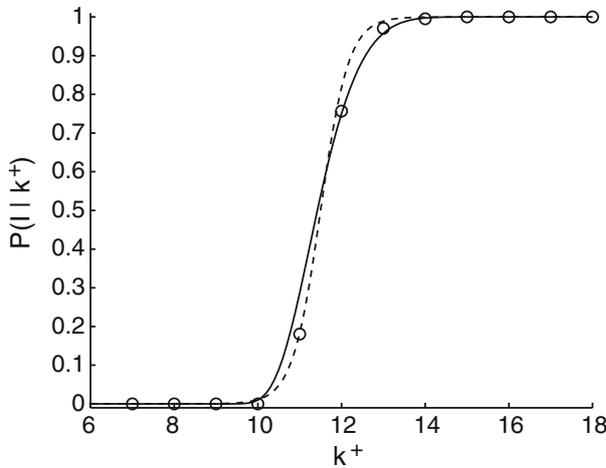


Fig. 9 The firing probability of an I-cell is a sigmoid-shaped function. Circles represent the I-cell firing probability vs. number of received EPSCs k^+ , estimated over repeated simulations as the one in Fig. 3 (50 runs in total). The plain curve represents (7) and the dashed curve is (3). The global threshold of 11.5 ($\langle k^+ \rangle = 15$ and $\Theta' = -3.5$) in (3) and the slope of the sigmoid function have been adjusted so as to match data points as much as possible

The output of the network is provided by the E-cells only. At iteration n , the network activity is given by the fraction of phase-locked E-cells or, equivalently, by the probability P_E that an E-cell is phase-locked

$$P_E(n) = \sum_{k^-} P(E|k^-)P_{k^-}(n-1)$$

$P(E|k^-)$ is given by Eq. (2) and $P_{k^-}(n-1)$ is the probability that an E-cell has received k^- IPSCs at the previous cycle $n-1$, given by a binomial distribution (Eq. 10 in Appendix A3).

At each oscillatory cycle, the competition mechanism due to inhibitory synapses between I-cells insures that only a fraction of the I-cells are active. The probability that an I-cell fires at iteration n is given by

$$P_I(n) = \sum_{k^+} P(I|k^+)P_{k^+}(n)$$

$P(I|k^+)$ is given by Eq. (3) and $P_{k^+}(n)$ is the probability that an I-cell receives k^+ EPSCs at the current cycle n , given by a binomial distribution (Eq. 13 in Appendix A3).

We have compared the dynamics of the binary model, described by Eqs. (8)-(15) in appendix A3, to the one of the spiking model. As shown in Fig. 10, the dynamics of the two model converges to similar fixed points. The convergence is also very fast, obtained in only one oscillatory cycle in most cases.

5 Discussion : Early olfactory systems as phase-locked associative memories

In this paper, we considered an excitatory-inhibitory spiking neural network exhibiting oscillatory synchronization as a model of the antennal lobe, the insect early olfactory

system. We presented a procedure to reduce this spiking model to a more abstract model using binary neurons. The process of abstraction consists in determining single neuron response probability, conditional to field potential oscillations. We found that phase-locking probability for excitatory cells is given by an inverted-U function while firing probability for inhibitory cells is well described by a sigmoid.

Based on these neural response functions, we then considered a network of binary-state neurons and discrete-time dynamics. We compared the dynamics of this binary model to the one of the spiking model, for different probabilities of connection. The binary network dynamics converges to fixed-point attractors similar to those obtained with the spiking neural network. Moreover, different simulations of the network for each stimulus converge to the same attractors, despite noise and different initial conditions, and that different attractors are obtained for different stimuli. In Fig. 7, the three clusters of points, corresponding to the attractors obtained from three different stimuli, could be easily separated by simple discriminant functions. This view is compatible with a perceptron-readout mechanism in the mushroom body or the olfactory cortex (Gálan et al., 2004). Similarly to our study, experimental and modeling works have shown that neuronal activity in early olfactory systems follows trajectories converging onto odor-specific attractors (Gálan et al., 2004; Raman and Gutierrez-Osuna, 2004; Sachse and Galizia, 2003; Linster and Masson, 1996). These trajectories need about one second to reach the attractors (Gálan et al., 2004). They cannot be seen in our model because convergence towards the stable state is too rapid, within one oscillatory cycle only. Nevertheless, they become apparent when a temporal dynamics is added to the external stimulus, e.g. when the stimulus is taken as a current step with a rise time constant instead of a constant external current. Therefore, the temporal dynamics of the glomerulus activities, lasting for several seconds (Sachse and Galizia, 2003; Linster and Masson, 1996), could explain the trajectories seen in AL activity.

A property of neural networks working as associative memory is their ability to converge to fixed point attractors representing the stored memories. Previous associative memory models of early olfactory systems are based on a firing rate assumption (Hendin et al., 1997; Li and Hopfield, 1989; Li and Hertz, 2000). In contrast, the binary and the spiking models developed in this paper can be seen as phase-locked associative memory models in which memories are embedded in a phase-locked E-cell activity relative to the LFP oscillation. Because E-cells have a binary response (1 if phase-locked and 0 otherwise) and connections are binary (either present or absent), a related associative memory model is the Willshaw network (Willshaw et al., 1969; Palm, 1980). In the Willshaw model, neurons are thresholded binary units and long-term storage is enabled by binary synaptic connections adjusted with a clamped version of Hebbian learning. The model is efficient if training patterns are sufficiently sparse, i.e. have a small number of active components, and if the threshold is properly adjusted, e.g. according to the global activation of the network (Graham and Willshaw, 1995; Schwencker et al., 1996; Hirase and Recce, 1996).

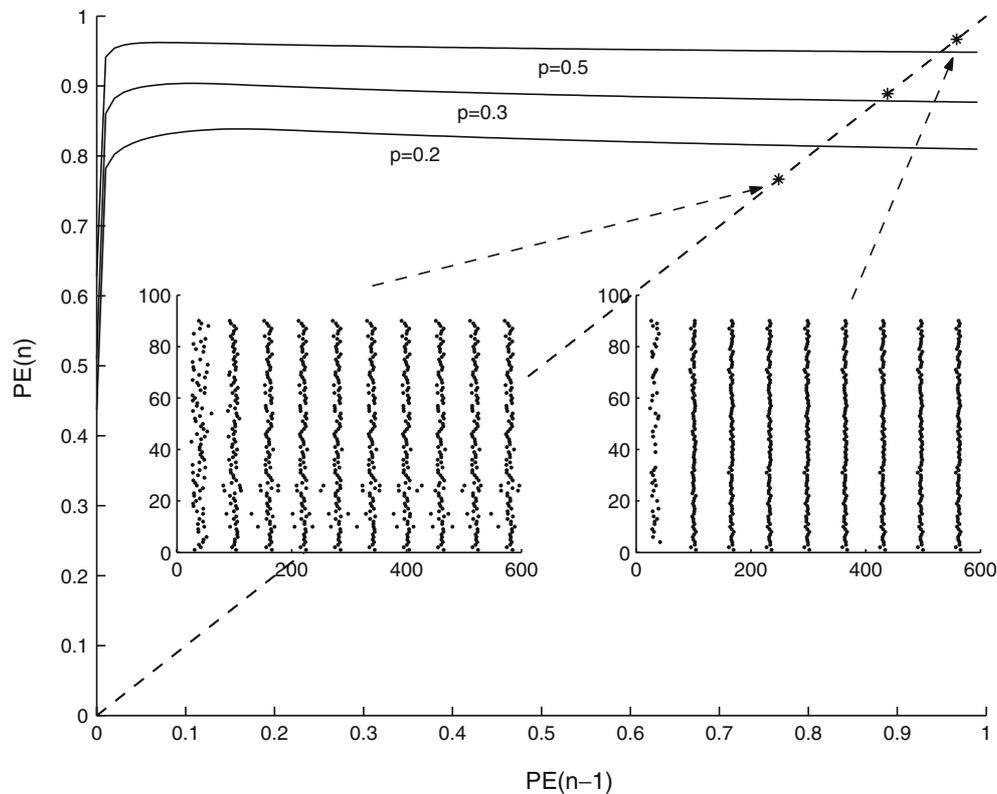


Fig. 10 Iterative map $P_E(n) = f(P_E(n-1))$ for the binary model described by Eqs. (8)-(15) in appendix A3, for different probabilities of connection $p = 0.2, 0.3$ and 0.5 , and for 100% of the neurons activated by the stimulus. Gaussian approximations of the binomial have been used in Eqs. (10) and (13). The fixed points for the binary model, obtained at the intersections with the identity line, correspond to the fixed points obtained for the spiking model and indicated by the stars. Two typical raster plots of E-cell spiking activity are given for the spiking model ($p = 0.2$ and $p = 0.5$)

Like the global adjustment of the threshold in the Willshaw model, we found, in our model, that the activity of single neurons is regulated by the mean activity of the other cells. For example, if the inhibition received by an E-cell does not deviate too much from the mean inhibition, then the E-cell will be phase-locked. Such a regulation of the neural response could act as a gain control mechanism. More neurons are activated when odor concentration increases (Stopfer et al., 2003; Sachse and Galizia, 2003). Let us assume that twice of the I-cells are activated by the same odor but at a higher concentration. Both the inhibitory drive k^- received by an E-cell and the mean inhibition $\langle k^- \rangle$ will be doubled. However, their ratio will remain unchanged in Eq. (2), leading to the same value of the phase-locking probability. Thus, E-cells that are phase-locked at low odor concentrations will also be phase-locked at higher concentrations of the same odor. Different codings of the same odor at different concentrations will then share common phase-locked E-cells. Recognizing an odor regardless of its concentration could be accomplished by decoding the set of common phase-locked E-cells, with downstream neurons acting as coincidence detectors. This concentration-invariant coding scheme is consistent with experimental data (Perez-Orive et al., 2002; Stopfer et al., 2003). In the binary and spiking models, the output is provided by the E-cells only. A possible role for the I-cells is to transform the initial E-cell pattern into a sparse

code so as to facilitate stored pattern retrieval. This is particularly true for the mammal OB for which the ratio of I-cells to E-cells is about 150:1. Projecting input patterns into the high dimensional space defined by the I-cells (dimension in the order of 10^6 for the OB) shares similarities with kernel methods like Support Vector Machines (Cortes and Vapnik, 1995).

Ongoing research is aimed at developing a kernel associative memory based on the ideas derived from this paper. An associative memory built on such ideas is expected to be applicable for artificial olfaction using electronic noses. Other lines of research as extensions to this work are interesting to pursue. In particular, we have shown in (Martinez, 2005) that the spiking neural network does not converge to fixed-point attractors when an adaptation current is added to the I-cells. Indeed, each time an I-cell fires, it increases its adaptation current and makes subsequent firing harder. Such a unit fatigue mechanism, introduced into the Hopfield network, has been shown to generate oscillatory activities (Horn and Usher, 2003, 1991). In our spiking model, the firing probability for inhibitory cells has been described by a sigmoid-shaped function whose threshold is dependent on the adaptation current. Similar to (Horn and Usher, 2003, 1991), our binary model can then be turned into an oscillatory associative memory by using, for each I-cell in Eq. (3), an individual dynamic threshold.

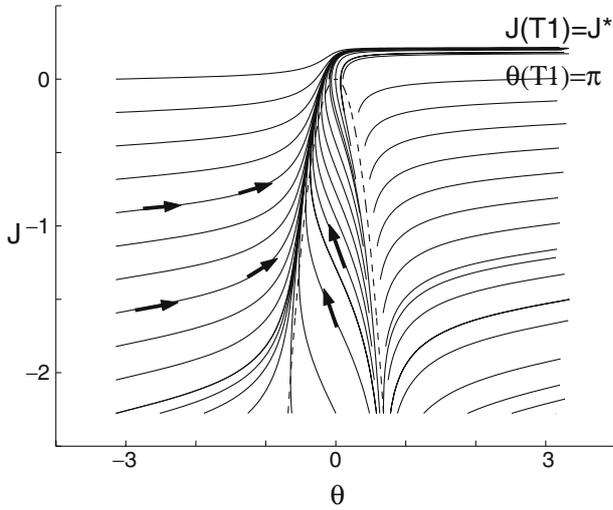


Fig. 11 Trajectories in the phase plane for E-cells receiving inhibition. Phase plane (θ, J) for Eqs. (1) and (4) with $t^f = 0$ and for the E-cells considered in the paper. For strong $g_{I \rightarrow E}$, trajectories converge towards a given trajectory, called a 'river' in (Börgers and Kopell, 2003), so that they all reach approximately the same point $(\theta, J) = (\pi, J^*)$ at the firing time T_1 . Trajectories initialized at θ far from π (indicated with arrows) are rapidly attracted to the river and, thus, the first spike approximately occurs at T_1 . Trajectories initialized close to $\theta = \pi$ induce a first spike soon after $t^f = 0$. Right after that, $\theta = -\pi$ and the trajectories are attracted to the river so that the second spike approximately occurs at T_1

Appendix

A1. Phase-locking probability of an E-cell

We first consider an E-cell receiving a single IPSC at a time t^f via a simple exponential synapse of strength $g_{I \rightarrow E}$ with decay time τ_I . The total input current, for $t \geq t^f$, is then given by

$$J(t) = I - g_{I \rightarrow E} e^{-(t-t^f)/\tau_I} \quad (4)$$

where $I = I^{ext} - I^{th}$.

Börgers and Kopell (2003) have shown that the firing time T_1 of a theta neuron receiving a single strong IPSC is relatively independent of the initial condition $\theta(t=0)$. Provided $g_{I \rightarrow E}$ is large enough, trajectories in the phase plane (θ, J) are all attracted towards a given trajectory so that they all reach approximately the same point (π, J^*) at firing time. This is shown in Fig. 11 for our theta model of E-cell. The result is an almost complete loss of the initial condition $\theta(t=0)$. Whatever the initial condition might be, the total input current is approximately equal to J^* at the firing time T_1

$$J^* \approx J(T_1) = I - g_{I \rightarrow E} e^{-(T_1-t^f)/\tau_I}$$

and thus

$$T_1 \approx \tau_I \ln g_{I \rightarrow E} - \tau_I \nu(I) + t^f$$

where $\nu(I) = \ln(I - J^*)$.

In order to obtain expressions for the firing time of an E-cell, we have generalized Börgers and Kopell's (2003) study

to the case of an E-cell receiving k^- IPSCs of strength $g_{I \rightarrow E}$ at time t_i^f , $i = 1, 2, \dots, k^-$. Without loss of generality, we consider that the neuron fires after receiving the k^- IPSC. At the firing time T_{k^-} , the total input current is approximately equal to J^*

$$J^* \approx J(T_{k^-}) = I - \sum_{i=1}^{k^-} g_{I \rightarrow E} e^{-(T_{k^-}-t_i^f)/\tau_I}$$

and thus, the firing time of an E-cell which received k^- inhibitions is

$$T_{k^-} \approx \tau_I \ln(k^- g_{I \rightarrow E}) - \tau_I \nu(I) + \tau_I \ln \frac{1}{k^-} \sum_{i=1}^{k^-} e^{t_i^f/\tau_I}$$

We consider that the E-cell is phase-locked at the current cycle if its firing time T_{k^-} is within a temporal window of $\pm \epsilon$ ms around the ensemble mean \bar{T} (see Fig. 5). The probability $P(E|k^-)$ that an E-cell receiving k^- IPSC is phase-locked at the current cycle is then given by

$$P(E|k^-) = P(|T_{k^-} - \bar{T}| < \epsilon)$$

Using the Chebychev inequality, we have previously derived in (Martinez, 2005) a lower bound on $P(E|k^-)$

$$\begin{aligned} P(E|k^-) &\geq 1 - \frac{E[(T_{k^-} - \bar{T})^2]}{\epsilon^2} \\ &\approx 1 - \frac{1}{\epsilon^2} \left(\frac{\sigma_I^2}{k^-} + \tau_I^2 \left(\ln \frac{k^-}{\langle k^- \rangle} \right)^2 \right) \end{aligned} \quad (5)$$

where $\langle k^- \rangle$ is the mean number of IPSCs received by the E-cells on average at the previous cycle and σ_I is the IPSC temporal jitter (standard deviation in the occurring times of the inhibitory events). Note that the bound is valid for $k^- > 0$ only. An E-cell without inhibition is not phase-locked, i.e. $P(E|k^- = 0) = 0$.

A2. Firing probability of an I-cell

In our model, the external current received by the I-cells is below threshold and additional synaptic current coming from the E-cells is needed to induce I-cell firing. We first consider an I-cell receiving a single EPSC at a time t^f via a simple exponential synapse of strength $g_{E \rightarrow I}$ with decay time τ_E . The total input current, for $t \geq t^f$, is then given by

$$J(t) = I + g_{E \rightarrow I} e^{-(t-t^f)/\tau_E} \quad (6)$$

where $I = I^{ext} - I^{th}$.

The I-cell will fire if the received synaptic current brings its trajectory in the phase plane above the *go curve*, as defined in (Rubin and Bose, 2004) for the theta model and shown in Fig. 12A. This will happen if $g_{E \rightarrow I}$ is high enough. As shown in Fig. 12B, an I-cell receives many almost synchronized synaptic events during LFP oscillations. We consider here it receives k^+ phase-locked EPSCs at the current cycle. Because the period of silence that separates consecutive

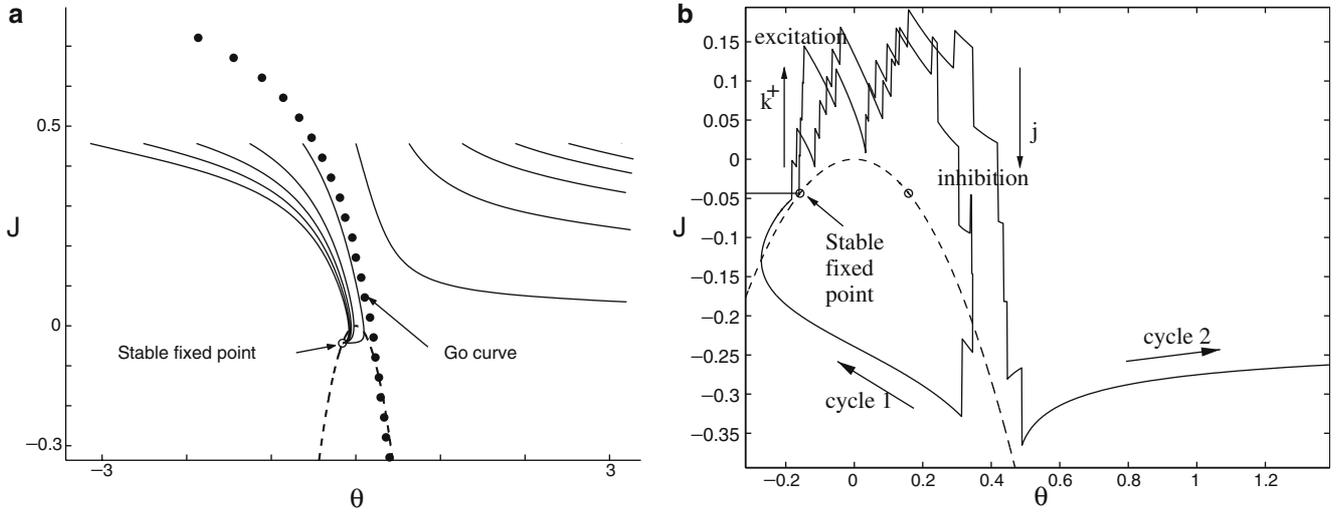


Fig. 12 Trajectories in the phase plane for I-cells receiving excitation. (a) Phase plane (θ , J) for Eqs. (1) and (6) with $t^f = 0$ and for the I-cells considered in the paper. Trajectories initialized above the go curve induce firing while trajectories initialized below the go curve converge to the stable fixed point. (b) An I-cell receives many almost synchronized synaptic events during LFP oscillations (k^+ EPSCs from the E-cells and j IPSCs from the other I-cells). The figure shows the trajectory of an I-cell in the phase plane during two consecutive LFP oscillations (extracted from simulation in Fig. 3). In both cases, the excitation is large enough to bring the trajectory above the go curve. However, firing only occurs at the second oscillatory cycle. At the first cycle, inhibition from other I-cells brings back the trajectory below the go curve and prevents I-cell firing

cycles of activity is large compared to the synaptic time constants (50 ms vs. 5 and 6 ms), the I-cells are at rest when E-cells fire (see Fig. 12B). The resting state is defined as the stable fixed point for the theta model. The I-cells receiving the strongest excitatory drive will cross firing threshold first. They will inhibit other I-cells and may even suppress their activity. This leads to a complex competition between I-cells (see Fig. 12B). For simplicity, we consider an I-cell connected to $n_I = p(N_I - 1)$ other I-cells on average and receiving k^+ EPSCs and j IPSCs at the same time. The cell will fire when $k^+ g_{E \rightarrow I} - j g_{I \rightarrow I} \geq \Theta$ where Θ is the additional current needed to move from the stable fixed point to the go curve. From Fig. 12A, $\Theta \approx 0.3$. The number j of received IPSCs is a binomial random variable with parameter (n_I, p_{k^+}) (p_{k^+} is the probability that another I-cell receiving the external current also receives more than k^+ EPSCs and therefore fires earlier). Using the gaussian approximation of the binomial, we can find the firing probability of the I-cell as

$$P(I|k^+) = p\left(j \leq \frac{k^+ g_{E \rightarrow I} - \Theta}{g_{I \rightarrow I}}\right) \approx \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\frac{Z}{\sqrt{2}}\right) \quad (7)$$

where erf is the error function and $Z = (a - n_I p_{k^+}) / \sqrt{n_I p_{k^+} (1 - p_{k^+})}$ with $a = (k^+ g_{E \rightarrow I} - \Theta) / g_{I \rightarrow I}$.

A3. Derivation of the binary model

We consider a network of binary neurons, with probability of connection p . At iteration n , the network activity is given by

the fraction of phase-locked E-cells or, equivalently, by the probability P_E that an E-cell is phase-locked

$$P_E(n) = \sum_{k^-} P(E|k^-) P_{k^-}(n-1) \quad (8)$$

with

$$P(E|k^-) = \begin{cases} 1 - \left(\ln \frac{k^-(n-1)}{\langle k^- \rangle}\right)^2 & \text{if } k^-(n-1) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

$$P_{k^-}(n-1) = \binom{N_I^a}{k^-} p^{k^-} (1-p)^{N_I^a - k^-} \quad (10)$$

$N_I^a = N_I P_I(n-1)$ represents the number of I-cells that have been activated at the previous cycle. At iteration n , the probability that an I-cell is active is given by

$$P_I(n) = \sum_{k^+} P(I|k^+) P_{k^+}(n) \quad (11)$$

with

$$P(I|k^+) = \sigma(k^+(n) - \langle k^+ \rangle - \Theta') \quad (12)$$

$$P_{k^+}(n) = \binom{N_E^a}{k^+} p^{k^+} (1-p)^{N_E^a - k^+} \quad (13)$$

$N_E^a = N_E P_E(n)$ represents the number of E-cells that are phase-locked at the current cycle. In Eqs. (9) and (12) the mean number of inhibitions $\langle k^- \rangle$ received on average by the E-cells and the mean number of excitations $\langle k^+ \rangle$ received on average by the I-cells are given by

$$\langle k^- \rangle = p N_I P_I(n-1) \quad (14)$$

$$\langle k^+ \rangle = p N_E P_E(n) \quad (15)$$

Instead of considering a random network as above, a particular connectivity matrix can be taken into account. Let us

denote by $W_{ij}^{E \rightarrow I}$ the binary weight from E-cell i to I-cell j ($=1$ if i and j are connected and 0 otherwise). Similarly, $W_{ij}^{I \rightarrow E}$ is the binary weight from I-cell i to E-cell j . Let us denote $S_i^{E,I}(n) \in \{0, 1\}$ the binary activity of a neuron i (excitatory or inhibitory) at the oscillation cycle n . An E-cell j generates $S_j^E = 0$ or 1 stochastically according to Eq. (9) in which

$$k^-(n-1) = \sum_{i=1}^{N_I} W_{ij}^{I \rightarrow E} S_i^I(n-1)$$

$$\langle k^- \rangle = p \sum_{i=1}^{N_I} S_i^I(n-1)$$

An I-cell j generates $S_j^I = 0$ or 1 stochastically according to Eq. (12) in which

$$k^+(n) = \sum_{i=1}^{N_E} W_{ij}^{E \rightarrow I} S_i^E(n)$$

$$\langle k^+ \rangle = p \sum_{i=1}^{N_E} S_i^E(n)$$

Note that a deterministic dynamics can also be obtained by using neurons whose activation functions are thresholded versions of Eqs. (9) and (12).

Acknowledgements This work was funded by the European Network of Excellence GOSPEL (<http://www.gospel-network.org>). The author thanks A. Tonnelier for stimulating discussions.

References

- Bazhenov M, Stopfer M, Rabinovich M, Huerta R, Abarbanel H, Sejnowski T, Laurent G (2001). Model of transient synchronization in the locust antennal lobe. *Neuron* 30:553–567
- Börgers C, Kopell N (2003) Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neur Comput* 15:509–538
- Buonviso N, Amat C, Litaudon P, Roux S, Royet J.-P, Farget V (2003) Rhythm sequence through the olfactory bulb layers during the time window of a respiratory cycle. *European J Neurosci* 17:1811–1819
- Chow C, White J, Ritt J, Kopell N (1998) Frequency control in synchronized networks of inhibitory neurons. *J Comput Neurosci* 5:407–420
- Cortes C, Vapnik V (1995) Support vector networks. *Machine Learning* 20:1–25
- Davison A, Feng J, Brown D (2003) Dendrodendritic inhibition and stimulated odor responses in a detailed olfactory bulb network model. *J Neurophysiology* 90:1921–1935
- Eeckman F, Freeman W (1991) Asymmetric sigmoid non-linearity in the rat olfactory system. *Brain Res* 557:13–21
- Ermentrout B (1996) Type 1 membranes, phase resetting curves, and synchrony. *Neural Comput* 8:979–1001
- Friedrich R, Habermann C, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7:862–871
- Friedrich R, Laurent G (2001) Dynamic optimization of odor representations by slow temporal patterning of mitral cells activity. *Sci* 291:889–894
- Gálan R, Sachse S, Galizia C, Herz A (2004) Odor-specific attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput* 16:999–1012
- Gerstner W, Kistler W (2002) *Spiking Neuron Models - Single Neurons, Populations, Plasticity* Cambridge Univ. Press.
- Graham B, Willshaw D (1995) Improving recall from an associative memory. *Biol Cybernet* 72:337–346
- Hendin O, Horn D, Tsodyks M (1997) The role of inhibition in an associative memory model of the olfactory bulb. *J Comput Neurosci* 4:173–182
- Hirase H, Recce M (1996) A search for the optimal thresholding sequence in an associative memory. *Network: Computation in Neural Systems* 7:741–756
- Holub A, Laurent G, Perona P (2002) A digital antennal lobe for pattern equalization: analysis and design, Volume 15. *Advances in Neural Information Processing Systems*
- Hoppenstead F, Izhikevich, E (2002) *Canonical neural models in Brain theory and neural networks*, 2nd Ed. MIT Press, Cambridge, MA.
- Horn D, Usher M (1991) Parallel activation of memories in an oscillatory neural network. *Neural Comput* 3:31–43
- Horn D, Usher M (2003) Neural networks with dynamical thresholds. *Phys Rev A* 40:1036–1044
- Hosler J, Buxton K, Smith B (2000) Impairment of olfactory discrimination by blockade of gaba and nitric oxide activity in the honey bee antennal lobes. *Behavioral Neurosci* 114:514–525
- Kashiwadani H, Sasaki Y, Uchida N, Mori K (1999) Synchronized oscillatory discharges of mitral/tufted cells with different molecular receptive ranges in the rabbit olfactory bulb. *J Neurophysiology* 82:1786–1792
- Lagier S, Carleton A, Lledo P-M (2004) Interplay between local gabaergic interneurons and relay neurons generate γ oscillations in the rat olfactory bulb. *The J Neurosci* 24:4382–4392
- Laurent G (1996) Dynamical neural assemblies. *TINS* 19:489–496
- Laurent G (1999) A systems perspective on early olfactory coding. *Sci* 286:723–728
- Laurent G, Davidowitz H (1994) Encoding of olfactory information with oscillating neural assemblies. *Sci* 265:1872–1875
- Laurent G, Stopfer M, Friedrich R, Rabinovich M, Volkovskii A, Abarbanel H (2001). Odor coding as an active, dynamical process: experiments, computation and theory. *Annu. Rev. Neurosci* 24:263–297
- Laurent G, Wehr M, Davidowitz D (1996) Temporal representations of odors in an olfactory network. *The J Neurosci* 16:3837–3847
- Li Z, Hertz J (2000) Odor recognition and segmentation by a model olfactory bulb and cortex. *Network: Comput in Neural Systems* 11:83–102
- Li Z, Hopfield J (1989). A model of the olfactory bulb and its oscillatory processing. *Biol Cybernet* 61:379–392
- Linster C, Masson C (1996) A neural model of olfactory sensory memory in the honeybee's antennal lobe. *Neural Computation* 8:94–114
- Margrie T, Schaefer A (2003) Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J Physiol* 546:363–374
- Martinez D (2005) Oscillatory synchronization requires precise and balanced feedback inhibition in a model of the insect antennal lobe. *Neural Comput*, In press
- Ng M, Roorda R, Lima S, Zemelman B, Morcillo P, Miesenbock G (2002) Transmission of olfactory information between three populations of neurons in the antennal lobe of the fly. *Neuron* 36:463–474
- Nusser Z, Kay L, Laurent G, Homanics G, Mody I (2001) Disruption of gaba_A receptors on gabaergic interneurons leads to increased oscillatory power in the olfactory bulb network. *J Neurophysiol* 86:2823–2833
- Palm G (1980) On associative memory. *Biol Cybernet* 36:19–31
- Perez-Orive J, Mazor, Turner G, Cassenaer S, Wilson R, Laurent G (2002) Oscillations and sparsening of odor representations in the mushroom bodies. *Sci* 297:359–365
- Quenet B, Horn D (2003). The dynamic neural filter: a binary model of spatiotemporal coding. *Neural Comput* 15:309–329
- Raman B, Gutierrez-Osuna R (2004) Chemosensory processing in a spiking model of the olfactory bulb: chemotopic convergence and center surround inhibition. *Advances in Neural Information Processing Systems* 17 (NIPS 2004), Vancouver, BC, December, 2004 pp 13–16
- Ravel N, Chabaud P, Martin C, Gaveau V, Hugues E, Tallon-Baudry C, Bertrand O, Gervais R (2003) Olfactory learning modifies the expression of odour-induced oscillatory responses in the gamma

- (60-90 Hz) and beta (15-40 Hz) bands in the rat olfactory bulb. *European J Neurosci* 17:350-358
- Rubin J, Bose A (2004) Localized activity patterns in excitatory neuronal networks. *Networks: Comput Neural Syst* 15:133-158
- Sachse S, Galizia C (2003) The coding of odour-intensity in the honeybee antennal lobe: local computation optimizes odour representation. *European J Neurosci* 18:2119-2132
- Schein A, Saul L, Ungar L (2003) A generalized linear model for principal component analysis of binary data. *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics (matlab code available at <http://www.cis.upenn.edu/~ais>)* pp 14-21
- Schwencker F, Sommer F, Palm G (1996) Iterative retrieval of sparsely coded associative memory patterns. *Neural Networks* 9:445-455
- Stopfer M, Bhagavan S, Smith B, Laurent G (1997) Impaired odor discrimination on desynchronization of odor-encoding neural assemblies. *Nat* 390:70-74
- Stopfer M, Jayaraman V, Laurent G (2003) Intensity versus identity coding in the olfactory system. *Neuron* 39:991-1004
- Stopfer M, Laurent G (1999) Short-term memory in olfactory network dynamics. *Nat* 402:664-668
- Strausfeld N, Hildebrand J (1999) Olfactory systems: common design, uncommon origins? *Curr Opin in Neurobiol* 9:634-639
- Wehr M, Laurent G (1996) Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nat* 384:162-166
- Willshaw D, Buneman O, Longuet-Higgins H (1969) Non-holographic associative memory. *Nat* 222:960-962