ORIGINAL PAPER

Biological Cybernetics

Techniques for temporal detection of neural sensitivity to external stimulation

Francisco B. Rodríguez · Ramón Huerta

Abstract We propose a simple measure of neural sensitivity for characterizing stimulus coding. Sensitivity is defined as the fraction of neurons that show positive responses to n stimuli out of a total of N. To determine a positive response, we propose two methods: Fisherian statistical testing and a data-driven Bayesian approach to determine the response probability of a neuron. The latter is non-parametric, data-driven, and captures a lower bound for the probability of neural responses to sensory stimulation. Both methods are compared with a standard test that assumes normal probability distributions. We applied the sensitivity estimation based on the proposed method to experimental data recorded from the mushroom body (MB) of locusts. We show that there is a broad range of sensitivity that the MB response sweeps during odor stimulation. The neurons are initially tuned to specific odors, but tend to demonstrate a generalist behavior towards the end of the stimulus period, meaning that the emphasis shifts from discrimination to feature learning.

F. B. Rodríguez (✉) · R. Huerta
Grupo de Neurocomputación Biológica (GNB),
Dpto. de Ingeniería Informática., Escuela Politécnica Superior,
Universidad Autónoma de Madrid, 28049 Madrid, Spain
e-mail: f.rodriguez@uam.es

R. Huerta
Institute for Nonlinear Science, UCSD,
9500 Gilman Dr., LaJolla, CA 92093-0402, USA
e-mail: rhuerta@ucsd.edu

1 Introduction

With recent advances in multi-unit recording techniques, there is a growing need for quantitative methods that can accurately detect stimulus-specific neural activity. In this paper, we propose a quantification of the sensitivity of neural responses to external stimuli. We consider two approaches: one based on Fisherian statistical testing, and another estimating a bound on the probability of a neural response.

The concept of sensitivity can be better understood by the contrast between "specialist" and "generalist" of neurons (Wilson et al. 2004). Specialists respond to specific stimuli and are vital for discrimination, when, for example, the animal needs to decide between odors. Generalists, on the other hand, respond to multiple stimuli and play a key role in extracting and discovering common features.

We formulate sensitivity measures and apply them to data provided by Perez-Orive et al. (2002) to assess the sensitivity of Kenyon cell responses over time in the mushroom body (MB). This area is the second processing layer in the insect brain and is known to be critical for memory formation and the discrimination of odors (de Belle and Heisenberg 1994; Zars et al. 2000; Pascual and Preat 2001; Dubnau et al. 2001; Connolly et al. 1996; Heisenberg 2003; Huerta et al. 2004; Nowotny et al. 2005).

The insect olfactory pathway starts at the antenna, where a massive number of receptors encode the odor stimulus in a high-dimensional code. In locusts, this number is approximately 90,000. This information is then sent to the antennal lobe (AL) for additional processing. In the locust AL approximately 1,000 neurons perform this task. The AL exhibits complex dynamics produced by the interaction of its excitatory and inhibitory neural populations (Daly et al. 2004; Galan et al. 2004; Laurent 2002). The excitatory cells are called projection neurons because only they transmit the
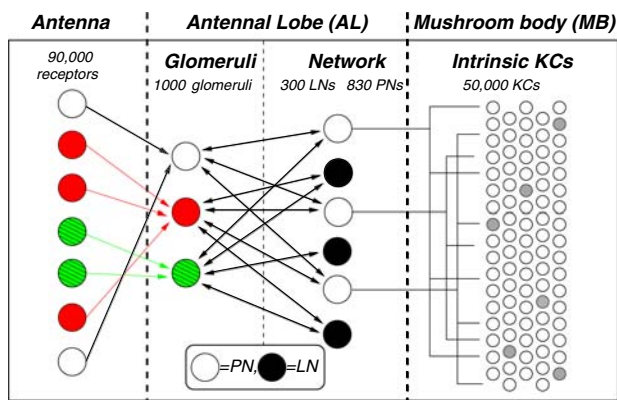
Fig. 1 Description of the structural organization of the first few processing layers in the olfactory system of the locust

Fig. 2 Stimulation protocol: the first odor pulse was initiated at $t = 3$ s and lasted for 1 s. Each odor was repeated ten times. We estimate the probability of response to spike recordings of 43 Kenyon cells (KC) from 6 experiments reported in Perez-Orive et al. (2002). The response of each KC was recorded in the presence of each of the 17 odors. The time frames Δt_3 and $\Delta t_{3.5}$ are explained in Sect. 4

result of AL processing to deeper regions. The projection neurons deliver the AL output to the 50,000 cells of the MB. This is the area where the data under analysis was obtained (Perez-Orive et al. 2002). We present a schematic representation of the insect olfactory pathway in the locust in Fig. 1.

The paper is organized as follows. First, we outline the analyzed data set. Next, we define neural sensitivity and describe a method to estimate the response probability, which is required in quantifying the sensitivity. Finally, we apply the proposed method to real data from the olfactory system of the locust and show how it can be used to analyze the temporal evolution of sensitivity.
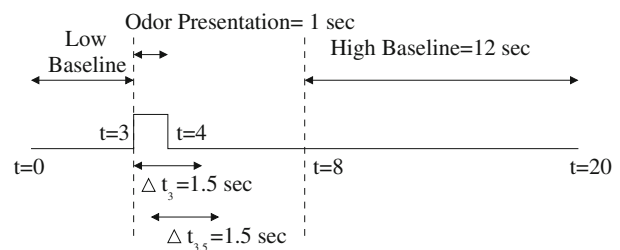
2 Description of experimental data

The recorded MB neurons are the Kenyon cells (KCs). In the experimental setup, 17 different odors were presented to the antenna of the locust for 1 s each. On average each odor was presented ten times and odor presentations were spaced 20 s apart (see Fig. 2 for a typical trial time protocol in the odor presentation). Tetrode recordings were obtained from KCs and a spike-sorting algorithm was used to detect individual spikes (Perez-Orive et al. 2002).

3 Sensitivity and detection of the neural response

We define sensitivity as the distribution of neurons that respond to n out of N stimuli. This definition requires the detecting whether or not a neuron responds to a particular stimulus. The straightforward approach is to use statistical testing by discriminating between two mutually exclusive hypotheses: response and the absence of it.

To code the neural output, we count the number of neural spikes $s_1, s_2, \ldots, s_n$ in n different presentations of the same odor in a given time window. Thus, the observable

is the neural activity represented by the number of spikes observed in certain equal-length intervals. The probability distribution of the number of spikes, $s_1, s_2, \ldots, s_n$, depends on whether a neuron is responsive to a presented stimulus $r = R$ or not $r = \overline{R}$. The null hypothesis $H_0 : r = \overline{R}$ is the default activity which is the case when the neuron does not respond. The alternative hypothesis, $H_1 : r = R$, represents the response to the stimulus. We now consider testing the simple null hypothesis $H_0 : r = \overline{R}$ versus an alternative hypothesis $H_1 : r = R$. Now, we have a decision problem between two alternative hypotheses.

Let us discuss different approaches that can help by tackling the problem from different angles. The simplest method is the Fisherian test (Christensen 2005), which is known as the standard tail test of the null hypothesis. It is possible to use this because the probability distribution of the null hypothesis, $P(s_1, s_2, \ldots, s_n | r = \overline{R})$, is known. Basically, by using this test, one can examine whether the number of neural spikes $s_1, s_2, \ldots, s_n$ is uncommon or infrequent with respect to what is measured as normal to a given confidence level α. The results of this test are in Sect. 6 and will be compared to others.

One can argue that this test misses all the information available during the stimulation window, although Fisherian supporters would say this extra information is not necessary. There is some debate about Fisherian testing versus other tests (Christensen 2005), nevertheless it is not our goal to enter the debate here.

The second statistical test, we consider derived from the Neyman–Pearson lemma (Neyman and Pearson 1933). Unfortunately, this test cannot be used on our experimental data but is worth explaining to state what is missing in our formulation. Neyman and Pearson show that the optimal test function is the ratio of the likelihood functions corresponding to the hypotheses. This statistical test is named the likelihood ratio test and in terms of our problem can be formulated as

$$\lambda = \frac{L(r = R | s_1, s_2, \ldots, s_n)}{L(r = \overline{R} | s_1, s_2, \ldots, s_n)} \geq \lambda_c.$$

We define the probability of having $s_1, s_2, \ldots, s_n$ responses in n different presentations of the same odor when we know for a fact that there is response to the presented odor as $P(s_1, s_2, \ldots, s_n | r = R)$. The likelihood function $L(r = R | s_1, s_2, \ldots, s_n)$ corresponds to this conditional probability function, $P(s_1, s_2, \ldots, s_n | r = R)$, considered as a function of its second argument with its first argument held fixed in the observed sample of the experiment. We can use the analogous conditional probability, $P(s_1, s_2, \ldots, s_n | r = \overline{R})$, in order to calculate the likelihood function $L(r = \overline{R} | s_1, s_2, \ldots, s_n)$ in the context of the null hypothesis. Within traditional tail-testing, we can reject the null hypothesis when $\lambda \geq \lambda_c$, where λ_c is determined by a confidence level $P(\lambda \geq \lambda_c | \overline{R}) = \alpha \leq 1/\lambda_c$ (Dempster 1997).

In order to apply the likelihood ratio test, we need to calculate the likelihoods of the hypotheses. We need to estimate the conditional probabilities $P(s_1, s_2, \ldots, s_n | r = \overline{R})$ and $P(s_1, s_2, \ldots, s_n | r = R)$. While we can comfortably assume that $P(s_1, s_2, \ldots, s_n | r = \overline{R})$ is equivalent to the baseline activity, the density $P(s_1, s_2, \ldots, s_n | r = R)$ is out of reach because it is not necessarily possible to force a neural response by means of external stimulation. This is the main problem that makes impracticable the application of a likelihood ratio test to decide if a neuron is responsive or not to a given stimulus.

To avoid these problems, we develop a different approach which is based on a data-driven Bayesian framework. The Bayesian approach enables the estimation of a bound on the probability of having a response. This is an additional advantage in contrast to straightforward Fisherian testing. In order to infer a possible positive response to perturbation, we observe some random variable of this system, i.e. an observable that we will name O. An observable is a property of the system state that can be determined by some sequence of physical measures. Therefore, the density $P(r = R | O = o, S = s)$ is the probability of a neural response given a specific observation $O = o$ of the system and the specific stimulus applied $S = s$. When this probability is close to one, the system has a positive response to the applied stimulus. The complementary probability or negative response probability, $r = \overline{R}$, to the stimulus is given by $P(r = \overline{R} | O = o, S = s) = 1 - P(r = R | O = o, S = s)$. Let us apply Bayes' theorem ($P(A|B)P(B) = P(B|A)P(A)$) to the conditional or posterior probability $P(r = \overline{R} | O = o, S = s)$ to obtain

$$
\begin{aligned}
&P(r = \overline{R} | O = o, S = s) \\
&= \frac{P(O = o | r = \overline{R}, S = s) P(r = \overline{R} | S = s)}{P(O = o | S = s)},
\end{aligned} \tag{1}
$$

where $P(r = \overline{R} | S = s)$ the "prior" probability for the non-response random variable. It is "prior" in the sense that it does not take into account any information about the stimulus. The estimation of $P(O = o | r = \overline{R}, S = s)$ represents

the probability distribution of the observable when we know for a fact that the system does not respond to the stimulus. At this point, we count on the probability distribution of the observable in the absence of the stimulus (i.e. $S = no$), which is the resting state or baseline activity of the system and is available in our data set. Then, we can write $P(O = o | r = \overline{R}, S = s) = P(O = o | S = no)$. This equality follows from two assumptions: if the system is not responding to a specific stimulus, it may be because (i) the stimulus is actually absent (i.e. $P(O = o | r = \overline{R}, S = s) = P(O = o | r = \overline{R}, S = no)$); or, (ii) because the observation o is indeed irrelevant whether the system responds or not in the absence of stimulus, which is to say that $P(O = o | r = \overline{R}, S = no) = P(O = o | S = no)$.

In a great variety of problems, another unknown probability which is difficult, if possible, to estimate is the "prior" probability $P(r = \overline{R} | S = s)$. However, we know that $P(r = \overline{R} | S = s) \leq 1$ such that

$$
P(r = \overline{R} | O = o, S = s) \leq \frac{P(O = o | S = no)}{P(O = o | S = s)}. \tag{2}
$$

Finally, we can express the complementary probability for no response as

$$
P(r = R | O = o, S = s) \geq 1 - \frac{P(O = o | S = no)}{P(O = o | S = s)}. \tag{3}
$$

We name the right-hand side of this equation as the lower bound of the probability of positive responses to the stimulus. In the following section we specify the lower bound for the problem which is needed to estimate the sensitivity of the neuron responses.

4 Derivation of the lower bound of the response probability to odors

We are going to describe the estimation of the response probability for the example of insect odor stimulation. The probabilistic framework can be adapted to detect generic neural responses to any external stimulation. Let us start by defining the window where we can measure the conditional response to a external stimulation. This probability estimation will be used later to determine whether a neuron responds or not to the stimulus and will be compared to Fisherian testing and to another standard method.

A given neuron can spike s times in a time window of Δt_p seconds where $s = [0, \infty)$, and the index p denotes the time placement of the window measured in seconds from the beginning of the trial. For instance, in Fig. 2, two examples of $\Delta t_p = 1.5$ s located at times $p = 3$ and 3.5 s, respectively (Δt_3 and $\Delta t_{3.5}$) are shown. To determine whether a given activity is the result of noise or indeed due to the applied stimulus, we first analyze the baseline activity of

individual neurons. We denote by R the event that represents the response of the neuron to the odor, and $\overline{R}$ is the lack of response to the external stimulation. We then proceed to estimate the probability of having s responses given that the odor is not present for all the KCs under analysis. We calculate the probability that a neuron responds given a specific odor and a set of spike observations $s_1, s_2, \ldots, s_n$ in n different presentations of the odor, that is, $P(R|s_1, s_2, \ldots, s_n, \text{odor})$. Hence *odor* in the formula represents a particular odor from the space of all possible odorant stimuli. As we will see below it is easier to tackle the problem if we calculate $P(\overline{R}|s_1, s_2, \ldots, s_n, \text{odor})$ and then use it to estimate a lower limit of the response probability.

In order to calculate the probability $P(\overline{R}|s_1, s_2, \ldots, s_n, \text{odor})$, we apply Bayes' theorem ($P(A|B)P(B) = P(B|A)P(A)$) to obtain

$$P(\overline{R}|s_1, s_2, \ldots, s_n, \text{odor})$$
$$= \frac{P(s_1, s_2, \ldots, s_n|\overline{R}, \text{odor}) \, P(\overline{R}|\text{odor})}{P(s_1, s_2, \ldots, s_n|\text{odor})}.$$

As explained in the previous section, the estimation of $P(s_1, \ldots, s_n|\overline{R}, \text{odor})$ captures the activity of the neuron when we know for a fact that there is no response to the odor. Since we count on the baseline activity, we can assume that it is equivalent to estimating $P(s_1, s_2, \ldots, s_n|\overline{R}, \text{odor} = \text{baseline})$. In other words, we take the baseline activity as an estimator of the typical activity to $\overline{R}$ to an odor. Let us continue by using $P(s_1, s_2, \ldots, s_n|\overline{R}, \text{odor}) = P(s_1, s_2, \ldots, s_n|\text{baseline})$. The event s_i depends on the previous events $(s_{i-1}, \ldots, s_1)$, which means that $P(s_1, s_2, \ldots, s_n|\text{odor}) = P(s_1|\text{odor}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \ldots, s_1|\text{odor})$. It is the case that for a given odor presentation we have a unique sequence $(s_n, \ldots, s_1)$ which is not enough to build an estimation of the conditional probability. If the stimulus presentations are well separated in time the assumption of statistical independence can be appropriate. Nonetheless, there can be some other processes in the system as a result of adaptation and learning that can carry over correlations. In order to account for this situation we can use a first order dependence $P(s_1, s_2, \ldots, s_n|\text{odor}) = P(s_1|\text{odor}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{odor})$. As we will show on the analyzed data set, there may be no need for this addition, but we include it in the formulation for the sake of the generality of the response probability derivation. One could use higher order dependencies as well, but this requires larger amounts of data to populate the conditional probabilities.

Using the first order dependence we obtain:

$$P(\overline{R}|s_1, s_2, \ldots, s_n, \text{odor})$$
$$= \frac{P(s_1|\text{baseline}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{baseline}) \, P(\overline{R}|\text{odor})}{P(s_1|\text{odor}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{odor})}.$$
$$(4)$$

The probability $P(s_i|s_{i-1}, \text{odor})$ will be estimated for different time window sizes, Δt_p, from different starting points after the odor presentation in 10 different trials of the same odor. The a priori probability of having a response given a specific odor, $P(R|\text{odor}) = 1 - P(\overline{R}|\text{odor})$, is unknown. However, we do not need this prior, because the lower bound on the probability distribution masks it. We know that

$$P(\overline{R}|s_1, s_2, \ldots, s_n, \text{odor})$$
$$\leq \frac{P(s_1|\text{baseline}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{baseline})}{P(s_1|\text{odor}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{odor})}. \quad (5)$$

Since we want to determine whether a given KC responds to the odor, we fix a probability response bound p_r such that

$$P(R|s_1, s_2, \ldots, s_n, \text{odor}) \geq p_r,$$

which together with (5) and the fact $P(R|s_1, s_2, \ldots, s_n, \text{odor}) = 1 - P(\overline{R}|s_1, s_2, \ldots, s_n, \text{odor})$ leads to

$$P(R|s_1, s_2, \ldots, s_n, \text{odor})$$
$$\geq 1 - \frac{P(s_1|\text{baseline}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{baseline})}{P(s_1|\text{odor}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{odor})} \geq p_r.$$

As a result, every KC that fulfills the above inequality is declared as being responsive to the odor with at least p_r probability value. We define the lower bound estimator of $P(R|s_1, s_2, \ldots, s_n, \text{odor})$ as:

$$\Phi(R|s_1, s_2, \ldots, s_n, \text{odor})$$
$$= 1 - \frac{P(s_1|\text{baseline}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{baseline})}{P(s_1|\text{odor}) \prod_{i=2}^{n} P(s_i|s_{i-1}, \text{odor})} \geq p_r.$$
$$(6)$$

In other words, first we set the probability response bound p_r and then we build up the probability distributions in the absence of the odorant $P(s_i|s_{i-1}, \text{baseline})$ and under its presence $P(s_i|s_{i-1}, \text{odor})$. Then, we calculate $\Phi(R|s, \text{odor})$ at a specific odor presentation for the given KC.

For our particular data set, the correlations from one presentation to the next are very small. The time separation between presentations is long enough and the activity of the neurons is very low. In addition, when the Kolmogorov–Smirnov test is run to reject the hypothesis that $P(s_i|s_{i-1}) = P(s_i)$ for all the KCs we do not find a KC that can reject the hypothesis at the 10% significance level. In fact, the average p-value is 0.8 for the baseline activity and 0.7 during odor stimulation.

Therefore, the inequality (6) can be expressed in a simpler form as

$$\Phi(R|s_1, s_2, \ldots, s_n, \text{odor})$$
$$= 1 - \frac{\prod_{i=1}^{n} P(s_i|\text{baseline})}{\prod_{i=1}^{n} P(s_i|\text{odor})} \geq p_r. \quad (7)$$

The main advantage of the lack of correlations across trials is that the conditional probabilities can be populated with
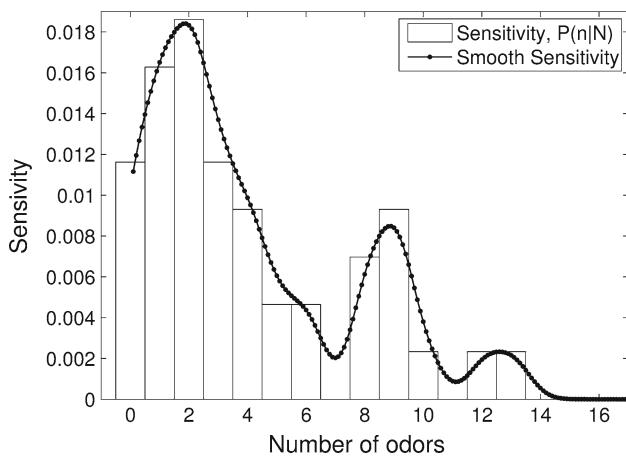
Fig. 3 Sensitivity curve calculated from the recorded KCs when 17 different odors were presented to the antennae of the locust for 1 s each (see Sect. 2). The discrete sensitivity distribution for a specific high baseline value of 12 s is presented as *bars*. We used for this analysis a time window of $\Delta t_3 = 1$ s. The response bound was $p_r = 0.99$. It is the lower bound of the probability of having the neuron respond to a given odor. We use a Gaussian distribution with dispersion equal to 0.6 to smooth the discrete sensitivity distribution

Fig. 4 Sensitivity surface as a function of the time window Δt_3. We use a Gaussian distribution with dispersion equal to 0.6 to smooth the sensitivity distribution. One can see the difference between the transient stage (for values of Δt_3 between 0.1 and 1.0 s) and the stationary state (see *blow up* in figure for values of Δt_3 between 1.0 and 2.0 s)

fewer trials. To calculate the lower bound estimator in the above equation we need to estimate the probability distributions $P(s_i|\text{baseline})$ and $P(s_i|\text{odor})$. Although it is common to assume that neurons follow a Poisson process, the MB neurons do not follow a Poisson distribution. As suggested by Kass et al. (2005) and Brown et al. (2004), instead of assuming a parametric probability distribution (as used, for example, in Bayesian decoding schemes) we use direct bootstrapping to obtain non-parametric distributions.

Bayesian bootstrap with non informative priors has been used to estimate posterior distributions (Donald 1981). Parametric probability distributions can be handled theoretically, but for this paper, we opted out of that path and decided to pursue the avenue opened by Kass et al. (2005). In Fig. 3 we calculate a typical sensitivity curve from the recorded KCs when 17 different odors were presented to the antennae of the locust for 1 s each (see Sect. 2). As we defined in Sect. 3, neural sensitivity is the probability that a neuron responds to n out of N odors, $P(n|N)$. In other words, $P(n|N)$ is the number of neuron responses divided by the total number of available neurons in the experiment. The neuron responses used in this sensitivity curve have been calculated using the above formalism for a specific time window of $\Delta t_3 = 1$ and probability response bound $p_r = 0.99$. In this example, the maximal sensitivity corresponds to a two-odor response. The next significant bump in this curve corresponds to a positive response of 9 odors out of the 17 presented during the experiment. One can see that this event is less likely than the response to two odors. Basically, we can see that most of the KCs are sensitive to two odors, while there are others that have a broader tuning.
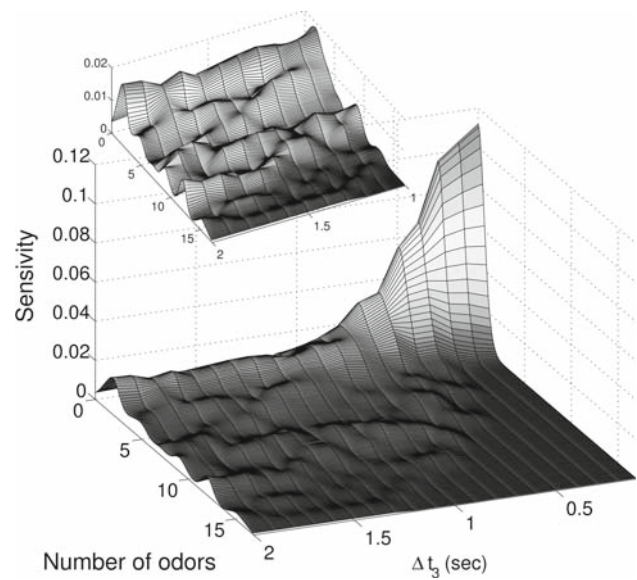
5 Estimating the conditional response to Kenyon cell data

We have used several KC recordings published in Perez-Orive et al. (2002) in order to build a distribution of the sensitivities of the KCs to odor's presentation. Each neuron responded to a percentage of the presented odors. At every data point in the sensitivity distribution, we placed a Gaussian distribution to densely populate the curve. We analyzed spike recordings of 43 KCs from 6 experiments, in which 17 odors were presented to the animal. For each odor 10 trials were recorded.

The goal of the analysis is to determine the level of sensitivity of the KCs as a function of the time interval Δt_p (see Sect. 4). In Fig. 4 we show that for small sizes of Δt_3 (between 0.1 and 1.0 s) there is a transient in the sensitivity curves. At the early stages of the odor presentation the KCs are very selective. After half-a-second, the KCs widen their selectivity to quite a few odors. Thus the response curves become quite stationary for sizes, Δt_3, between 1.0 and 2.0 sec. To analyze the transient response after odor stimulation we run the response window, Δt_p, at different starting points between $p = 3$ and $p = 4.6$ with constant $\Delta t_p = 1$. This helps in understanding the temporal coding properties of the KC response. In the initial stages, the KCs are very selective to the odors (see Fig. 5), but after half-a-second, the selectivity to odors broadens (see right panel in Fig. 5). Thus, the KCs become more sensitive to more odors as time passes. After the odor is removed, the selectivity distribution returns

$$\prod_{i=1}^{n} P(s_i|\text{odor}) = 4.089897e{-}07$$

$$\Phi(R|s_1, s_2, \ldots, s_n, \text{odor}) = 1 - \frac{\prod_{i=1}^{n} P(s_i|\text{baseline})}{\prod_{i=1}^{n} P(s_i|\text{odor})}$$

$$\prod_{i=1}^{n} P(s_i|\text{baseline}) = 0.0$$
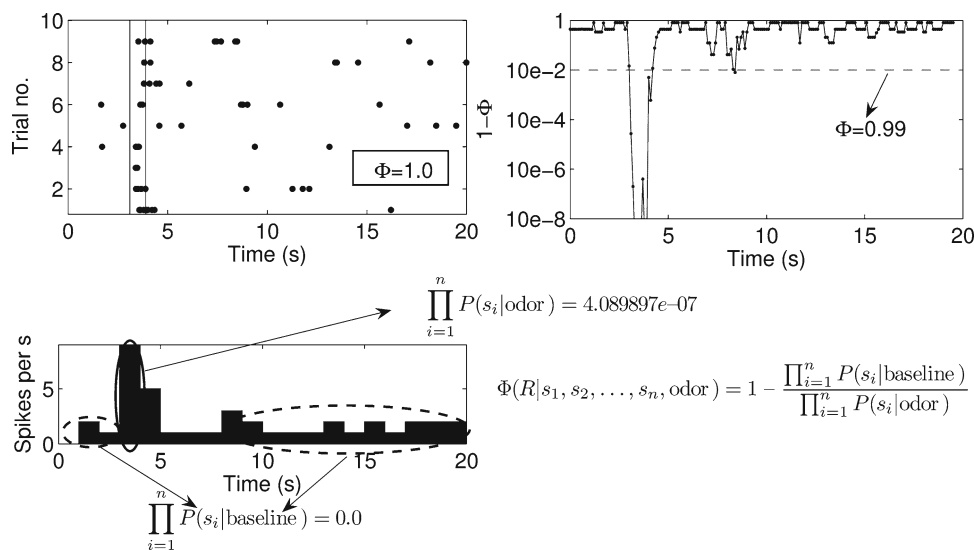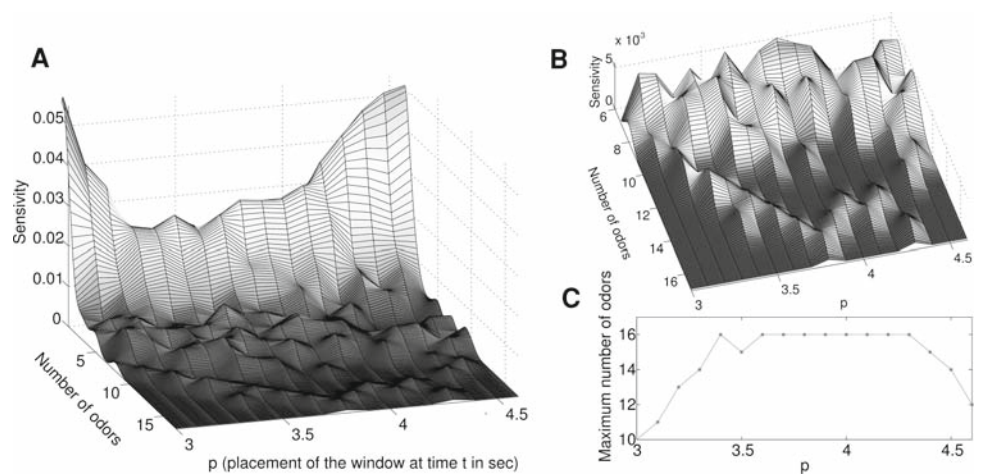
Fig. 6 An example of a spike raster for one KC. The time window used was $\Delta t_3 = 1$ s (indicated by *vertical bars*). In the PSTHs, we show the probability distributions used to calculate the lower bound estimator. Φ on (7). The values of the factorized probabilities can be seen in Table 1. These probabilities are estimated by means of a bootstrap procedure (see for details Efron 1979). This neuron responds significantly for this particular odor with the lower bound of the probability response to 1.0.

The *right panel* is a baseline control for the same neuron. $1 - \Phi$ is represented on a logarithmic scale as a function of the placement of the time window. It can be seen that the maximal reliability occurs precisely on the window of the stimulation. The *dashed line* corresponds to a probability response bound of $p_r = 0.99$. All the points that are below this line are classified as responses

slowly to normal. The reason is that the AL is known to produce complex spatio-temporal patterning after odor removal (see Mazor and Laurent 2005 for details).

5.1 An example of the calculation of neuron response probability

To illustrate the neural response calculation, let us show an example of our method applied to a specific neuron. In this example, we choose a time window of $\Delta t_3 = 1$. In Fig. 6, the spike raster and its peristimulus time histograms (PSTHs) are depicted. The calculation for the lower bound estimator of response probability gives us a value of

$\Phi(R|s_1, s_2, \ldots, s_n, \text{odor}) = 1$, which means that this KC responds with high probability to the specific odor presentation. One can see that the KC is significantly excited by the stimulus. In the corresponding PSTHs, one can see the probability distributions used to obtain the lower bound estimator, Φ. The factors of these probability distributions of (7) are calculated by means of a bootstrap procedure and some of them are shown in Table 1. The right panel in the same figure represents the control for the baseline activity. We calculate the lower bound estimator $(1-\Phi)$ as a function of the placement of the time window response at time t. The maximal probability of the response is precisely situated on the stimulation period.

Table 1 Non-parametric probability distributions calculated through a bootstrap procedure from the spike rasters shown on Fig. 6

# of spikes in $\Delta t_3 = 1s$	$P(s|\text{odor})$	$P(s|\text{baseline})$
$s = 0$	0.1019	0.846327
$s = 1$	0.2045	0.119967
$s = 2$	0.3976	0.026
$s = 3$	0.0962	0.00680667
$s = 4$	0.1998	0.0

We use these distributions to calculate the lower bound estimator, Φ, of $P(R|s_1, s_2, \ldots, s_n, \text{odor})$

6 Method comparison

A method commonly used to detect whether a neuron responds to a stimulus or not consists of measuring the mean firing rate over fixed periods of time during baseline activity. This set of mean firing rates is used to estimate the standard deviation of the baseline activity. For example, in Perez-Orive et al. (2002), neural baseline statistics were used as a reference to determine a yes/no response. When the stimulus is presented, if the firing rate becomes significantly higher than the baseline activity then it is said that the neuron has fired. The arbitrary part of the method is to determine how much significance is needed. The strategy consists of fixing n times the standard deviation and then by "visual inspection" determining whether it is good enough. This standard method resembles Fisherian or tail testing under the assumption that the firing distribution respect to the baseline activity is Gaussian. Then the confidence level to reject the null hypothesis (baseline activity) is $\frac{1}{2}(1 - erf(n/\sqrt(2)))$.

This method itself does not provide an estimate or probability of how accurate the arbitrary determination of n is. Rather than relying on visual inspection, the most straightforward method is Fisherian testing directly calculated from $P(s_1, s_2, \ldots, s_n | r = \overline{R})$ as mentioned in Sect. 3. In addition, one can use our lower bound estimation (7) as an alternative that provides more information about the probability of response of a neuron.

In the Fig. 7, we present the overall comparison among all methods. We use the same critical values for all methods. For our method, we chose the different probability response bounds of $p_r = 0.9, \ 0.99$. In the standard method, we use 1.281 and 2.326 standard deviations which correspond to 90 and 99% of Gaussian area respectively. Last, we have introduced the standard tail test (see Sect. 3) as the lower bound in the comparisons using values of $\alpha = 0.1, \ 0.01$. Black boxes in each panel mean that a response to the odor exists. One can see that the probability bound method is the most restrictive of the three in assessing whether a neuron responds to certain stimulus. For instance, our method gives a total of 197 responses for $p_r = 0.99$. The standard method has 236 for the equivalent case and 250 for the Fisherian test. All the responses produced by the lower bound limit for $p_r = 0.99$ are present in the Fisherian testing, while there are three cases that differ with respect to the standard method (nSDs $= 2.326$). In the same figure, it is observed that the Fisherian test is least restrictive. Remember that the Fisherian test only has information about the probability distribution of spikes in the baseline, $P(s_1, s_2, \ldots, s_n | r = \overline{R})$. In this sense, the Fisher test is a mask for the lower bound method. However, the lower bound method with additional information during the stimulation window makes the deci-
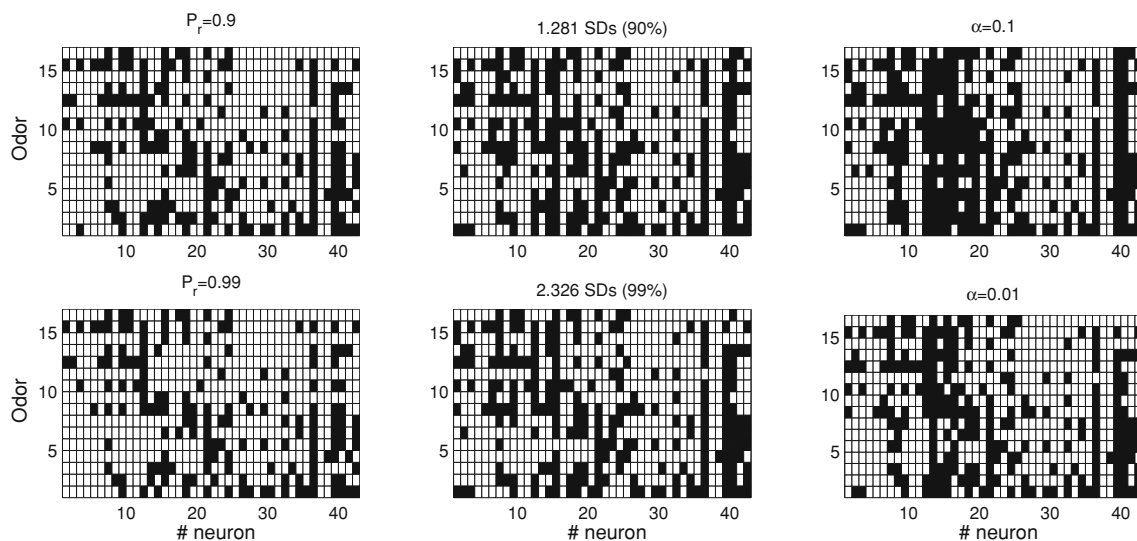
Fig. 7 Here we provide a direct comparison among the lower bound estimation ($p_r = 0.9, \ 0.99$), the standard method of n standard deviations ($nSDs = 1.281, \ 2.326$) and the Fisher test ($\alpha = 0.1, \ 0.01$). The sensitivity curve for $p_r = 0.99$ is plotted in Fig. 3. The stimulus time window used was $\Delta t_3 = 1$ sec

de Belle JS, Heisenberg M (1994) Associative odor learning in Drosophila abolished by chemical ablation of mushroom bodies. Science 263:692–695

Dempster AP (1997) The direct use of likelihood for signicance testing. Stat Comput 7:242–252

Donald BR (1981) The Bayesian bootstrap. Ann Stat 9(1):130–134

Dubnau J, Grady L, Kitamoto T, Tully T (2001) Disruption of neurotransmission in Drosophila mushroom body blocks retrieval but not acquisition of memory. Nature 411(6836):476–480

Efron B (1979) Bootstrap methods: another look at the jackknife. Ann Stat 7:1–26

Friedman HS, Priebe CE (1998) Estimating stimulus response latency. J Neurosci Methods 83(2):185–194

Galan RF, Sachse S, Galizia CG, Herz AV (2004) Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid odor classification. Neural Comput 16(5):999–1012

Heisenberg M (2003) Mushroom body memoir: from maps to models. Nat Rev Neurosci 4(4):266–275

Huerta R, Nowotny T, Garc-Sanchez M, Abarbanel HD, Rabinovich MI (2004) Learning classification in the olfactory system of insects. Neural Comput 16(8):1601–1640

Kass RE, Ventura V, Brown EN (2005) Statistical issues in the analysis of neuronal data. J Neurophysiol 94:8–25

Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. Nat Rev Neurosci 3(11):884–895

Mazor O, Laurent G (2005) Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. Neuron 48(4):661–673

Neyman J, Pearson E (1933) On the problem of the most efficient tests of statistical hypotheses. Phil Trans R Soc Lond Ser A 231:289–337

Nowotny T, Huerta R, Abarbanel HD, Rabinovich MI (2005) Self-organization in the olfactory system: one shot odor recognition in insects. Biol Cybern 93(6):436–446

Pascual A, Preat T (2001) Localization of long-term memory within the Drosophila mushroom body. Science 294(5544):1115–1117

Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G (2002) Oscillations and sparsening of odor representations in the mushroom body. Science 297(5580):359–365

Ratnam R, M Goense JB, Nelson ME (2003) Change-point detection. In neuronal spike train activity. Neurocomputing 52:849–855

Ritov Y, Raz A, Bergman H (2002) Detection of onset of neuronal activity by allowing for heterogeneity in the change points. J Neurosci Methods 122(1):25–42

Wilson RI, Turner GC, Laurent G (2004) Transformation of olfactory representations in the Drosophila antennal lobe. Science 303(5656):366–370

Zars T, Fischer M, Schulz R, Heisenberg M (2000) Localization of a short-term memory in Drosophila. Science 288(5466):672–675