

Generating sparse and selective third-order responses in the olfactory system of the fly

Sean X. Luo^a, Richard Axel^{a,b,c,1}, and L. F. Abbott^{a,d,1}

^aDepartment of Neuroscience, ^bDepartment of Biochemistry and Molecular Biophysics, ^cHoward Hughes Medical Institute, and ^dDepartment of Physiology and Cellular Biophysics, Columbia University, New York, NY 10032

Contributed by Richard Axel, April 28, 2010 (sent for review March 4, 2010)

In the antennal lobe of *Drosophila*, information about odors is transferred from olfactory receptor neurons (ORNs) to projection neurons (PNs), which then send axons to neurons in the lateral horn of the protocerebrum (LHNs) and to Kenyon cells (KCs) in the mushroom body. The transformation from ORN to PN responses can be described by a normalization model similar to what has been used in modeling visually responsive neurons. We study the implications of this transformation for the generation of LHN and KC responses under the hypothesis that LHN responses are highly selective and therefore suitable for driving innate behaviors, whereas KCs provide a more general sparse representation of odors suitable for forming learned behavioral associations. Our results indicate that the transformation from ORN to PN firing rates in the antennal lobe equalizes the magnitudes of and decorrelates responses to different odors through feedforward nonlinearities and lateral suppression within the circuitry of the antennal lobe, and we study how these two components affect LHN and KC responses.

antennal lobe | decorrelation | lateral horn | normalization | olfaction

In the olfactory system, as in other sensory systems, signals from primary receptors are processed and transformed before being relayed to higher brain areas. In *Drosophila melanogaster*, olfactory receptor neurons (ORNs) located in the antennae and maxillary palps synapse onto projection neurons (PNs) within the glomeruli of the antennal lobes (Fig. S1). ORNs expressing a given receptor converge on an anatomically invariant glomerulus. PNs innervate a single glomerulus and thus receive their primary input from ORNs expressing the same olfactory receptor (1), although cross-glomerular interactions mediated by interneurons are also present (2–5). The PNs send their axons to two distinct regions of the fly brain, the lateral horn and the mushroom body. By constructing models of neurons in these regions, we study the effects of the transformations arising from circuitry within the antennal lobe on the capacity of neurons in the lateral horn and the mushroom body to represent and discriminate odors.

Any functional interpretation of the transformation from ORN to PN responses depends on the nature of the processing being performed by the third-order neurons that receive PN input. The lateral horn and mushroom body appear to be involved in different forms of sensory processing. The lateral horn is believed to be important in generating innate behaviors (6), including those elicited by pheromones (7). PN projections to lateral horn neurons (LHNs) are spatially stereotyped (8, 9) and clustered (10), suggesting that circuits in this region may be “hardwired” for the detection of specific odors that elicit innate behavioral responses. The mushroom body is implicated in decision making (11) and in the formation of associative memories (12, 13). Both calcium imaging (14) and electrophysiology (15) indicate that the responses of Kenyon cell (KCs) in the mushroom body are sparse, with each odorant eliciting responses in a few percent of the KCs and individual KCs responding to a small number of tested odorants. Electrophysiological experiments suggest that connections from PNs to Kenyon cells may be random (16, 17), although zonal specificity of PN projections and KC dendrites in the mushroom bodies has been reported (10, 18). The distinct

characteristics of LHNs and KCs lead us to consider the implications of the ORN-to-PN transformation with respect to two different downstream tasks: odor discrimination in the lateral horn and general-purpose representation of a wide variety of odors in the mushroom body. We begin by discussing how model PN responses are constructed, then show how LHN and KC responses with the desired properties can be generated, and finally analyze the roles played by different elements of the model.

Results

We build our third-order neuron models by taking measured ORN firing rates (19, 20) (Fig. 1A) and then determining the corresponding PN responses on the basis of an experimental characterization of the transformations occurring within the antennal lobe (see below) (21, 22). Finally, we use the computed PN responses to drive model LHNs and KCs. To model the responses of LHNs or KCs, the computed firing rate for each PN in response to a particular odor is multiplied by a weight representing the strength of a synapse from that PN onto the third-order neuron being modeled. The weighted PN rates are then summed, and the total is compared to a fixed threshold (23) to determine whether or not the third-order neuron generates a response to that odor. Although selectivity is likely to arise from some form of input summation and thresholding, the computations performed by neurons in the lateral horn and the mushroom body are undoubtedly more complex than those of our model. Using a minimal model allows us to identify the problems that more complex cellular and circuit mechanisms must solve, and it allows us to focus on the role of the antennal lobe circuitry.

Transformation from ORN to PN Firing Rates. The transformation from ORN to PN responses consists of two functional components: a feedforward nonlinearity and lateral suppression. The relationship between the firing rate of an ORN and that of a PN is nonlinear with strong saturation at high ORN rates (21), a result due at least in part to short-term depression at ORN-PN synapses (24). We refer to this transformation as a feedforward nonlinearity, although lateral circuitry may also contribute to this effect. This nonlinearity was described using an exponential function, but it can also be fit by an equation of the form $r_{PN} = R_{max} r_{ORN}^{1.5} / (\sigma^{1.5} + r_{ORN}^{1.5})$, where r_{PN} and r_{ORN} are the firing rates of a corresponding PN and ORN pair, and R_{max} and σ are constants. The advantage of using this latter form is that it fits in well with more recent work in which the effects of lateral suppression have been incorporated (22). Lateral suppression, which is the dominant effect of ORNs that do not drive a given PN directly, can be described by adding a term proportional to the sum

Author contributions: S.X.L., R.A., and L.F.A. designed research, performed research, contributed analytic tools, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

¹To whom correspondence may be addressed. E-mail: ra27@columbia.edu or lfa2103@columbia.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1005635107/-DCSupplemental.

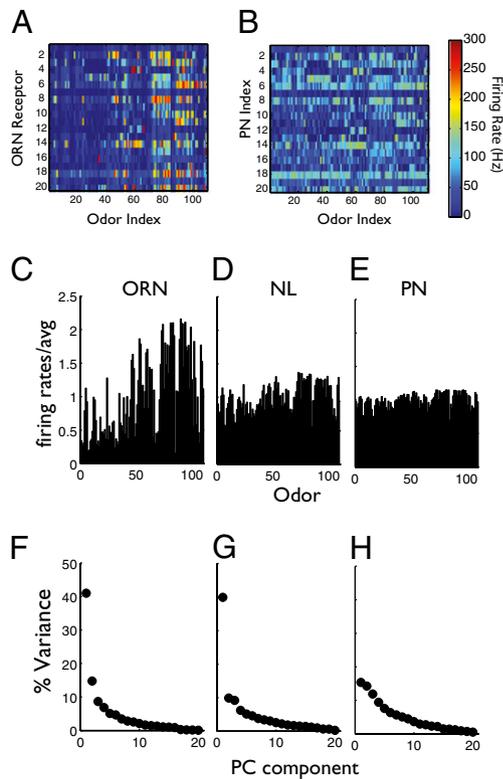


Fig. 1. ORN and PN responses. (A) A representation of the dataset we use as input to our models from Hallem and Carlson (20). Colors (scale bar) show the firing rates of ORNs expressing 20 different olfactory receptors (horizontal axis) to 110 different odorants (vertical axis). From the full dataset, we used responses generated by 20 nonspecialized receptors (we do not include pheromone receptors) to 110 tested odors. (B) The same data as in A in the same format, but converted to model PN responses through the transformation of Eq. 1. (C–H) Response magnitude equalization and decorrelation. ORN, NL, and PN indicate that ORN, PN with no lateral suppression, and PN response were used for the corresponding column. (C) Sums of ORN responses to 110 odorants from the dataset plotted relative to their average across odors. (D) Sums of PN responses to 110 odors generated by Eq. 1, but with no lateral suppression ($m = 0$ in Eq. 1). (E) Sums of PN responses to 100 odors generated by Eq. 1. (F–H) Percentage of variances from a PCA analysis of the response used in C–E.

of the firing rates of all of the ORNs to the denominator of the expression for r_{PN} . Defining s_{ORN} as the sum of all ORN firing rates (independent of the glomerulus to which they project), the equation

$$r_{PN} = \frac{R_{\max} r_{ORN}^{1.5}}{\sigma^{1.5} + r_{ORN}^{1.5} + (m s_{ORN})^{1.5}}, \quad [1]$$

where m is another constant (see *Methods* for parameter values), provides an accurate description of both the feedforward nonlinearity and lateral suppression generated by the antennal lobe circuitry (22). Eq. 1 refers to average responses, but trial-to-trial variability is included in the PN responses we use (Fig. 1B) by adding randomly generated fluctuations modeled after the data (21) (*Methods*).

Before modeling LHNs and KCs, it is instructive to examine how the transformation of Eq. 1 modifies olfactory responses. The sum of all ORN firing rates, s_{ORN} , introduced above and the analogous sum over all PN firing rates, s_{PN} , are measures of the total magnitude of the response to a given odor, which has a strong impact on LHN and KC performance. A third-order neuron tuned to respond to an odor that elicits a weak overall PN response (small s_{PN}) will tend to respond nonselectively to

odors that elicit stronger total responses (large s_{PN}). This problem can be avoided if the magnitudes of PN responses are equalized, meaning that the total activity across the population of PNs (i.e., s_{PN}) is roughly the same for each odor. We use the term equalization, rather than normalization, because normalization has been used to refer to the entire transformation of Eq. 1 in the visual system literature where such a transformation was first proposed and has been used extensively (25, 26). The transformation of Eq. 1 has a strong magnitude-equalizing effect; the sum of ORN rates (s_{ORN} , Fig. 1C) is much more variable across odors than the sum over PN rates (s_{PN} , Fig. 1E). To isolate the origin of this equalization of the response magnitude, we also computed s_{PN} when lateral suppression was eliminated by setting the parameter m in Eq. 1 to zero. The results indicate that much, but not all, of the response magnitude equalization arises from the feedforward nonlinearity (Fig. 1D).

Selectivity is also difficult to achieve if the patterns of PN activation across odors are too similar, that is, too highly correlated. We performed a principal component analysis (PCA) of the ORN and computed PN firing rates to compare their degree of correlation. For totally correlated data, a single principal component accounts for 100% of the data variance, whereas for completely uncorrelated data, all of the components account for an equal percentage of the variance. For the olfactory responses, >40% of the ORN response variance is accounted for by a single principal component projection, indicating a high degree of correlation (Fig. 1F). The maximal principal component projection for the PN responses accounts for only 15% of the total variance (Fig. 1H), so PNs are significantly decorrelated by the transformation of Eq. 1. This decorrelation is a consequence of lateral suppression in the antennal lobe, because it is almost entirely eliminated if we set the parameter m in Eq. 1 to zero (Fig. 1G). A previous mathematical characterization of PN responses did not show decorrelation (21, 27) because it did not include lateral suppression. In agreement with more recent data (22), the more accurate characterization we use here does decorrelate.

Discrimination from Odor-Specific Weights (Lateral Horn). We hypothesize that neurons in the lateral horn are connected in a manner that generates highly selective responses to specific odors, allowing them to evoke innate behaviors. To study selectivity in the lateral horn from the data we are using, we constructed 110 LHNs, each selective for a different one of the 110 odorants in the dataset. In other words, we assigned a single target odor from this set to each LHN and attempted to make it respond exclusively to this target odor. We are not suggesting that each of these odors generates an innate behavior. Instead, the model LHNs are used as a proof of principle to demonstrate how highly specific odor-selective LHNs can be constructed.

We assign synaptic weights for the LHNs using Fisher's linear discriminant, a standard technique (*Methods*). Importantly, this approach uses only PN responses to the targeted odor and general statistical properties of the PN responses across odors such as the mean and covariance to determine the weights, and avoids overfitting. The weights of inputs onto LHNs produced in this way are both positive and negative. We interpret negative weights as arising from inhibitory input to the LHN from unmodeled interneurons driven by the corresponding PN. Using these weights, we can generate LHNs tuned to respond almost exclusively to a single odorant for all 110 odors in the dataset (Fig. 2).

LHNs will respond exclusively to their target odor if the input for this odor is consistently higher than the inputs for any other odors. Histograms of the sum of 20 weighted PN inputs (Fig. 2A) for targeted (red) and nontargeted (blue) odors for all 110 LHNs over repeated trials reveal two roughly Gaussian distributions that are well separated. This allows us to divide the inputs into those producing and those not producing a response with a fixed threshold (dashed vertical line in Fig. 2A). Using this threshold,

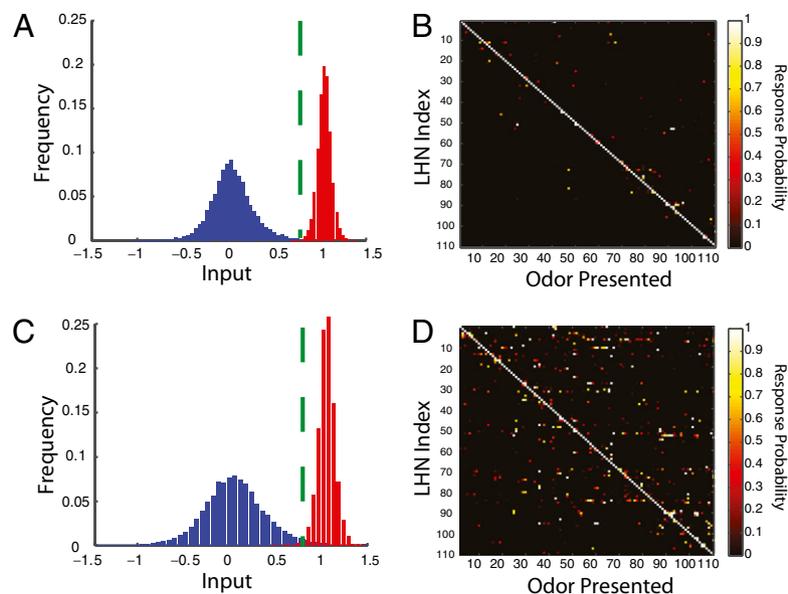


Fig. 2. Responses of model LHNs. (A) Histograms of the distribution of LHN inputs from 20 PNs. Nontargeted odors produce one distribution (blue) and targeted odors the other (red). The green dashed line indicates the response threshold. Input is in units that set the mean input for the targeted odor equal to 1. (B) Response probabilities for 110 model LHNs to 110 odors, with 20 PN inputs per LHN. Each row shows response probabilities for LHNs tuned to respond to a specific odor (indicated by the LHN index), and each column shows the responses of all LHNs to a particular presented odor. The diagonal gives the response of each LHN to its targeted odor, and the off-diagonal elements are responses to nontargeted odors. Colors denote the fraction of trials producing a response (color bar). (C) Input histogram for a reduced connectivity model in which each LHN receives input from only 5 PNs. (D) Response probabilities for 5 PN connections per LHN.

we computed response probabilities across trials (PN responses are different in each trial due to the response variability built into our PN model) for 110 differently tuned LHNs, each labeled by the odor to which it is targeted (LHN index in Fig. 2B). The diagonal elements of the matrix of response probabilities (Fig. 2B) correspond to each LHN responding to its targeted odor, and the nondiagonal entries represent responses to nontargeted odors. For perfect performance of the LHNs, response probabilities would be 100% along the diagonal and zero off the diagonal, meaning that the LHNs respond every time to the odor for which they are targeted and never to any other odor. The selectivity of the LHNs is quite close to being perfect in this sense (Fig. 2B). This excellent performance is in agreement with a theoretical analysis of the capacity of a system of this type (28).

Performance of the LHNs can also be characterized by the fraction of nontargeted odors that generate a response (the false-positive rate) and the fraction of targeted odors that fail to generate a response (the false-negative rate). These quantities depend on the threshold, but if we adjust this so that the false-positive and false-negative rates are equal, the value they take is 0.4%. The area under a receiver operating characteristic (ROC) curve provides a threshold-independent measure of discrimination performance, and this area is 0.999 for model LHNs driven by 20 PN inputs.

We also considered what happens if the number of PN inputs to the model LHNs is reduced to five, a number suggested from analysis of KCs (15). In this case, the weights are constructed again using Fisher's linear discriminant, but for a reduced set of inputs (*Methods*). With only five PN connections per LHN, the input distributions for targeted and nontargeted odors are slightly more overlapping (Fig. 2C), and there are more nonzero response probabilities for nontargeted odors (Fig. 2D). Nevertheless, the overall performance is still good. The equal false-positive and false-negative probabilities are 2%, and the area under the ROC curve is 0.987. Thus, acceptable discrimination can be obtained with as few as five PN inputs per LHN.

Sparse KC Responses from Random Weights (Mushroom Body). Our modeling work for KCs is based on the assumption that they provide a sparse representation across a wide variety of odors. Of course the LHNs, being highly selective, also provide a sparse representation, but not across nearly as broad a range of odors as the KCs do. Our goal is to construct a population of 2,500 model KCs, the number estimated in the mushroom body (29), such that on average 5% of the KCs respond to each odor, and each KC responds to ~5% of the odors presented, corresponding to experimental estimates (14, 15). We allow fluctuations around this 5% value from cell to cell or from odor to odor, but we associate large fluctuations with a low-quality sparse representation.

Our procedure for constructing a KC response is to randomly choose PNs from n glomeruli and to set the strengths of the synapses connecting them to a KC randomly and independently from a uniform distribution over positive values (*Methods*). Good performance can be achieved with such random connections (Fig. 3A) even though we used the same threshold for all of the KCs and did not individually adjust the threshold to achieve a precise 5% sparseness in each KC. The scattering of color shows that the sparse KC responses are well distributed across all 2,500 cells and 110 odorants. Fig. 3 shows KCs receiving excitatory input from PNs connected to five different glomeruli, along with global inhibition (30) (*Methods*). In this case, all 110 odors elicit responses above the 95% threshold in, on average, 125 cells and, at a minimum, at least two KCs.

One measure of the quality of the sparse representation for different numbers of PN inputs is the number of missed odors (Fig. 3B Upper). Such an odor would not be represented in the mushroom body and, presumably, could not be detected by the fly unless it was represented separately in the lateral horn. We define a neuron as responding if it receives an above-threshold input in at least 50% of trials. Even with this fairly stringent criterion, none of the 110 odors is missed for most values of n , the number of PNs connected to each KC (Fig. 3B Upper). The best performance is for $n < 15$. A high-quality sparse representation should also

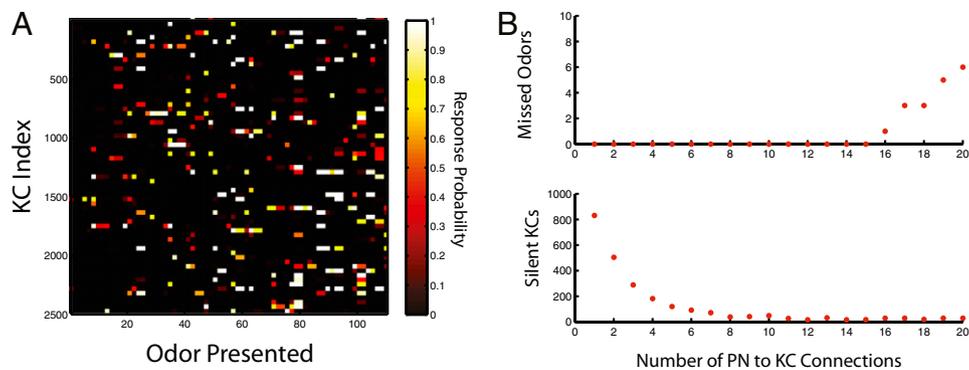


Fig. 3. Model KC responses. (A) Response probabilities of 2,500 model KCs, each receiving five random PN connections and global inhibition. All of the odors elicit responses. (B) The number of missed odors (odors producing no response across the entire KC population, *Upper*) and the number of silent KCs (KCs responding to none of the 110 odorants, *Lower*).

efficiently use its neuronal resources, which means there should not be a large number of KCs that never respond to any odor. Whereas 1 or 2 PN inputs are sufficient to obtain a good sparse representation, this leads to a large number of silent KCs (Fig. 3B *Lower*). With only 1 PN input, about 800 of the 2,500 KCs do not respond to any odor, according to the 50% criterion. The number of nonresponding KCs drops progressively as the number of PN to KC connections increases. Combining the results of Fig. 3B *Upper* and *Lower*, the best performance seems to arise for between 5 and 15 PN inputs per KC, numbers consistent with estimates obtained from electrophysiology (15).

The synaptic weights connecting PNs to our model KCs are all positive, representing pure excitation, but the model KCs also receive global inhibition (*Methods*). Inhibition is useful because it allows the KCs to take advantage of the decorrelation of the PN inputs. However, we also examined the performance of model KCs when all inhibition was removed and they were driven solely by PN input (Fig. S2). Performance is still good with regard to missed odors (Fig. S2 *Upper*), but there are significantly more silent KCs when inhibition is eliminated (Fig. S2 *Lower*).

Concentration Dependence. Changing odor concentration is a particularly stringent test of a model like ours in which responses are determined by a fixed threshold. Without response magnitude equalization across concentrations, selectivity that is present at low concentrations will be lost at high concentrations or, alternatively, neurons that are selective at high concentrations will fail to respond at low odor concentrations. Imaging and electrophysiological experiments (14, 15) indicate that at least some of the KCs retain their selectivity to the same odors across a range of different concentrations. It is reasonable to assume that at least some LHNs share this feature, as this would have obvious behavioral advantages. Hallem and Carlson measured ORN firing rates at four different concentration levels for 10 odorants (20). We studied the effects of varying concentration by replacing the ORN data used in the previous sections with this collection of 40 response sets, representing 10 odors presented at four different concentrations, 10^{-2} , 10^{-4} , 10^{-6} , and 10^{-8} dilutions.

For the model LHNs, we generalized Fisher's linear discriminant for multiple stimuli across different concentrations (*Methods*) and used five PN inputs per LHN. The performance of these model LHNs is shown as four response probability matrices (Fig. 4A), one for each concentration. In this case, the bright diagonal is recapitulated four times, illustrating that these cells respond reliably to their target odors at all of the different concentrations. The off-diagonal response probabilities, representing responses to nontargeted odors, remain small, indicating that the false-positive rate is low across concentrations. The main error comes, not surprisingly, from discrimination at low concentration, where

the false-positive rate increases somewhat. Over all odors and concentrations, the false-positive and false-negative rates are 7% and the area under the ROC curve is 0.92.

To construct KCs, we again used random positive synaptic weights with global inhibition and adjusted the single threshold used for all KCs so that 5% of the neurons respond, averaged across concentrations. Note that in contrast to previous models of concentration invariance (31), the KCs in our model do not receive compensatory inputs or change their thresholds adaptively for different concentrations. These fixed-threshold KCs respond to odors sparsely at both low (10^{-8} SV) and high (10^{-2} SV) concentrations, and an analysis of the results shown in Fig. 4B indicates that, collectively, the population responds to every odor in the dataset regardless of its concentration. In other words, none of the odors is missed across this wide range of concentrations, so a full representation of the odors is retained.

Role of the ORN-to-PN Response Transformation. To study the effect of the nonlinear transformation that takes place between the ORN and PN responses, we repeated our analysis of model LHNs and KCs but driven either by direct ORN input or by PN input with lateral suppression removed ($m = 0$ in Eq. 1).

The results are summarized in Fig. 5, and further details can be found in Figs. S3–S5.

For LHNs driven by either 20 (Fig. 5A) or 5 (Fig. 5B) inputs, the false-positive and false-negative error rates (when the threshold is adjusted to make them equal) rise only slightly when the effect of lateral suppression is removed (bars marked “NL” in Fig. 5A and B), but there is a larger increase in error percentage when the feedforward nonlinearity is also eliminated (equivalent to using ORN rather than PN input, bars marked “ORN” in Fig. 5A and B; also Fig. S3). KCs driven by PN input lacking lateral suppression perform basically as well as those using full PN input, as in neither case are any odors missed from the KC representations with 5 inputs per KC (Fig. 5C). Performance degrades significantly in the absence of the feedforward nonlinearity when ORN input is used (Fig. 5C and Fig. S4). In this case, ~30 odors are missing from the representation.

Lateral suppression, and hence decorrelation, plays a larger role in buffering the concentration dependence of LHN responses. The equal false-positive and false-negative error percentages rise more significantly when lateral interactions are eliminated, although most of the decrement in performance associated with using ORN input still arises from the absence of the feedforward nonlinearity (Fig. 5D and Fig. S5A). KCs driven by ORN rather than PN input (removing the effects of Eq. 1) fail to respond to many of the tested odors at low concentrations (Fig. 5B) and thus lose the ability to represent odors across the wide range of concentrations seen in Fig. 4A.

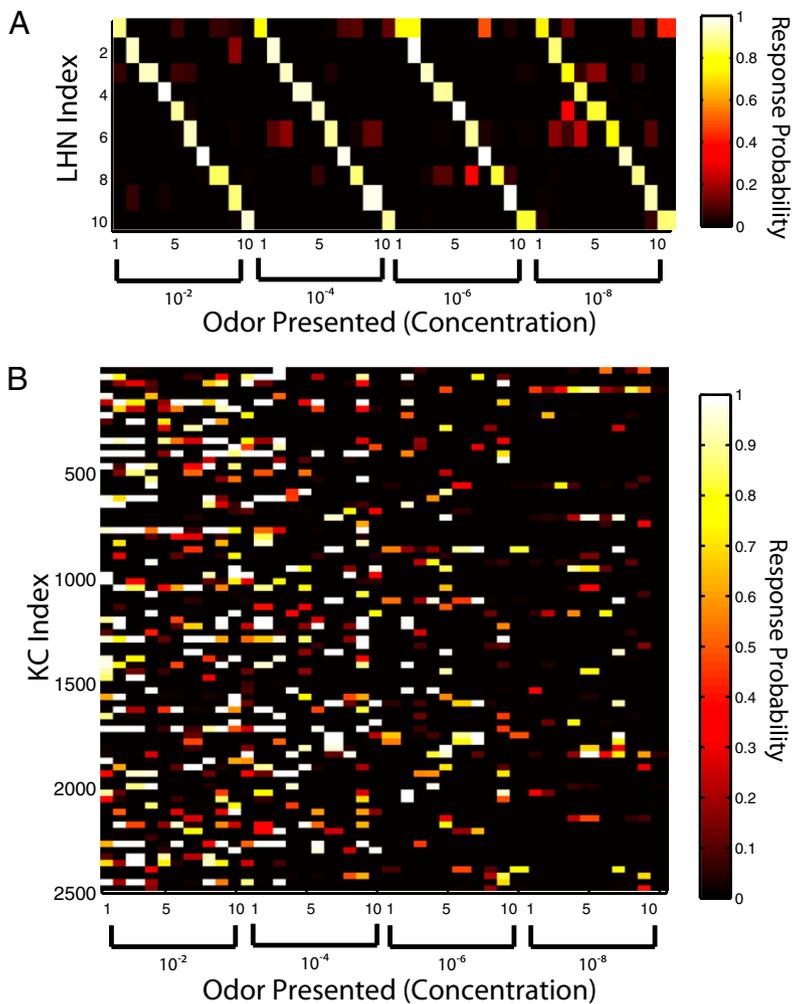


Fig. 4. LHN and KC responses for different concentrations. (A) Response probabilities for LHNs receiving input from five PNs and targeted to 10 different odors at four different concentrations. Four 10×10 matrices are shown, corresponding to the four concentrations, 10 trial odors, and LHNs targeted to each of the 10 odors. As in Figs. 2 and 3, each row shows a particular LHN, labeled by its target odorant, and each column gives the responses across LHNs to a specific odor. The bright diagonals repeated four times indicate that the LHNs respond reliably to their targets at different concentrations. Off-diagonal elements show false-positive responses. (B) Response probabilities of 2,500 KCs with five PN inputs each to 10 odorants at four concentrations. Four $10 \times 2,500$ blocks are shown, corresponding to the four concentrations, 10 odors, and 2,500 KCs. Otherwise, each block is as in Fig. 3A.

In summary, the most important factor in generating selective (LHNs) and sparse (KCs) third-order responses is the response magnitude equalization provided by the feedforward nonlinearity of the antennal lobe. The effects of lateral suppression, which are responsible for response decorrelation, are also important, although smaller than those of the feedforward nonlinearity, for retaining selectivity over a range of odor concentrations.

Discussion

A transformation in the representation of odors occurs in the antennal lobe that equalizes the magnitudes of and decorrelates the responses to different odors. Equalizing sensory response magnitudes might appear to be an unwise strategy because it diminishes a difference that can distinguish sensory stimuli. For this reason, discussions of sensory processing often stress the expansion of representations rather than the compression imposed by the compressive nonlinearity of the ORN-to-PN transformation. We have shown, however, that this compressive nonlinearity allows highly selective and sparse responses to be generated in higher-order neurons through the application of a fixed threshold. Response magnitude equalization is crucial for all of the results we have obtained: highly selective LHN responses, sparse KC responses, and concentration invariant LHN and KC responses. Given the importance of selectivity and sparseness in sensory representations, magnitude equalization may be a widespread coding strategy.

Decorrelation has also been postulated to be an important aspect of sensory processing (32–34). Decorrelation within the antennal lobe contributes less significantly to our results than does

response magnitude equalization. This is, in part, because our model LHNs and KCs can themselves remove correlations present on their inputs. The Fisher linear discriminant we used to set the specific connections from PNs to LHNs can perform this function, as can the global inhibition we introduced to model KCs. Global inhibition can decorrelate ORN responses because their largest variance principal component (Fig. 2B), like the global inhibition (*Methods*), is aligned with their odor-averaged response (20).

The normalization model (25) that describes the transformation from ORN to PN responses (Eq. 1) has been used to account for an impressive range of properties of visually responsive neurons (25, 26, 35, 36). It has been recently noted that normalization affects the population coding of orientation in the visual system in ways related to what we have found here for olfaction (37). The fact that the normalization model accounts for the effects of both feedforward nonlinearities and lateral suppression in olfactory responses supports the idea that it describes a fundamental computation in sensory processing.

The models we have generated to explore how olfactory representation may result in either innate or learned behaviors invoke anatomically distinct circuits for the two types of behavior. Innate behaviors, we postulate, result from determined connections from PNs to LHNs that elicit stereotyped responses in downstream circuits. This is in accord with the observation that a topographic map of olfactory information is retained in the lateral horn, but the character of the map differs from that of the antennal lobe (8, 9). In our model of the mushroom body, we propose a second transformation that invokes random, feedfor-

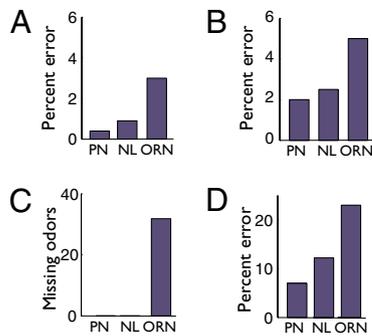


Fig. 5. Effect of feedforward nonlinearity and lateral suppression on LHN and KC responses. PN indicates that PN input was used with the full transformation of Eq. 1, NL indicates that no lateral suppression was included by setting $m = 0$ in Eq. 1, and ORN means that ORN rather than PN input was used, eliminating both components of the antennal lobe transformation. Percentage of error indicates the equal false-positive and false-negative error rates. (A) Percentage of errors for LHNs receiving 20 inputs and tuned as in Fig. 2. (B) Same as A, but for 5 inputs per LHN. (C) The number of odors to which 2,500 randomly tuned KCs (constructed as in Fig. 3) are unresponsive. (D) Percentage of errors for LHNs responding to odors at a variety of concentrations as in Fig. 4A.

ward, convergent connectivity between the glomeruli and the KCs. Neurons responsive to a given odor would thus be randomly distributed in the mushroom body (17). The valence of an odor and the ultimate behavioral output of a representation in the mushroom body would therefore be imposed by experience, rendering this a purely associative structure. If the mushroom body is indeed

an associative olfactory center, spatial order is conceptually superfluous. In contrast, innate behavioral responses mediated by the lateral horn are the consequence of stereotyped and determined inputs imposed by evolution rather than experience. Recent data suggest that a similar organizational logic may be operative in mammals, where information from the olfactory bulb bifurcates to project to limbic structures and the piriform cortex. Stereotyped connections to the amygdala may mediate innate olfactory responses and random inputs to the piriform cortex (38) may mediate more measured or learned olfactory behavior.

Methods

To avoid confounding receptors dedicated to pheromone detection with more general sensors, receptors Or33b, Or47b, Or65a, and Or88a were removed from the dataset of Hallem and Carlson (20) (SI Text). PN responses were generated from the ORN firing rates using Eq. 1 with $R_{\max} = 165$ Hz, $\sigma = 12$ Hz, and $m = 0.05$ Hz (22). Values of these parameters differ between PNs, so we have used values that provide the best overall fit. Noise was included by making the transformation $r_{\text{PN}} \rightarrow r_{\text{PN}} + \delta \tanh(\alpha r_{\text{PN}} \eta)$, where η is a random variable with zero mean and unit variance, $\delta = 10$ Hz, and $\alpha = 0.025$ Hz. We used Fisher's linear discriminant to set the weights from PNs to LHNs (SI Text).

For model KCs, we chose the n synaptic connections for the KCs randomly and drew their weights, denoted by the vector \mathbf{w} , from a uniform distribution between 0 and 1. Model KCs received a global inhibition constructed as described in SI Text.

ACKNOWLEDGMENTS. We thank R. Wilson, S. Olsen, E. Hallem, and J. Carlson for extremely valuable discussion concerning their data; L. Vosshall and S. Fusi for insights and encouragement; and R. Datta, V. Ruta, S. Ganguli, T. Toyozumi, Z. Pitkow, and M. Vidne for their help. This research was supported by a National Institutes of Health Director's Pioneer Award, part of the National Institutes of Health Roadmap for Medical Research, through Grant 5-DP1-OD114-02.

- Vosshall LB, Amrein H, Morozov PS, Rzhetsky A, Axel R (1999) A spatial map of olfactory receptor expression in the *Drosophila* antenna. *Cell* 96:725–736.
- Olsen SR, Bhandawat V, Wilson RI (2007) Excitatory interactions between olfactory processing channels in the *Drosophila* antennal lobe. *Neuron* 54:89–103.
- Shang Y, Claridge-Chang A, Sjulson L, Pypaert M, Miesenböck G (2007) Excitatory local circuits and their implications for olfactory processing in the fly antennal lobe. *Cell* 128:601–612.
- Root CM, et al. (2008) A presynaptic gain control mechanism fine-tunes olfactory behavior. *Neuron* 59:311–321.
- Olsen SR, Wilson RI (2008) Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature* 452:956–960.
- Tanaka NK, Awasaki T, Shimada T, Ito K (2004) Integration of chemosensory pathways in the *Drosophila* second-order olfactory centers. *Curr Biol* 14:449–457.
- Datta SR, et al. (2008) The *Drosophila* pheromone cVA activates a sexually dimorphic neural circuit. *Nature* 452:473–477.
- Marin EC, Jefferis GS, Komiyama T, Zhu H, Luo L (2002) Representation of the glomerular olfactory map in the *Drosophila* brain. *Cell* 109:243–255.
- Wong AM, Wang JW, Axel R (2002) Spatial representation of the glomerular map in the *Drosophila* protocerebrum. *Cell* 109:229–241.
- Jefferis GS, et al. (2007) Comprehensive maps of *Drosophila* higher olfactory centers: Spatially segregated fruit and pheromone representation. *Cell* 128:1187–1203.
- Zhang K, Guo JZ, Peng Y, Xi W, Guo A (2007) Dopamine-mushroom body circuit regulates saliency-based decision-making in *Drosophila*. *Science* 316:1901–1904.
- de Belle JS, Heisenberg M (1994) Associative odor learning in *Drosophila* abolished by chemical ablation of mushroom bodies. *Science* 263:692–695.
- Davis RL (2005) Olfactory memory formation in *Drosophila*: From molecular to systems neuroscience. *Annu Rev Neurosci* 28:275–302.
- Wang Y, et al. (2004) Stereotyped odor-evoked activity in the mushroom body of *Drosophila* revealed by green fluorescent protein-based Ca²⁺ imaging. *J Neurosci* 24:6507–6514.
- Turner GC, Bazhenov M, Laurent G (2008) Olfactory representations by *Drosophila* mushroom body neurons. *J Neurophysiol* 99:734–746.
- Masuda-Nakagawa LM, Tanaka NK, O'Kane CJ (2005) Stereotypic and random patterns of connectivity in the larval mushroom body calyx of *Drosophila*. *Proc Natl Acad Sci USA* 102:19027–19032.
- Murthy M, Fiete I, Laurent G (2008) Testing odor response stereotypy in the *Drosophila* mushroom body. *Neuron* 59:1009–1023.
- Lin HH, Lai JS, Chin AL, Chen YC, Chiang AS (2007) A map of olfactory representation in the *Drosophila* mushroom body. *Cell* 128:1205–1217.
- Hallem EA, Ho MG, Carlson JR (2004) The molecular basis of odor coding in the *Drosophila* antenna. *Cell* 117:965–979.
- Hallem EA, Carlson JR (2006) Coding of odors by a receptor repertoire. *Cell* 125:143–160.
- Bhandawat V, Olsen SR, Gouwens NW, Schlieff ML, Wilson RI (2007) Sensory processing in the *Drosophila* antennal lobe increases reliability and separability of ensemble odor representations. *Nat Neurosci* 10:1474–1482.
- Olsen SR, Bhandawat V, Wilson RI (2010) Divisive normalization in olfactory population codes. *Neuron* 66:287–299.
- Huerta R, Nowotny T, García-Sánchez M, Abarbanel HD, Rabinovich MI (2004) Learning classification in the olfactory system of insects. *Neural Comput* 16:1601–1640.
- Kazama H, Wilson RI (2008) Homeostatic matching and nonlinear amplification at identified central synapses. *Neuron* 58:401–413.
- Heeger DJ (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9:181–197.
- Carandini M, Heeger DJ, Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17:8621–8644.
- Abbott LF, Luo SX (2007) A step toward optimal coding in olfaction. *Nat Neurosci* 10:1342–1343.
- Itskov V, Abbott LF (2008) Pattern capacity of a perceptron for sparse discrimination. *Phys Rev Lett* 101:018101.
- Crittenden JR, Skoulakis EM, Han KA, Kalderon D, Davis RL (1998) Tripartite mushroom body architecture revealed by antigenic markers. *Learn Mem* 5:38–51.
- Tanaka NK, Tanimoto H, Ito K (2008) Neuronal assemblies of the *Drosophila* mushroom body. *J Comp Neurol* 508:711–755.
- Assisi C, Stopfer M, Laurent G, Bazhenov M (2007) Adaptive regulation of sparseness by feedforward inhibition. *Nat Neurosci* 10:1176–1184.
- Barlow HB (1961) Possible principles underlying the transformation of sensory messages. *Sensory Communication*, ed Rosenblith WA (MIT Press, Cambridge, MA).
- Uttley AM (1979) *Information Transmission in the Nervous System* (Academic, London).
- Atick JJ, Redlich AN (1990) Towards a theory of early visual processing. *Neural Comput* 2:308–320.
- Simoncelli EP, Schwartz O (1999) Modeling non-specific suppression in V1 neurons with a statistically-derived normalization model. *Advances in Neural Information Processing Systems*, eds Kearns MS, Solla SA, Cohn DA (MIT Press, Cambridge, MA), Vol 11, pp 153–159.
- Reynolds JH, Heeger DJ (2009) The normalization model of attention. *Neuron* 61:168–185.
- Ringach DF (2010) Population coding under normalization. *Vision Res*, 10.1016/j.visres.2009.12.007.
- Stettler DD, Axel R (2009) Representations of odor in the piriform cortex. *Neuron* 63:854–864.