

Sensory processing in the *Drosophila* antennal lobe increases reliability and separability of ensemble odor representations

Vikas Bhandawat¹, Shawn R Olsen^{1,2}, Nathan W Gouwens^{1,2}, Michelle L Schliefl^{1,2} & Rachel I Wilson¹

Here we describe several fundamental principles of olfactory processing in the *Drosophila melanogaster* antennal lobe (the analog of the vertebrate olfactory bulb), through the systematic analysis of input and output spike trains of seven identified glomeruli. Repeated presentations of the same odor elicit more reproducible responses in second-order projection neurons (PNs) than in their presynaptic olfactory receptor neurons (ORNs). PN responses rise and accommodate rapidly, emphasizing odor onset. Furthermore, weak ORN inputs are amplified in the PN layer but strong inputs are not. This nonlinear transformation broadens PN tuning and produces more uniform distances between odor representations in PN coding space. In addition, portions of the odor response profile of a PN are not systematically related to their direct ORN inputs, which probably indicates the presence of lateral connections between glomeruli. Finally, we show that a linear discriminator classifies odors more accurately using PN spike trains than using an equivalent number of ORN spike trains.

Each glomerulus in the olfactory system receives synaptic input from many ORNs, all of which express the same odorant receptor gene. Each second-order neuron sends a dendrite into a single glomerulus, so for each odorant receptor gene there is an identifiable ORN type and a corresponding type of second-order neuron. An odorant typically activates multiple ORN types, and so each odor is represented as a population code across different glomerular processing channels^{1,2}. What happens to olfactory signals as they move through these channels? It is technically challenging to address this question in vertebrates because there are so many glomeruli. In *Drosophila*, the problem is comparatively simpler because the antennal lobe contains only ~50 glomeruli. Each of these glomeruli has a stereotyped position that is identifiable across flies, and almost all have been matched to an identified ORN type^{3–6}. For these reasons, it may be easier to discover the basic principles of early olfactory processing in this model organism.

In general, effective information transmission requires that the response evoked by a stimulus should be highly reliable, and that the responses evoked by different stimuli should be distinctive. Therefore, we have focused on two fundamental questions. First, how reproducible is the number of spikes evoked by repeated presentations of the same odor? There has been remarkably little attention paid to the reproducibility of olfactory responses, and the small number of previous studies on this issue have been concerned with the precision of spike timing rather than the reproducibility of spike counts^{7,8}. Response reproducibility is a central issue in sensory processing because the signal-to-noise ratio of a neural response limits the rate of information transmission by that neuron.

The second fundamental question concerns the distinctiveness of neural responses to different stimuli. How selective are ORNs, and how does their selectivity compare with that of second-order olfactory neurons? Three studies published more than 20 years ago in vertebrates reached conflicting conclusions on this issue, but it was not feasible for these investigators to directly compare the selectivity of pre- and postsynaptic neurons corresponding to the same glomerulus^{9–11}. More recently, three studies made this direct comparison in the *Drosophila* antennal lobe, but again the results were conflicting^{12–14}. Two of these studies used genetically encoded sensors, which may not report spike trains faithfully owing to their limited dynamic range^{15,16}; the third study recorded spike trains directly, but examined only one glomerulus.

Here we aim to resolve these issues with a systematic analysis of the inputs and outputs of seven glomeruli in the *Drosophila* antennal lobe (**Supplementary Fig. 1** online). Our results show that there is a major transformation of olfactory representations in this region of the brain. The most important effects of this transformation are to improve the signal-to-noise ratio of individual spike trains and to distribute odor representations more uniformly in neuronal coding space.

RESULTS

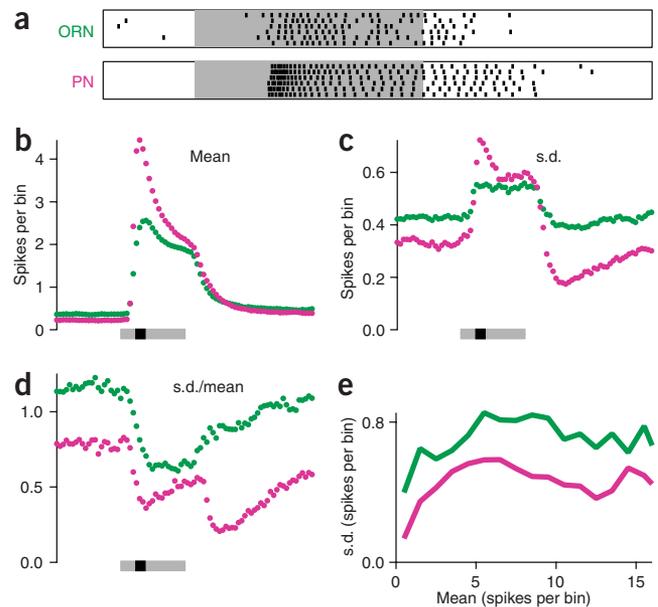
Odor responses are more reliable in PNs than in ORNs

The variability of a neuronal response can be quantified by assessing the variability in the number of spikes evoked by a sensory stimulus. In most sensory systems, the spike-count variability of stimulus-evoked responses increases at each successive level of processing in a sensory

¹Department of Neurobiology, Harvard Medical School, 220 Longwood Avenue, Boston, Massachusetts 02115, USA. ²These authors contributed equally to this work. Correspondence should be addressed to R.I.W. (rachel_wilson@hms.harvard.edu).

Received 16 July; accepted 16 August; published online 7 October 2007; doi:10.1038/nn1976

Figure 1 Odor responses are more reliable in PNs than in ORNs. (a) Odor responses of an ORN and PN pre- and postsynaptic to the glomerulus (glomerulus VA2, odor is geranyl acetate). Each tick represents a spike, and each row in a raster represents a different trial. The gray bar indicates a 500-ms odor stimulus period. (b) Mean odor responses are larger in PNs (magenta) than in ORNs (green). Spikes were counted in 50-ms bins and averaged across five trials with the same odor, then averaged across all blocks of trials (all odors and all experiments). The gray bar indicates the stimulus period; the black bar indicates a 100-ms period when average PN firing rates are maximal. (c) Standard deviations (s.d.) of spike counts in five trials with the same odor, averaged across all blocks of trials (all odors and all experiments). (d) Coefficient of variation (s.d./mean) of spike counts in five trials with the same odor, averaged across all blocks of trials. Note that the coefficient of variation of PN responses drops again after odor offset. This is because some responses contain zero spikes for an epoch following odor offset, so the s.d. in these bins is zero for some responses. (e) The average s.d. of spike counts is lower for PNs than for ORNs even when mean firing rates are matched. s.d. values were measured for all counting windows in all blocks of trials, binned according to mean firing rate and averaged across all counting windows in the same bin. Note that because the s.d. deviation depends sublinearly on the mean, the average coefficient of variation is larger than (the average s.d.)/(the average mean).



system^{17–19}. However, in the *Drosophila* antennal lobe, ~40 ORNs with the same receptive field converge onto ~4 PNs in each glomerulus^{20–22}. Thus, by pooling across these inputs, PNs might be able to reduce their response variability. We therefore compared the reliability of odor-evoked spike counts in ORNs and PNs.

We presented an odor stimulus in multiple consecutive trials to the same cell (a ‘block’ of trials; **Fig. 1a**). To quantify the spike-count reliability across trials, we divided each set of repeated responses into 50-ms windows that overlapped by 25 ms. In each time window, we computed the mean and the standard deviation of the spike count across repeated responses by the same cell to the same odor. Odors typically evoked more vigorous responses in PNs than in ORNs (**Fig. 1a,b** and **Supplementary Fig. 2** online). So, although the typical standard deviation of PN responses is slightly greater than that of ORN responses (**Fig. 1c**), PN responses are less variable in proportion to the magnitude of the response ($P < 10^{-15}$, whether comparing over the entire stimulus period or the 100-ms epoch at the response peak, Mann-Whitney U -test, $n = 779$ ORN responses and 843 PN responses; **Fig. 1d**). Thus, individual PNs are more reliable than individual ORNs, which should tend to make their responses more informative.

We also compared the standard deviations of ORN and PN spike counts as a function of the mean spike count for each time window. For all mean spike counts, PNs have a lower standard deviation than ORNs (**Fig. 1e**). Furthermore, the standard deviation is not strongly dependent on the mean, and so stronger responses have a lower coefficient of variation. Because PN responses are on average stronger than ORN responses (**Fig. 1b**), this also tends to make PNs more reliable than ORNs.

PNs preferentially transmit the rising phase of ORN signals

ORN responses typically do not peak until 100–300 ms after odor onset²¹. This is probably because spiking is coupled to odorant receptor activation by the generation of second messengers. However, odors can trigger rapid behavioral responses in flies, with a total latency from stimulus to motor reaction of less than 300 ms²³. This suggests that neurons in the brain are preferentially tuned to detect the rising phase of ORN signals, rather than the response peak. This motivated us to compare the onset kinetics of odor responses in synaptically connected ORNs and PNs.

Comparing peri-stimulus time histograms averaged across all odor responses in all cells, we noted that PN responses rise more rapidly than ORN responses (**Fig. 2a**). Furthermore, PN responses begin to decay while ORN responses are still growing. This is also clear in most direct comparisons between synaptically connected ORNs and PNs (**Figs. 1a** and **2b,c**). Overall, PN responses peak significantly faster than ORN responses ($P < 10^{-7}$, paired t -test, $n = 69$ odor-glomerulus combinations; **Fig. 2d**), and the time to half-decay of the response is shorter for PNs than for their presynaptic ORNs ($P < 10^{-5}$, paired t -test; **Fig. 2e**). Taken together, faster rise and faster decay mean that a more excitatory drive to third-order neurons occurs within an early epoch of the odor response ($P < 10^{-11}$, paired t -test; **Fig. 2f**). Therefore, PNs act as high-pass filters that preferentially signal the rising phase of the ORN response.

Because PN responses accommodate rapidly, we chose to quantify response magnitudes in PNs by measuring the average firing rates during an early epoch of the response (a 100-ms time window beginning 100 ms after odor onset; **Fig. 1b**). Because a fruitfly can respond rapidly after encountering an odor, this early epoch should be particularly informative to downstream neurons. Throughout this study, we also quantified PN responses in a different way: following other investigators^{21,24}, we measured the average spike rates over the entire 500-ms stimulus period. The main conclusions from this study are the same for both of these response metrics.

ORNs and PNs differ in odor selectivity and odor preferences

Setting aside the issues of trial-to-trial reliability and response kinetics, we examined the average response magnitudes for each cell type to our odor stimuli (**Fig. 3a**). How does the response profile of each PN type compare with the response profile of its corresponding ORNs? We began by asking simply whether these responses are linearly correlated. For each glomerulus we found a statistically significant correlation between the ORN and PN response profile ($P < 0.05$ for all seven glomerular comparisons, Pearson’s correlation; **Supplementary Table 1** online), but r^2 values are only in the range of 0.26–0.81. This means that a linear scaling of ORN responses explains only 26–81% of the odor-dependent variance in PN responses.

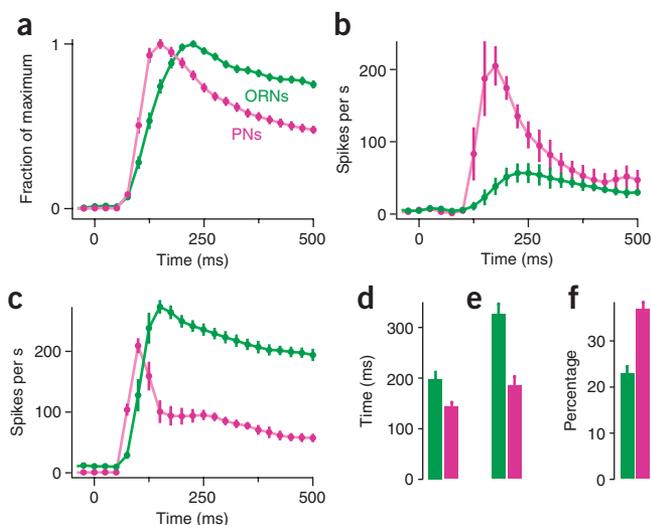


Figure 2 PNs preferentially transmit the rising phase of ORN signals. (a) Average peak-normalized peri-stimulus time histograms (PSTHs), averaged across all odors and all glomeruli (\pm s.e.m.). Note that PN responses rise and decay more rapidly than ORN responses. Odor stimulation begins at 0 ms and ends at 500 ms. (b) An example comparing the responses of pre- and postsynaptic neurons to the same odor. PSTHs show the average response of ORNs and PNs in glomerulus VA2 to geranyl acetate (mean \pm s.e.m., averaged across experiments). Note that the PN response is robust at a time point when the ORNs have just begun to respond, and the PN response begins decaying before the ORNs have peaked. (c) Another example of PSTHs for ORNs and PNs in glomerulus DM1 showing responses to ethyl butyrate. The PN response rises faster and peaks earlier, even though in this case the PN peak is smaller. (d) Compared with ORN responses, PN responses have a shorter latency to reach 90% of the response peak (mean \pm s.e.m., across all blocks of trials; see **Supplementary Methods**). (e) PN responses have a faster decay from peak to half-peak. (f) A larger percentage of the total spike count occurs in the first 200 ms after odor onset for PN responses compared with ORN responses.

Two features of PN odor responses diminish this linear correlation. First, for each glomerulus, PNs are less selective than their presynaptic ORNs ($P < 0.05$, Wilcoxon signed-rank test; **Fig. 3b**, similar results in **Supplementary Fig. 3** online). To test the generality of this result, we also compared ORN and PN selectivity for glomerulus DM4 at three odor concentrations. As with our standard concentration (1:1,000 dilution), weaker stimuli (1:10,000 and 1:100,000 dilutions) produce PN response profiles that were less selective than the corresponding ORN response profiles (**Fig. 4**; see also **Supplementary Fig. 4** online). Other investigators who used identical stimulus conditions have shown that ORN responses are very sparse at the 1:100,000 dilution, indicating that this concentration is near the bottom of the dynamic range of this system^{21,22,24,25}. These results show that broad PN tuning is a phenomenon that is not limited to high odor concentrations.

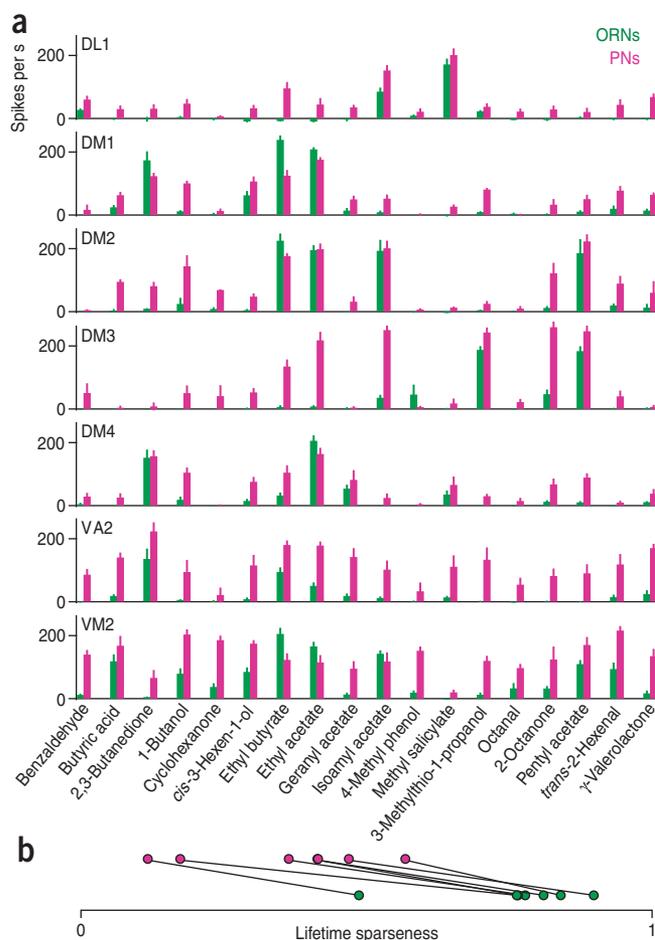
Another factor that diminishes this linear correlation is that the rank order of odor preferences differs for ORNs and PNs. For example, whereas ethyl butyrate is the 3rd-ranked odor of DL1 PNs, it is only ranked 16th among the odor responses of DL1 ORNs (**Fig. 3a**). Some of this difference is due to errors in estimating each average response profile on the basis of a limited sample of individual experiments. However, sampling error cannot completely account for this result. This can be shown by pairing an individual ORN with a corresponding individual PN and computing the correlation between their odor ranks, and then comparing the distribution of these correlations with the correlations obtained from ORN-ORN or PN-PN pairings. Because we were not able to test every odor in every experiment, we assembled many simulated response profiles by drawing randomly from a normal distribution defined by the mean and standard deviation of each average response profile (**Fig. 5a**; see also **Supplementary Methods**

Figure 3 ORNs and PNs differ in their odor selectivity. (a) Response profiles of synaptically connected ORNs (green) and PNs (magenta) for seven glomeruli. Bars show averages across all experiments (\pm s.e.m.; see **Supplementary Table 2** for n). Responses are measured as the mean spike rate during the 100-ms epoch when firing rates are peaking (black bar in **Fig. 1b–d**), minus the baseline firing rate. Results are similar over the entire 500-ms stimulus period (**Supplementary Fig. 3**). (b) The selectivity of each response profile is quantified as lifetime sparseness^{29,48} (see **Supplementary Methods**; 0 = nonselective, 1 = maximally selective). ORNs and PNs that correspond to the same glomeruli are connected. PNs are consistently less selective than their corresponding ORNs. The highest ORN sparseness value is for glomerulus DL1 and the lowest is for glomerulus VM2.

online). The median correlation between ORN and PN ranks was only 0.47, which is substantially lower than the correlation between ORNs of the same type or between PNs of the same type (0.65 and 0.61, respectively; **Fig. 5**). The simplest explanation for this result is that the odor preferences of a PN are influenced by lateral connections between glomeruli^{26,27}.

A nonlinear transformation function for each glomerulus

So, the output of a glomerulus is not a simple linear scaling of its inputs. Furthermore, because ORNs and PNs differ in their ranked



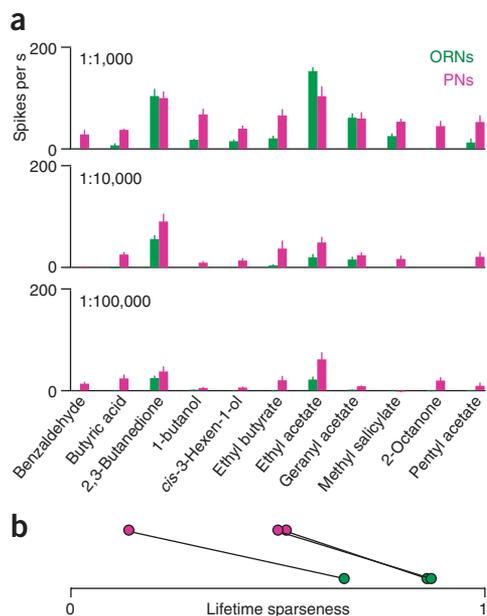


Figure 4 ORNs and PNs differ in their odor selectivity even at low stimulus intensities. **(a)** Response profiles for DM4 ORNs and PNs to a panel of 11 odors at three concentrations. Bars show averages across all experiments (\pm s.e.m.; see **Supplementary Table 3** for n). Because the response peak tends to occur later for more dilute stimuli, we measured responses as the mean spike rate during the entire 500-ms stimulus period, minus the baseline firing rate (as in **Supplementary Fig. 3**). **(b)** The selectivity of each response profile for the three odor dilutions. ORNs and PNs that correspond to the same dilution are connected. Note that DM4 PNs are consistently less selective than DM4 ORNs at all three concentrations. (Selectivity at the 1:1,000 dilution is slightly different from the selectivity value plotted in **Supplementary Figure 3** for this glomerulus because here we used only a subset of our 18 test odors.)

odor preferences, no monotonic function can describe the relationship between the ORN and PN response profile for a glomerulus. We therefore asked whether there is any systematic relationship between ORN and PN responses. For each glomerulus, we plotted PN responses to each odor as a function of ORN responses to the same odor (**Fig. 6a**; see also **Supplementary Fig. 5** online). This revealed a consistent transformation function for each glomerulus, albeit with some scatter. These functions have a similar shape for most glomeruli: they initially slope steeply, meaning that the gain of the transformation function is high for weak inputs. As ORN input levels increase these curves flatten, meaning that the gain of the transformation function decreases. (This is true for all the glomeruli we tested apart from DL1.) Therefore, these plots show that PNs inherit much of their tuning from their presynaptic ORNs, but the transformation is nonlinear.

This type of transformation function may be useful because it makes better use of the available response range of a PN. This is illustrated by projecting the points in each plot onto both the x and y axes (**Fig. 6a**). ORNs do not use all parts of their dynamic range with equal frequency in response to our stimuli. However, two odors that elicit similarly

weak activity in an ORN can elicit different levels of activity in a postsynaptic PN, because each glomerular transformation function shows high gain at low ORN input levels. This tends to distribute the responses to these stimuli more uniformly throughout the response range of each PN.

This type of sensory transformation has been termed ‘histogram equalization’²⁸ because it produces a flatter histogram of response intensities. To examine whether this is the case across the entire population of cells in our data set, we plotted the distribution of response intensities for ORNs and PNs, accumulated across all odors and all glomeruli (**Fig. 6b**). The ORN response histogram shows a large peak at low response intensities, and a long, flat tail covering the rest of the ORN dynamic range. The PN response histogram, by contrast, is much flatter, indicating that all available response intensities are used with more uniform frequency. In this sense, PNs encode our odor stimuli more efficiently than ORNs do.

Although there is a clear overall relationship between ORN and PN responses for each glomerulus, it is also important to note that these functions do not predict PN odor responses completely. Because ORNs and PNs differ in their ranked odor preferences, no monotonic function will account for all these data.

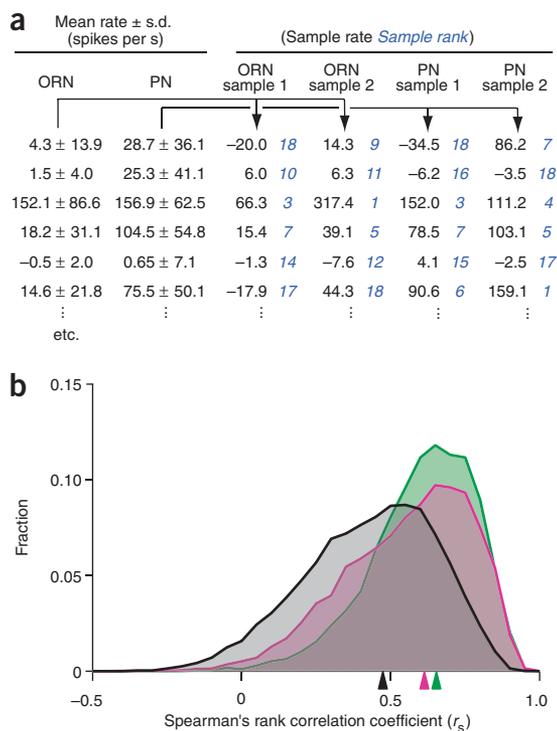


Figure 5 The rank order of ORN and PN odor preferences is different. **(a)** An example illustrating how we computed correlations between the odor ranks of individual cell response profiles. Here we show the mean and standard deviation of the ORN and PN response profiles for glomerulus DM4. (Note the table is truncated after six odors.) We drew randomly from these distributions to produce representative simulated profiles for two individual ORNs and two individual PNs (arrows). Next we ranked the odors in each individual response profile (blue). In this example, the correlation coefficient between the 18 odor ranks of ORN sample 1 and ORN sample 2 (r_s) is 0.79. Correlation coefficients are lower for PN-PN comparisons (0.58 in this example). In comparison with each of these, ORN-PN correlations are much lower (0.33, 0.39, 0.42 and 0.46 in this particular example). **(b)** Histograms showing the distribution of Spearman's rank correlation coefficients (r_s), accumulated across 2,000 runs of the simulation procedure for each glomerulus. Arrowheads indicate the median of each distribution. The ORN-PN correlations (gray) do not lie between the ORN-ORN (green) and PN-PN (magenta) correlations, indicating that ORN and PN odor ranks are not drawn from the same underlying mean distribution.

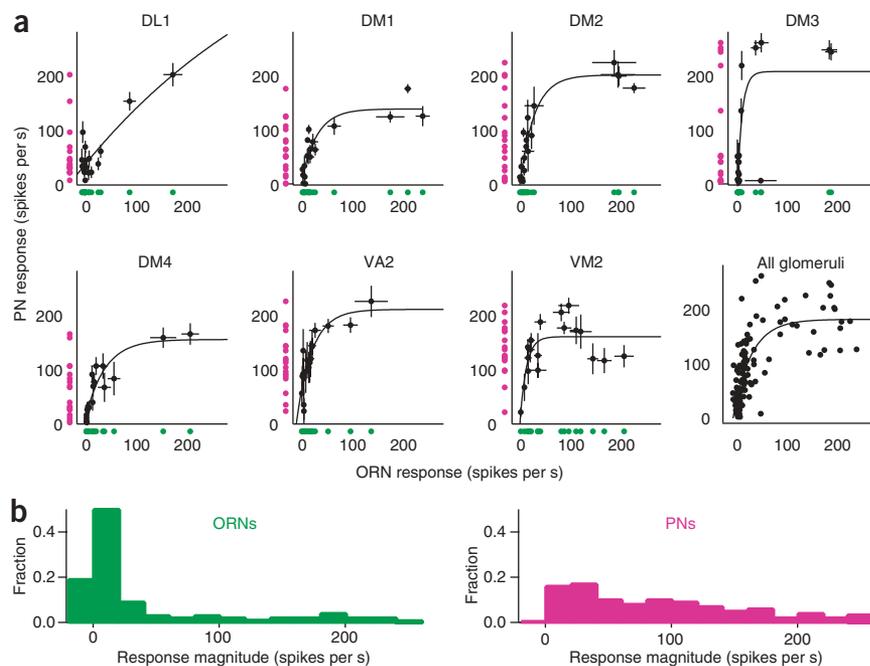


Figure 6 PN odor responses are partly explained by a highly nonlinear transformation of their direct ORN inputs. **(a)** For each glomerulus, the average PN response to an odor is plotted against the average ORN response to that odor (black symbols, \pm s.e.m.). Curves are exponential fits ($y = y_0 + A \cdot e^{kx}$). Green and magenta symbols are projections of the data onto the x and y axes, showing that odor responses generally occupy the dynamic range of a PN more evenly than they occupy that of an ORN. Responses are measured as the mean spike rate during the 100-ms epoch when firing rates are peaking (with no baseline subtraction), but results are similar if responses are measured as the mean spike rate during the entire 500-ms stimulus period (**Supplementary Fig. 5**). **(b)** Histograms of ORN and PN response magnitudes. Each histogram is accumulated across all 126 response magnitudes ($= 7$ glomeruli $\times 18$ odors). The PN histogram is flatter than the ORN histogram, indicating that PNs use their dynamic range more efficiently.

glomeruli. We achieved this by independently shuffling the odor labels on each glomerular response profile. We then measured inter-odor distances in seven-dimensional space

for all possible pairwise combinations of odors. When we repeated this simulation many times, the range of distances we obtained was indistinguishable from the distances we measured in our real data set (**Fig. 7c,d**). This means the separation between ensemble odor representations is roughly what we would predict, based solely on histogram equalization in each glomerulus individually.

Correlations between cell types and odors

We have seen that PNs use all parts of their dynamic range with approximately equal frequency, and in this sense encode odors more efficiently than ORNs do²⁸. However, the term ‘efficient coding’ has also been applied to the idea that the responses of different neurons should be maximally independent from each other³³. We measured the independence of different glomerular coding channels by computing the percentage of the variance in the ensemble odor responses that is captured by each of the seven principal components of the seven-dimensional ORN or PN coding space. If all seven cell types were completely correlated, then the first principal component would account for 100% of the variance in the data. In other words, all the data would lie along a single line in multidimensional space. Conversely, if all cell types were perfectly decorrelated, and if the data were drawn from a multidimensional Gaussian distribution, then each principal component would account for an equal amount of the total variance ($100\% \div 7 = \sim 14\%$). (Even in this case, we would need a very large odor set to discern this perfect decorrelation.)

The principal components of our ORN data set fall between these hypothetical extremes (**Fig. 8a**). In part, this reflects the limited size of our odor set and the non-Gaussian distribution of the ORN response histograms (**Fig. 6b**). We demonstrated this by independently and randomly shuffling the odor labels on each of the seven ORN response profiles and re-computing the principal components of this simulated data set. These simulations always produced a first principal component that accounted for a disproportionately large share of the variance (usually 30–40%; **Fig. 8a**). However, the principal components of the real (non-shuffled) data set are even more skewed, with the first principal component accounting for 54% of the variance. This

Odor representations in multiglomerular coding space

Third-order neurons receive convergent input from multiple PN types^{29–32}. It is therefore important to examine odor representations in multiglomerular coding space. In the simplest case, histogram equalization in one dimension should also produce a more uniform distribution of odors in multiple dimensions. To visualize odor representations in seven-dimensional space, we reduced the dimensionality of this space by performing principal components analysis. The first two principal components define the two-dimensional projection that maximizes the variance of the data. In this projection, most odors are still clustered near the origin of the ORN space, with only a few odors located far from this cluster (**Fig. 7a**). In the equivalent PN space, odors fill the available coding space more uniformly (**Fig. 7b**).

Thus, as a result of the ORN-to-PN transformation, odor representations are distributed more efficiently in multiglomerular coding space. In concrete terms, the ensemble patterns of spiking activity elicited by any two odors become more different. We quantified this by measuring Euclidean distances between odors in seven-dimensional space for all possible pairwise combinations of odors ($[18 \text{ choose } 2] = 153$ pairs). During the early epoch of odor responses, distances are significantly larger in PN space than in ORN space ($P < 0.0001$, paired t -test, $n = 153$). Moreover, the distribution of distances is narrower for PNs than for ORNs over the entire stimulus period (note the differing interquartile ranges in **Fig. 7c,d**). This means that odors are distributed more uniformly in PN coding space. Some odor distances decrease, but others increase.

Is the separation of odors in multidimensional space larger or smaller than we would predict, based solely on the independent odor separation in each one-dimensional coding channel? The answer depends on the degree of correlation between the different glomeruli. Lateral connections shape PN odor responses^{26,27} (**Fig. 5**); if these connections increase correlations between different PN types, this would decrease inter-odor distances in multidimensional space. To address this issue, we constructed a simulated data set that preserves the distribution of response magnitudes for each glomerular cell type, but breaks any dependencies between odor responses in different

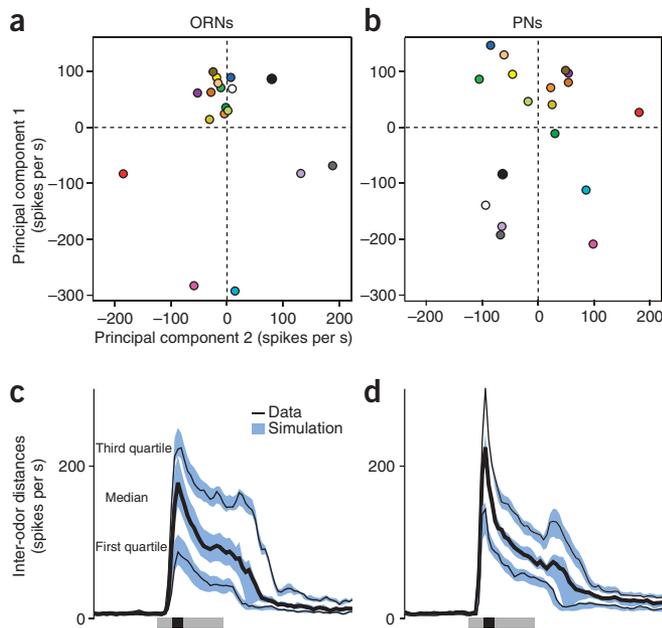


Figure 7 Odors are distributed more uniformly in ensemble PN coding space than in ensemble ORN coding space. **(a)** Average odor responses from seven ORN types projected onto the space defined by the first two principal components. Each point represents a different odor. **(b)** Same as **a** for PN data (with the same color conventions), showing a more uniform separation between odor representations. **(c)** The difference between ensemble ORN responses to different odors is quantified as the Euclidean distance between odor representations in seven-dimensional space. Distances are computed for all 153 pairwise combinations of the 18 odor stimuli, and the median and interquartile range of this distribution are plotted here for each time point. The interquartile range is wide because some odors are well separated in ORN space, but many are poorly separated. Blue bands indicate the range of results obtained by shuffling odor labels on each glomerular response profile (see **Supplementary Methods**). The gray bar indicates the 500-ms stimulus period and the black bar indicates the 100-ms period when firing rates were measured for **a** and **b**. **(d)** Same as **c** for PN responses. At the peak of the response (black bar), distances are significantly larger in PN space compared with ORN space. PN responses then quickly accommodate (**Fig. 2**), and so inter-odor distances shrink. However, the interquartile range of distances remains smaller than in ORN space. This indicates a more uniform distribution of distances. As in **c**, shuffling odor labels on each glomerular response profile produces a range of results (blue bands) that resembles the real data.

means that ORNs in different glomeruli have odor preferences that are more highly correlated than we would expect, based solely on the distribution of response magnitudes in each glomerular channel.

Similarly, about half the variance in the PN data is captured by the first principal component (51%; **Fig. 8b**). This is mainly due to the limited size of our odor set and the non-Gaussian distribution of the PN response histograms, as shuffling the odor labels on each PN response profile always produced a skewed distribution of principal component contributions (**Fig. 8b**). Because real PN data produced a distribution that was even more skewed than the simulated data, PNs (like ORNs) are more correlated than we would expect, based solely on the distribution of response magnitudes for each PN type.

In summary, sensory processing in the *Drosophila* antennal lobe does not change the degree of independence between different glomerular coding channels. The conclusions of this analysis are similar regardless of whether we measure spike rates around the response peak (**Fig. 8**) or over the entire stimulus period (**Supplementary Fig. 6** online).

PN responses are more linearly separable than ORN responses

Increased PN reliability and more uniform odor distances in PN coding space should mean that odors are more discriminable on the basis of PN spike trains than on the basis of an equivalent number of ORN spike trains. We tested this prediction by measuring the ability of an algorithm to identify the odor stimulus on the basis of the ensemble neural response elicited by that odor. Because our data come from single (not multiple) unit ORN and PN recordings, we simulated 'multi-unit' responses by assembling data from different glomerular classes. Each simulated data set consisted of 90 multi-unit responses (18 odors with 5 spike trains per odor per cell). We performed linear discriminant analysis to identify the linear combinations of input variables that best separated all 18 odor response clusters from each other. To evaluate the quality of these discriminations, we withheld 1 multi-unit odor response from the data set, trained the algorithm with the remaining 89, and predicted the odor corresponding to the one withheld response. The predicted odor was then compared

with the actual odor. We repeated this analysis with many independently assembled multi-unit responses at each time point in the odor response.

Before odor onset, the prediction success rate hovers near chance (**Fig. 9a**). (The success rate is slightly above chance because different cells have different spontaneous firing rates, and spontaneous firing rates sometimes drift during experiments; thus, spontaneous firing rates were slightly 'predictive' of the odor because successive trials with an odor were presented consecutively rather than interleaved). After odor stimulus onset, success rates rise rapidly. As expected, including more glomerular classes in the data set produced higher success rates

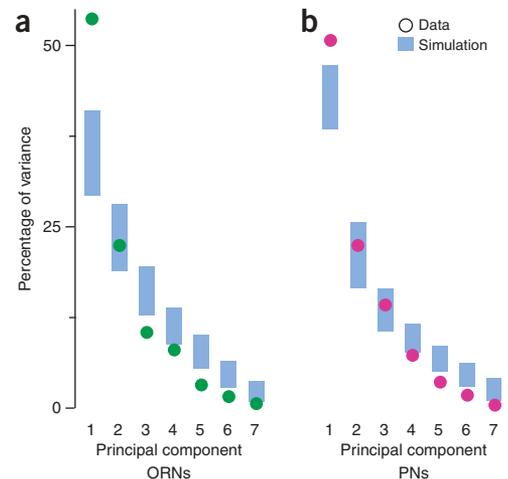


Figure 8 Correlations between different glomeruli are similar for ORNs and PNs. **(a)** Principal components analysis (PCA) was applied to the 18×7 ORN response matrix. The magnitude of the variance accounted for by each principal component (green circles) is a measure of the correlations between different ORN types. Blue bands indicate the range of results obtained by shuffling odor labels on each glomerular response profile (see **Supplementary Methods**). Comparing the data and the simulation shows that ORNs are less independent in their odor responses than we would expect, based solely on the distribution of response magnitudes within each glomerular coding channel. **(b)** Same as **a** for the 18×7 PN response matrix. Correlations between PN types are similar to correlations between ORN types.

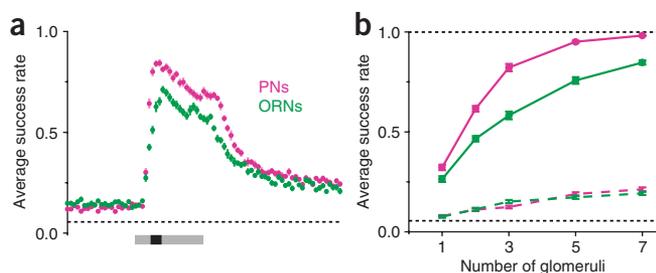


Figure 9 A linear discriminator can classify odors more accurately with responses from multiple PNs than with responses from the same number of ORNs. **(a)** Odor classification success rate from linear discriminator analysis with data sets that include cells from three glomerular classes. All possible combinations of three glomeruli were sampled. Points are the mean \pm s.e.m., averaged across 20 runs of the classification procedure. The dotted line represents chance performance. **(b)** Success rate is higher for PN data than for ORN data, regardless of how many glomerular classes are included in the data set. Points are the mean \pm s.e.m., averaged over the 100-ms window shown in **a**, and then averaged across 20 runs of the classification procedure. Dashed green and magenta lines plot the classification success rate during the baseline period before odor onset; this is an artifact of varying spontaneous activity rates (see text), and ORN and PN performance is similar. Dotted black lines indicate perfect and chance performance.

(Fig. 9b). Because success rates using PN data plateau at 100% for some classifications that use data from more than three glomeruli, this procedure underestimates the difference between ORN and PN responses. Nevertheless, success rates were significantly higher for PN data than for ORN data for all conditions in Figure 9b ($P < 0.005$, Mann-Whitney U -tests, $n = 20$ runs of the classification procedure for ORNs and PNs for each condition). This demonstrates that a linear discriminator can classify odors more accurately with responses from several PNs than with responses from the same number of ORNs.

DISCUSSION

An improved signal-to-noise ratio

Studies in other systems have implied that the variability of stimulus-evoked spike counts almost always increases at successively higher levels of sensory processing³⁴. For example, the visual responses of higher cortical neurons are often very noisy¹⁷, in contrast to the reliability of retinal ganglion cells³⁵. A direct comparison of the responses evoked by identical stimuli in the retina, thalamus and visual cortex has confirmed that spike-count reliability decreases at each successively higher level of the visual stream¹⁸. This is despite the fact that a simple cell in primary visual cortex pools signals from ~ 30 thalamic neurons³⁶, which should improve its reliability. Similarly, a direct comparison of spike trains at successive levels of an insect auditory circuit has found that noise increases at successively higher levels¹⁹. Our results show a different trend: spike counts in individual PNs are more consistent than spike counts in individual ORNs. This is partly because PNs tend to fire more vigorously than their presynaptic ORNs in response to the same stimulus, and stronger responses are more reliable for both ORNs and PNs. This may imply an increasingly deterministic control of spike timing at high firing rates owing to intrinsic refractoriness³⁷. However, even at the same firing rates, PN responses are more reliable than ORN responses. This may reflect the benefits of pooling: each PN is postsynaptic to many ORNs, and all these ORNs respond in a similar way to odors^{21,22,38}. If noise is uncorrelated across ORNs, then pooling these inputs should improve the reliability of PN responses.

On balance, the improvement in reliability is smaller than one might predict. Each glomerulus corresponds to ~ 40 ORNs and ~ 4 PNs; this

means the average PN pool inputs from 10–40 ORNs (depending on whether each ORN contacts all PNs in a glomerulus). Because pooling N ORN inputs should decrease the variability of the pooled average by \sqrt{N} , we would expect the coefficient of variation to improve by $\sqrt{10}$ to $\sqrt{40}$. The effect we describe is on the low end of this range, suggesting that each ORN contacts only a single PN, or that PNs receive additional noise from other neuronal sources.

High-pass filtering of olfactory signals

Our results show that PNs can be extremely sensitive to small differences between weak ORN inputs. Even a small increase in ORN spike rate above the baseline can produce a robust response in postsynaptic PNs. As a result, PN responses rise rapidly even when ORN responses build slowly. This is particularly useful because the onset kinetics of ORNs are intrinsically limited by the speed of the signal transduction cascades that link odorant receptor activation to spike initiation. PN responses then rapidly decline while ORN spike rates continue to rise. This means that PNs act as high-pass filters, transmitting the rising phase of ORN responses preferentially over the tonic component of ORN responses. This rapid accommodation might be due to any of several mechanisms, including short-term synaptic depression at the ORN-to-PN synapse.

Taken together, a faster rise and a faster decay should sharpen the estimate of odor arrival time by downstream neurons. For a fly in flight, this should translate to an improved estimate of odor plume location. Notably, *Drosophila* can turn in flight less than 300 ms after encountering an odor plume²³. A similar phenomenon operates in the visual system: sluggish photoreceptor responses trigger speedy depolarizations in downstream neurons³⁹ and ultimately rapid behavioral responses to visual stimuli.

We note that *Drosophila* PN responses differ from the responses of locust PNs, which typically show more complex temporal patterning^{29,40,41}. Locust PNs also show a higher average level of maintained activity throughout the odor response (relative to the response peak) and often show excitatory responses to odor offset⁴¹. By contrast, *Drosophila* PNs accommodate rapidly and typically do not burst after stimulus offset (but for some exceptions see Supplementary Fig. 2f).

A nonlinear transformation increases coding efficiency

An important finding from this study is that although PNs inherit a substantial portion of their odor tuning from their presynaptic ORNs, this relationship is nonlinear. This nonlinearity disproportionately amplifies small differences between weak ORN inputs. By contrast, small differences between strong ORN inputs are not amplified to the same degree. Most ORN odor responses cluster in the weak end of the dynamic range of the ORN. As a result of this nonlinear transformation, PNs use their dynamic range more uniformly than ORNs do. If all portions of the dynamic range of a neuron are used with equal frequency, the carrying capacity of that information channel is maximized because the entropy of the neuron's response is maximized. This tends to protect signals from contamination by noise added at later stages in the processing channel⁴². This has long been recognized as a useful computation in sensory processing²⁸. If broader tuning curves are useful, why has evolution not simply produced broadly tuned ORNs? ORN responses are directly linked to the way odorant receptor proteins interact with odor molecules; therefore, broadening ORN tuning might require changing the biophysics of odorant receptors in ways that are unfavorable for other reasons.

Broad PN tuning may seem counterintuitive: we tend to think of higher-order neurons as being more selective than their presynaptic inputs. These expectations are founded in part on the paradigm of

visual processing, in which successive layers of higher-order neurons are increasingly specialized to represent complex conjunctions of visual features. However, there is also a huge expansion in the number of higher-order neurons devoted to representing complex visual features compared with the number of retinal ganglion cells. Thus the dimensionality of higher visual representations is increasingly large, so these brain regions can afford to code information as ensembles of narrowly tuned neurons. By contrast, in early olfactory processing, the dimensionality of the second-order representation is the same as the dimensionality of the first-order representation. Therefore, total coding space cannot increase (unless time is used as another coding dimension⁴³).

In a truly efficient coding scheme, neurons should efficiently encode natural stimulus distributions, not arbitrary stimulus distributions²⁸. Although our odor set is chemically diverse and relatively large, it may not be representative of the odors a wild fly would encounter. In the future, it would be interesting to learn whether the principles of olfactory processing we describe here also apply to a more naturalistic distribution of odor stimuli.

Another caveat is that we have not sampled all antennal lobe glomeruli. However, because most of the glomeruli in our data set showed a similar nonlinear transformation function, our conclusions probably generalize to many glomeruli outside our data set. An interesting special case is the glomerulus DA1, which was not included in our data set. ORNs projecting to this glomerulus respond weakly but selectively to a *Drosophila* pheromone. Their postsynaptic PNs are also selective for this odor, but respond much more robustly⁴⁴. Thus, this processing channel shows amplification without a change in response selectivity.

Ensemble odor responses

Extending our analysis from one glomerular channel to multiglomerular ensembles, we found that odors are more uniformly separated in ensemble PN coding space than in ensemble ORN coding space. This separation is approximately what we would expect, based on the increased separation between odors within each glomerular coding channel. We demonstrated this using a simulation that made the odor preferences of each glomerulus independent from each other, while preserving the characteristic response magnitudes of each cell type.

We also found that sensory processing in the antennal lobe does not substantially alter the independence of different glomerular coding channels. The degree of correlation between different PN types is similar to the degree of correlation between different ORN types. Notably, ORNs have more correlated odor preferences than we would expect based solely on their tuning breadth and the size of our odor set. We also obtained the same result by re-analyzing a large published data set²⁴ comprising 24 *Drosophila* ORN types and 110 odors (result not shown). This result may be due to the common evolutionary origin of different *Drosophila* odorant receptors in gene duplication events⁴⁵. In addition, some odors are intrinsically more volatile than others, which will tend to produce similarities in the odor preferences of different glomeruli. Like ORNs, PNs are also more correlated than we would expect, based on their tuning breadth and the size of our odor set. This may reflect correlations that are inherited from ORNs.

The role of lateral connections between glomeruli

Glomeruli in the *Drosophila* antennal lobe are connected by GABAergic interneurons^{20,46} and cholinergic interneurons^{26,27}. What is the role of these connections in the transformations we have described?

We have noted that the rank order of PN odor preferences is different from the order of ORN odor preferences. We show that this difference is too large to be explained by the uncertainty in our estimates of each

average odor response profile (Fig. 5). Therefore, some of this difference is probably caused by lateral interglomerular connections, because lateral inputs would have an odor tuning that reflects the odor preferences of ORNs that are presynaptic to other glomeruli. In principle, either inhibitory or excitatory lateral connections could cause this phenomenon. The computational significance of this phenomenon is not clear, as it does not seem to decorrelate the responses of PNs in different glomeruli (Fig. 8).

It is easy to see how lateral connections could cause scatter around each glomerular transformation function. However, lateral connections may also play an important part in determining the underlying shape of these transformation functions. For example, all of these functions have a y intercept above 0 (Fig. 6). This reflects the tendency of PNs to respond weakly to an odor even when their presynaptic ORNs are not responding at all. Lateral excitatory connections are strong enough to trigger these responses^{26,27}. Moreover, lateral inhibition could act on ORN axon terminals to govern the probability of neurotransmitter release, and thereby contribute to a nonlinear relationship between pre- and postsynaptic activity. At the analogous synapse in the olfactory bulb, ORN-to-mitral cell synapses show strong frequency-dependent short-term plasticity that is modulated by presynaptic inhibition through GABAergic local neurons⁴⁷. Thus, we should not assume that the systematic relationships in Figure 6 are intrinsic to each glomerulus. More mechanistic experiments will be required to disentangle the role of intra- versus interglomerular mechanisms in shaping these transformation functions.

Odor discrimination

We have shown that a linear discriminator can identify odors more accurately on the basis of PN spike trains than on the basis of an equivalent number of ORN spike trains. This is probably due to both the increased distances between odors in PN space and the improved signal-to-noise ratio among PNs. It is important to point out that linear discriminant analysis is not meant to emulate a biologically plausible downstream neuron, and that real third-order neurons will be subject to more constraints than our algorithm is. Also, this is not an optimal decoder, and so its performance may not reflect the total amount of information in the responses it decodes. Finally, the total amount of information in the entire PN ensemble cannot, of course, exceed the total amount of information in the entire ORN ensemble. What we have shown here is that the information in a limited subset of the PN ensemble is more useful to a linear decoder than the information in an equivalent number of ORNs. This highlights the potential functional consequence of increased PN reliability, combined with increased inter-odor distances in PN coding space.

In conclusion, we have described two fundamental tasks that are accomplished by the first stage of the olfactory processing stream. On a neuron-to-neuron basis, our comparisons show that signal reproducibility is increased and distinctions between the responses to different stimuli are enhanced. The details of odor processing in the vertebrate olfactory bulb might be different, especially because the number of glomeruli in vertebrates is much larger. Nevertheless, most organisms share a common olfactory processing architecture, which suggests that some of the basic principles we have demonstrated in flies may also apply to vertebrates.

METHODS

Fly stocks. Flies were reared at room temperature on conventional cornmeal agar. All experiments were carried out on adult female flies 2–7 d after eclosion. **Supplementary Table 4** online lists genotypes for all experiments. See **Supplementary Methods** for stock origins.

ORN recordings. Flies were immobilized in the trimmed end of a plastic pipette tip under a $\times 50$ air objective mounted on an Olympus BX51WI microscope. A reference electrode filled with saline was inserted into the eye, and a sharp saline-filled glass capillary (tip diameter $< 1 \mu\text{m}$) was inserted into a sensillum. Recordings were obtained with an A-M Systems Model 2400 amplifier, low-pass filtered at 2 kHz and digitized at 10 kHz. ORN spikes were detected using routines in IgorPro (Wavemetrics). See **Supplementary Methods** for details.

PN recordings. Whole-cell recordings from PN somata were carried out *in vivo* as previously described⁴⁶. One neuron was recorded per brain and the morphology of each cell was visualized *post hoc* with biocytin-streptavidin and nc82 histochemistry as described previously⁴⁶, except that in the secondary incubation we used 1:250 goat anti-mouse-AlexaFluor633 and 1:1,000 streptavidin-AlexaFluor568 (Molecular Probes). See **Supplementary Methods** for details.

Olfactory stimulation. We chose a panel of 18 odors to maximize the chemical diversity of our stimuli, and to maximize overlap with odors used in other studies of the same ORNs^{22,24}. For all experiments (except in **Fig. 4**), odors were diluted 1:100 v/v in paraffin oil (J.T. Baker, VWR no. JTS894), except 3-methylthio-1-propanol, which was diluted 1:100 v/v in water, and 4-methyl phenol, which was diluted 1:100 w/v in water. In **Figure 4**, odors were diluted 1:100, 1:1,000, or 1:10,000 in paraffin oil. See **Supplementary Methods** for details.

Data analysis. See **Supplementary Methods** for data analysis details.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank K. Ito, L. Luo, L.B. Vosshall and L.M. Stevens for gifts of fly stocks. We benefited from helpful conversations with S.A. Baccus, V. Jayaraman, H. Kazama, A.W. Liu, J.H.R. Maunsell, O. Mazor, M. Meister, R.C. Reid, H. Sompolinsky, G.C. Turner and Y. Zhou. This work was funded by a grant from the US National Institutes of Health (1R01DC008174-01), a Pew Scholar Award, a McKnight Scholar Award, a Smith Family Foundation New Investigators Award and an Armenise-Harvard Junior Faculty Award (to R.I.W.). S.R.O. is supported by a US National Science Foundation Predoctoral Fellowship. N.W.G. is supported by a Howard Hughes Medical Institute Predoctoral Fellowship. M.L.S. is supported by a National Institutes of Health Postdoctoral Fellowship (1F32DC008741-01A1).

Published online at <http://www.nature.com/natureneuroscience>

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions>

- Bargmann, C.I. Comparative chemosensation from receptors to ecology. *Nature* **444**, 295–301 (2006).
- Mombaerts, P. Genes and ligands for odorant, vomeronasal and taste receptors. *Nat. Rev. Neurosci.* **5**, 263–278 (2004).
- Lai, S.S., P.P. et al. Three-dimensional reconstruction of the antennal lobe in *Drosophila melanogaster*. *J. Comp. Neurol.* **405**, 543–552 (1999).
- Hallem, E.A. & Carlson, J.R. The odor coding system of *Drosophila*. *Trends Genet.* **20**, 453–459 (2004).
- Couto, A., Alenius, M. & Dickson, B.J. Molecular, anatomical, and functional organization of the *Drosophila* olfactory system. *Curr. Biol.* **15**, 1535–1547 (2005).
- Fishilevich, E. & Vosshall, L.B. Genetic and functional subdivision of the *Drosophila* antennal lobe. *Curr. Biol.* **15**, 1548–1553 (2005).
- Stopfer, M. & Laurent, G. Short-term memory in olfactory network dynamics. *Nature* **402**, 664–668 (1999).
- Bazhenov, M., Stopfer, M., Sejnowski, T.J. & Laurent, G. Fast odor learning improves reliability of odor responses in the locust antennal lobe. *Neuron* **46**, 483–492 (2005).
- Derby, C.D. & Ache, B.W. Quality coding of a complex odorant in an invertebrate. *J. Neurophysiol.* **51**, 906–924 (1984).
- Mathews, D.F. Response patterns of single neurons in the tortoise olfactory epithelium and olfactory bulb. *J. Gen. Physiol.* **60**, 166–180 (1972).
- Duchamp, A. Electrophysiological responses of olfactory bulb interneurons to odor stimuli in the frog. A comparison with receptor cells. *Chem. Senses* **7**, 191–210 (1982).
- Ng, M. et al. Transmission of olfactory information between three populations of neurons in the antennal lobe of the fly. *Neuron* **36**, 463–474 (2002).
- Wang, J.W., Wong, A.M., Flores, J., Vosshall, L.B. & Axel, R. Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* **112**, 271–282 (2003).
- Wilson, R.I., Turner, G.C. & Laurent, G. Transformation of olfactory representations in the *Drosophila* antennal lobe. *Science* **303**, 366–370 (2004).
- Pologruto, T.A., Yasuda, R. & Svoboda, K. Monitoring neural activity and $[\text{Ca}^{2+}]$ with genetically encoded Ca^{2+} indicators. *J. Neurosci.* **24**, 9572–9579 (2004).
- Sankaranarayanan, S. & Ryan, T.A. Real-time measurements of vesicle-SNARE recycling in synapses of the central nervous system. *Nat. Cell Biol.* **2**, 197–204 (2000).
- Shadlen, M.N. & Newsome, W.T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
- Kara, P., Reinagel, P. & Reid, R.C. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* **27**, 635–646 (2000).
- Vogel, A., Hennig, R.M. & Ronacher, B. Increase of neuronal response variability at higher processing levels as revealed by simultaneous recordings. *J. Neurophysiol.* **93**, 3548–3559 (2005).
- Stocker, R.F., Heimbeck, G., Gendre, N. & de Belle, J.S. Neuroblast ablation in *Drosophila* PIGAL4 lines reveals origins of olfactory interneurons. *J. Neurobiol.* **32**, 443–456 (1997).
- de Bruyne, M., Clyne, P.J. & Carlson, J.R. Odor coding in a model olfactory organ: the *Drosophila* maxillary palp. *J. Neurosci.* **19**, 4520–4532 (1999).
- de Bruyne, M., Foster, K. & Carlson, J.R. Odor coding in the *Drosophila* antenna. *Neuron* **30**, 537–552 (2001).
- Budick, S.A. & Dickinson, M.H. Free-flight responses of *Drosophila melanogaster* to attractive odors. *J. Exp. Biol.* **209**, 3001–3017 (2006).
- Hallem, E.A. & Carlson, J.R. Coding of odors by a receptor repertoire. *Cell* **125**, 143–160 (2006).
- Hallem, E.A., Ho, M.G. & Carlson, J.R. The molecular basis of odor coding in the *Drosophila* antenna. *Cell* **117**, 965–979 (2004).
- Olsen, S.R., Bhandawat, V. & Wilson, R.I. Excitatory interactions between olfactory processing channels in the *Drosophila* antennal lobe. *Neuron* **54**, 89–103 (2007).
- Shang, Y., Claridge-Chang, A., Sjulson, L., Pypaert, M. & Miesenböck, G. Excitatory local circuits and their implications for olfactory processing in the fly antennal lobe. *Cell* **128**, 601–612 (2007).
- Laughlin, S. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch. [C]* **36**, 910–912 (1981).
- Perez-Orive, J. et al. Oscillations and sparsening of odor representations in the mushroom body. *Science* **297**, 359–365 (2002).
- Marin, E.C., Jefferis, G.S., Komiyama, T., Zhu, H. & Luo, L. Representation of the glomerular olfactory map in the *Drosophila* brain. *Cell* **109**, 243–255 (2002).
- Wong, A.M., Wang, J.W. & Axel, R. Spatial representation of the glomerular map in the *Drosophila* protocerebrum. *Cell* **109**, 229–241 (2002).
- Tanaka, N.K., Awasaki, T., Shimada, T. & Ito, K. Integration of chemosensory pathways in the *Drosophila* second-order olfactory centers. *Curr. Biol.* **14**, 449–457 (2004).
- Simoncelli, E.P. Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.* **13**, 144–149 (2003).
- Shadlen, M.N. & Newsome, W.T. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* **4**, 569–579 (1994).
- Berry, M.J., Warland, D.K. & Meister, M. The structure and precision of retinal spike trains. *Proc. Natl. Acad. Sci. USA* **94**, 5411–5416 (1997).
- Alonso, J.M., Usrey, W.M. & Reid, R.C. Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *J. Neurosci.* **21**, 4002–4015 (2001).
- Berry, M.J. II & Meister, M. Refractoriness and neural precision. *J. Neurosci.* **18**, 2200–2211 (1998).
- Shanbhag, S.R., Muller, B. & Steinbrecht, R.A. Atlas of olfactory organs of *Drosophila melanogaster*. 1. Types, external organization, innervation, and distribution of olfactory sensilla. *Int. J. Insect Morphol. Embryol.* **28**, 377–397 (1999).
- Armstrong-Gold, C.E. & Rieke, F. Bandpass filtering at the rod to second-order cell synapse in salamander (*Ambystoma tigrinum*) retina. *J. Neurosci.* **23**, 3796–3806 (2003).
- Stopfer, M., Jayaraman, V. & Laurent, G. Intensity versus identity coding in an olfactory system. *Neuron* **39**, 991–1004 (2003).
- Mazor, O. & Laurent, G. Transient dynamics vs. fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**, 661–673 (2005).
- Laughlin, S.B., Howard, J. & Blakeslee, B. Synaptic limitations to contrast coding in the retina of the blowfly *Calliphora*. *Proc. R. Soc. Lond. B* **231**, 437–467 (1987).
- Laurent, G. Olfactory network dynamics and the coding of multidimensional signals. *Nat. Rev. Neurosci.* **3**, 884–895 (2002).
- Schlieff, M.L. & Wilson, R.I. Olfactory processing and behavior downstream from highly selective receptor neurons. *Nat. Neurosci.* **10**, 623–630 (2007).
- Nozawa, M. & Nei, M. Evolutionary dynamics of olfactory receptor genes in *Drosophila* species. *Proc. Natl. Acad. Sci. USA* **104**, 7122–7127 (2007).
- Wilson, R.I. & Laurent, G. Role of GABAergic inhibition in shaping odor-evoked spatio-temporal patterns in the *Drosophila* antennal lobe. *J. Neurosci.* **25**, 9069–9079 (2005).
- Wachowiak, M. et al. Inhibition of olfactory receptor neuron input to olfactory bulb glomeruli mediated by suppression of presynaptic calcium influx. *J. Neurophysiol.* **94**, 2700–2712 (2005).
- Vinje, W.E. & Gallant, J.L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).