

Supplemental Experimental Procedures

ORN and LFP Recordings

Flies were immobilized in the trimmed end of a plastic pipette tip. An antenna or maxillary palp was stabilized using glass capillaries or coverslips positioned using micromanipulators. The recording electrode was a sharp saline-filled glass capillary was inserted into a sensillum for ORN recordings, or inserted into the third antennal segment for LFP recordings. A sharp saline-filled glass capillary in the eye served as a reference electrode. Signals were recorded on an A-M Systems Model 2400 amplifier with a 10-M Ω headstage, low-pass filtered at 2 kHz and digitized at 10 kHz. Recorded ORNs were identified based on sensillum morphology and size, sensillum position on the antenna or palp, spike amplitude, spontaneous spike frequency, and odor tuning of all cells in a sensillum. These properties form an unambiguous signature of ORN identity. To sort spikes from the DM1 ORN, it was necessary to kill another ORN housed in the same sensillum using diphtheria toxin light chain (DTI), so these recordings were made in flies with the genotype *UAS-DTI/CyO;Or92a-Gal4*. Other ORN recordings were performed in the following genotypes: VM7, *NP3481,UASCD8GFP*; DL5, *NP3062,UASCD8GFP*; DM4, *NP3062,UASCD8GFP*. ORN recordings in Figure S1A were performed in *w¹¹¹⁸* flies. For this figure, we recorded from a total of 223 ORNs, sampled randomly from all morphological classes (30 antennal large basoconic, 106 antennal small basoconic, 24 palp basoconic, 30 trichoid, 28 coeloconic). Of these 223 ORNs, 42 had spike waveforms which could not be sorted from the waveform of another ORN housed in the same sensillum; in these cases, the summed response was divided in half and each half assigned to one ORN. None of the unsortable sensilla showed a strong response to any private stimuli (range = -13.6 to 10.6 spikes/sec, interquartile range -3.8 to 0.2 spikes/sec).

For LFP measurements, we selected recording sites with the knowledge that the contribution of each ORN type to the LFP may depend on the recording site. We used an odor (ethyl acetate 10⁻⁸) that activates a single ORN type (DM1 ORNs) situated on the proximal end of the antenna to understand how the signal decays as a function of distance. We found that the signal decays to about 50% of the peak value as the recording site moved distally along the long axis of the third segment, with little decay of the signal along the short axis. Therefore, in order to weight proximal and distal ORN types roughly equally, we measured the LFP at one proximal site and one distal site, and then averaged these measurements together.

PN Recordings

The composition of the internal patch-pipette solution was (in mM): potassium aspartate 140, HEPES 10, MgATP 4, Na₃GTP 0.5, EGTA 1, KCl 1, biocytin hydrazide 13 (pH = 7.3, osmolarity adjusted to ~ 265 mOsm). The composition of the external saline was (in mM): NaCl 103, KCl 3, *N*-tris(hydroxymethyl) methyl-2-aminoethane-sulfonic acid 5, trehalose 8, glucose 10, NaHCO₃ 26, NaH₂PO₄ 1, CaCl₂ 1.5, and MgCl₂ 4. Osmolarity was adjusted to 270–275 mOsm. The saline was bubbled with 95% O₂/5% CO₂ and reached a final pH = 7.3. An Olympus BX51WI microscope with a 40x water-immersion objective and IR-DIC optics was used to obtain recordings under visual control. Recordings were targeted to specific PNs based on GFP visualization. We targeted 4 PN types in this study: DM1 in *NP5221-Gal4,UASCD8GFP*, DM4 and DL5 in *NP3062-Gal4,CD8GFP*, and VM7 in *NP3481-Gal4,UASCD8GFP*. Recordings were obtained with an A-M Systems Model 2400 amplifier in current-clamp mode (10-M Ω headstage), low-pass filtered at 5 kHz, and digitized at 10 kHz.

GABA receptor antagonists were prepared as concentrated stock solutions and a measured volume of stock was added to the saline perfusing the brain to achieve final concentrations of 5 μ M (picrotoxin, Sigma) and 10 μ M (CGP54626, Tocris). Occasionally these drugs caused the PN membrane potential to oscillate at ~2-4 Hz; when this occurred the recording was terminated.

Olfactory Stimuli

Test stimuli for VM7 PNs (Figure 4A,B) were the following monomolecular odors at 1:100 dilution: methyl salicylate, benzaldehyde, linalool, 1-octen-3-ol, 1-butanol, trans-2-hexenal, 2-heptanone, pentyl acetate, butyric acid, ethyl acetate, ethyl butyrate. We also tested 6 blend stimuli on VM7 PNs. Three blend stimuli were made by mixing methyl salicylate 10⁻², fenchone 10⁻⁴, methyl acetate 10⁻⁶, trans-2-hexenal 10⁻⁶, propionic acid 10⁻⁵ with three concentrations of butanone (10⁻⁴ to 10⁻⁶). Three more blend stimuli were generated by adding

pentyl acetate 10^{-3} to each of the above blends. Test stimuli for DL5 PNs (Figure 4C,D) were the following monomolecular odors at 1:100 dilution: acetophenone, butanal, benzaldehyde, butyric acid, 1-butanol, cadaverine, ethyl butyrate, ethanol, hexanal, linalool, methyl salicylate, 3-methylthio-1-propanol, 1-octen-3-ol, paraffin oil, valeric acid, pentyl acetate, 1-penten-3-ol.

Reported odor concentrations represent v/v dilutions in solvent and are 1:100 unless otherwise noted (always paraffin oil, J.T. Baker, VWR #JTS894, except for 3-methyl-thio-1-propanol, which was diluted 1:100 v/v in water, and ethanol, which was diluted 1:100 w/v in water). Odor stimulation was performed as described previously (Bhandawat et al., 2007) with modifications described here. As in previous studies, a constant stream of charcoal-filtered air (2.2 L/min) was directed at the fly throughout each experiment, and when triggered by a command pulse, a three-way solenoid valve redirected a portion (0.20 L/min) through the headspace of the odor vial for 500 msec. Thus all odors were diluted an additional 10-fold in air just before reaching the fly. In this study, in order to deliver odor blends, we added a second solenoid to direct a portion of the carrier stream (0.20 L/min) through a second odor vial. The two solenoids were triggered simultaneously, and they drew from the carrier stream at the same point. The two odorized streams both rejoined the carrier stream 15 cm from the end of the end of the delivery tube, which measured 3 mm in diameter and was positioned 8 mm from the fly. We made fresh dilutions in solvent every five days, except odors with high vapor pressures (methyl acetate and ethyl acetate) which were replaced every three days. These two odors were also made up in a larger volume (5 mL vs 2 mL), so as to deplete a smaller fraction of the odor molecules in the headspace on each trial.

Data Analysis and Modeling: alternative models of gain control

In addition to evaluating the input gain control and response gain control models (Figure 3), we also evaluated three other models of gain control. First, we evaluated two subtractive models of inhibition. In one, we modeled a rightward shift in the input-output function:

$$PN = R_{\max} \left(\frac{(ORN - s)^{1.5}}{(ORN - s)^{1.5} + \sigma^{1.5}} \right)$$

This produced a poor fit (mean squared error = 846 and 149 spikes²/sec² for VM7 and DL5, compare to Figure 3G). The second subtractive model was a downward shift in the input-output function:

$$PN = R_{\max} \left(\frac{ORN^{1.5}}{ORN^{1.5} + \sigma^{1.5}} \right) - s$$

This also produced a poor fit (mean squared error = 245 and 80 spikes²/sec² for VM7 and DL5).

Finally, we evaluated a model where we held R_{\max} and σ constant across all concentrations of pentyl acetate and instead we allowed the exponent in Equation (1) to vary across pentyl acetate concentrations. This produced the poorest fit to the data.

Data Analysis and Modeling: an alternative model of odor discrimination

In an additional analysis, we allowed perceptrons to have both positive and negative weights, and we allowed each perceptron to have a different threshold. To determine the weights and the thresholds, we used either a support vector machine with a linear kernel or else, in a separate analysis, Fisher's linear discriminant analysis. Performance rates were always higher for these perceptrons than for sign-constrained perceptrons. We do not describe these results in detail because (1) these perceptrons are less biologically realistic than sign-constrained perceptrons, and because (2) their higher performance rates leave less dynamic range for seeing the effects of the transformations we describe. However, we confirmed that our main conclusions are still true for these networks: the input gain model always performs better than the model without lateral inhibition, and it also outperforms the response gain model. The only difference was that the effects of these transformations were smaller, and the distributions of errors in the confusion matrices were different.

It should be noted that the performance of perceptrons trained and tested with the response gain model can be improved by changing the overall level of inhibition or changing the form of the relationship between s and total ORN activity. We did not explore these alternatives systematically because our goal was simply to compare the versions of these two gain models that were the best fits to our data.

Table S1: Candidate “private” odor stimuli.

	2a	7a	9a	10a	19a	22a	23a	33b	35a	42a	42b	43a	43b	47a	47b	49b	59b	65a	67a	67c	82a	85a	85b	85f	88a	98a
trans-2-hexenal 10^{-8}	6	23	-3	-1	-3	-1	0	2	4	°	°	1	2	0	2	-1	-5	3	1	4	-8	1	-8	-2	8	-10
trans-2-hexenal 10^{-6}	4	204	0	3	-5	10	-4	6	-7	°	°	6	3	-3	5	-2	1	1	0	-6	-5	2	-5	-1	5	-2
ethyl acetate 10^{-8}	4	-5	-4	6	14	-3	-4	1	-1	°	56	7	3	10	1	0	2	0	0	-4	-5	-1	-4	3	-2	-5
ethyl acetate 10^{-6}	-1	14	1	-2	4	2	0	0	-3	°	111	6	3	3	0	0	0	0	7	3	-7	5	2	-3	6	2
methyl acetate 10^{-2}	5	-6	24	9	-13	54	2	15	-6	°	°	-12	22	25	-21	-5	269	-3	7	7	15	-10	29	12	14	4
2-butanone 10^{-2}	8	9	94	7	-16	84	3	7	13	157	°	-8	25	16	-15	3	14	-2	67	21	17	1	39	41	-14	4

To identify “private” stimuli which activate just one ORN type, we first examined published data describing the spiking responses of many ORNs to a large panel of odor stimuli (Hallem & Carlson, 2006). Based on this data (reproduced in this table), it appeared that trans-2-hexenal, ethyl acetate, and methyl acetate were good candidates for “private” odors. In addition, our pilot experiments showed that 2-butanone strongly activated Or42a-expressing ORNs (VM7 ORNs) at low concentrations, suggesting this odor might also be a good candidate (see table).

Odorant receptors are listed in the top row of the table. Entries in the table represent mean firing rates, measured in units of spikes/sec over the 500-msec odor stimulus period, and averaged over several measurements in different flies. All entries in the table are from Hallem & Carlson (2006), except the three values in bold italics, which were measured by us. An open symbol (°) indicates that no data is available. Strong responses indicating candidate receptor-odor pairs are highlighted in color.

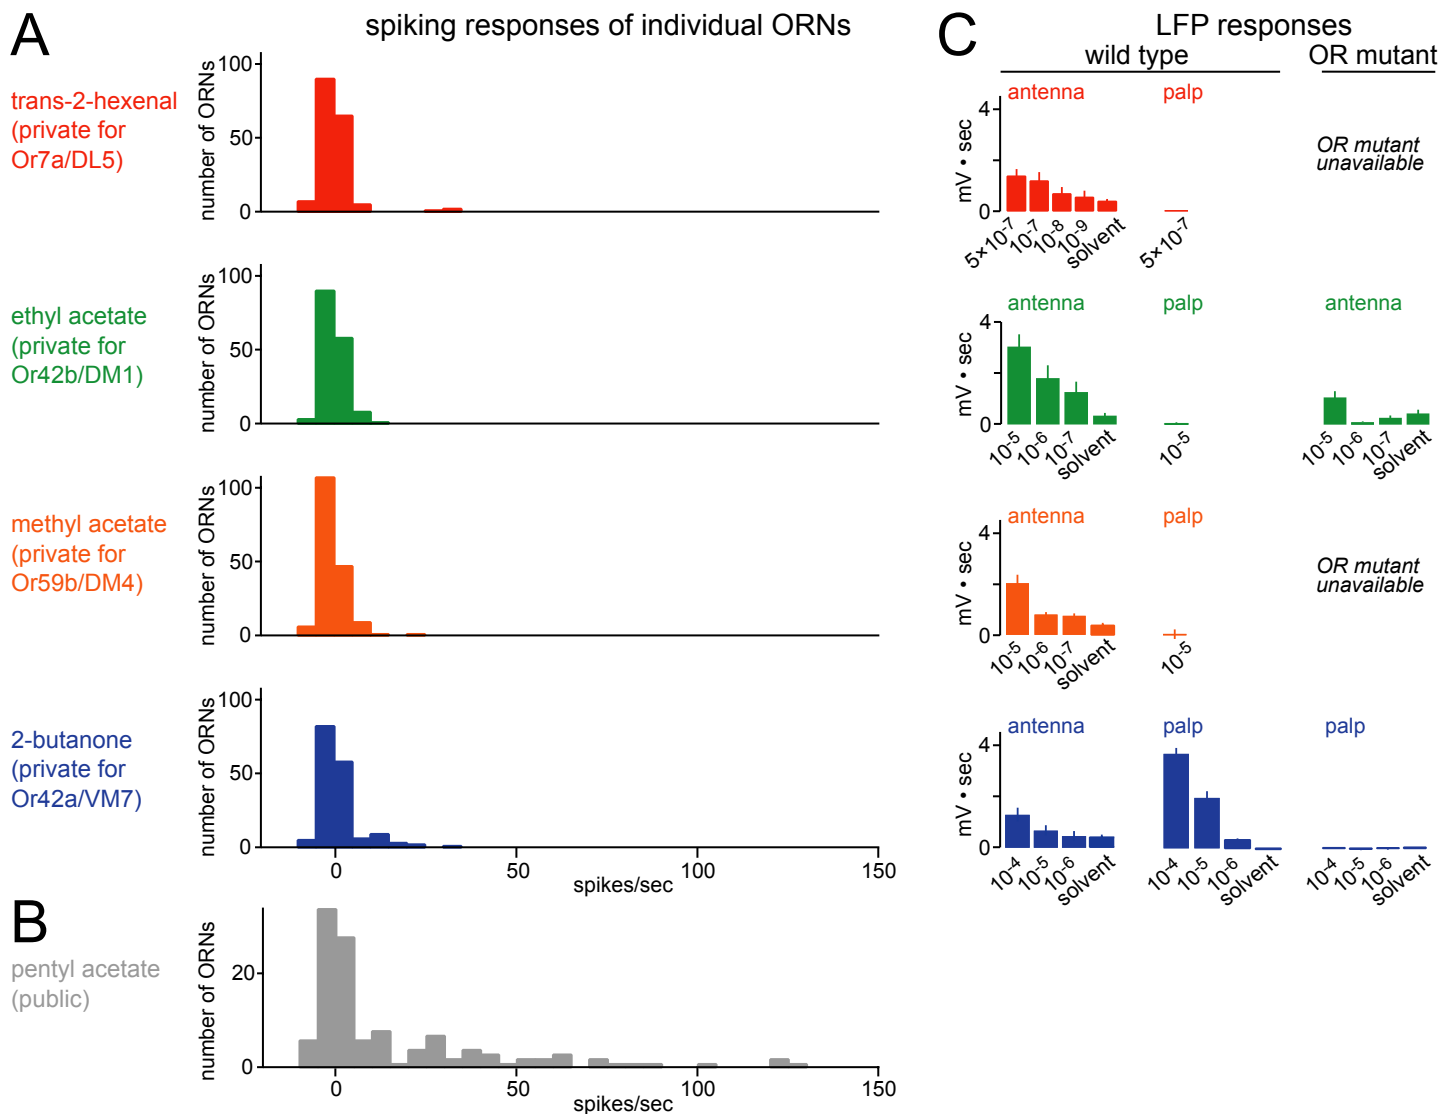


Figure S1: Evidence that “private” odor stimuli activate mainly a single ORN type

A. Histograms show the distribution of responses in non-cognate ORNs to the four private odors at the highest concentrations used in this study. Note that these stimuli evoke almost no response from non-cognate ORNs. The vast majority of responses were essentially zero (median = -0.25 spikes/sec, 10th percentile = -3.1 spikes/sec, 90th percentile = 5.2 spikes/sec). Here we recorded randomly from 223 ORNs, targeting all morphological types to ensure good sampling. The four cognate ORN types corresponding to these private odors were identified by physiological and morphological criteria, and their responses to their private odors were excluded from this sample. See Supplemental Experimental Procedures for details.

B. By contrast, the “public” odor (pentyl acetate) evokes widespread activity at the highest concentration used.

C. LFP responses in the antenna and the palp to the four private odors. In two cases, a mutation was available in the odorant receptor expressed by a cognate ORN (Or42b and Or42a), and here we verified that the LFP response to the corresponding private odor is virtually abolished by the mutation (see Experimental Procedures). Note that the cognate ORNs for the first three of these odors are located in the antenna (Or7a/DL5, Or42b/DM1, Or59b/DM4). For these odors, antennal LFP responses are small and palp responses are zero, implying that mainly a single antennal ORN type is active. Cognate ORNs for the last odor (2-butanone) are located in the palp (Or42a/VM7), where LFP responses to the private odor are larger, presumably because VM7 ORNs are a large fraction of all palp ORNs.

Note that the four “private” stimuli are not strictly private at the highest concentration we used. For example, it is clear the 2-butanone (10^{-4}) evokes a small amount of activity in antennal ORNs. In choosing the concentrations of “private” odors used in this study, we were constrained by the need to drive the cognate PNs for these odors to a level approaching saturation, and this necessitated some compromises for the highest of these concentrations. Nevertheless, the fact that a few non-cognate ORNs respond weakly to these stimuli does not affect our major conclusions.

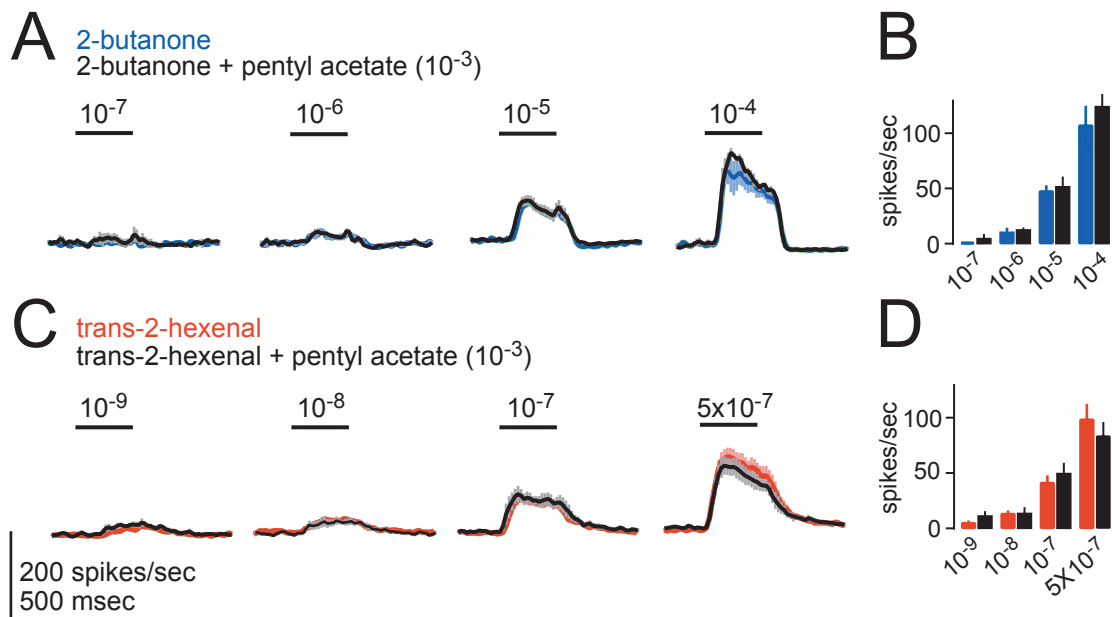


Figure S2: Pentyl acetate does not affect VM7 or DL5 ORNs.

A. The “public” odor pentyl acetate does not affect the response of VM7 ORNs to their private odor, 2-butanone. Recordings from VM7 ORNs show the response to either 2-butanone alone or 2-butanone blended with pentyl acetate (10^{-3} , the highest concentration used in this study). Bar indicates odor delivery period. Each peristimulus-time histogram is a mean of 4-8 recordings, \pm SEM.

B. Average firing rate over 500-msec odor stimulus window.

C-D. Analogous to (A-B) for DL5 ORNs. Private odor is trans-2-hexenal. Each PSTH is a mean of 5-6 recordings.

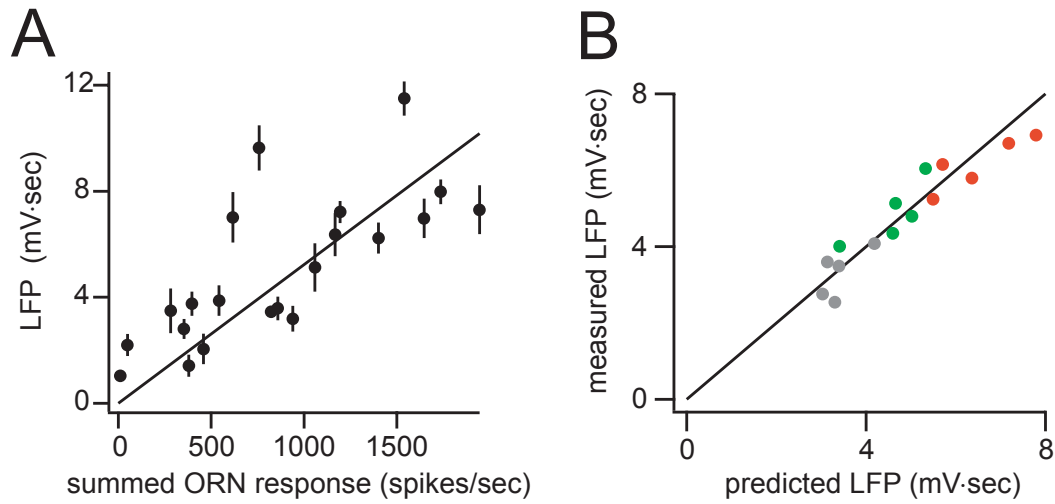


Figure S3: Antennal LFP is linearly related to total ORN spike rate.

A. LFP recordings from the antenna are correlated with the summed response of all 24 ORN types. ORN data is from Hallem and Carlson (2006). Each point is a different odor. The slope of the fitted line is $(1/190) \text{ mV}\cdot\text{sec}^2 / \text{spikes}$.

B. The antennal LFP summates linearly, at least within the response regime we are investigating in this study. We show this by measuring LFP responses to 6 stimuli that activate largely non-overlapping sets of ORNs, and also responses to three blends of these stimuli. LFP responses to the blends were predicted by summing the LFP responses to the individual components. Odors used were ethyl acetate (10^{-6}), methyl acetate (10^{-6}), trans-2-hexenal (10^{-6}), 2-heptanone (10^{-6}), 1-hexanol (10^{-6}), and methyl salicylate (10^{-4}). Each point is a separate experiment. Red points are responses to a blend of all 6 odors; green points are responses a blend of 5 odors (methyl acetate was held out); gray points are responses a blend of ethyl acetate (10^{-6}) and trans-2-hexenal (10^{-6}).

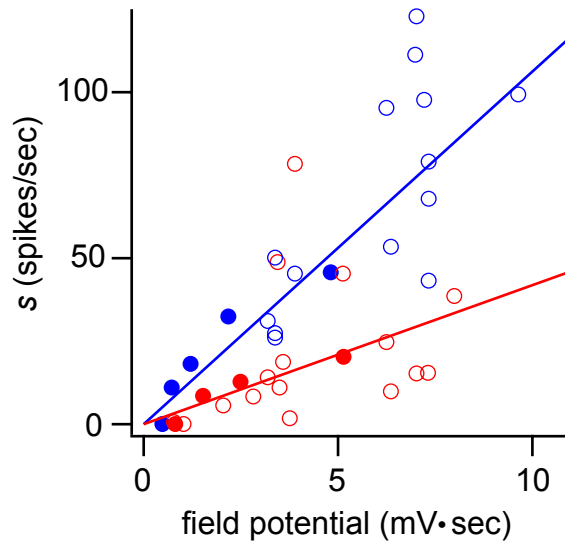


Figure S4: The suppression factor s and the LFP are linearly related.

Here, the solid circles and fitted line are reproduced from Fig. 3h. The open symbols represent values computed from the data in Fig. 4, as follows. For each of the test stimuli in Fig. 4, we obtained the value of the suppression factor s by taking the measured PN firing rate and the measured ORN firing rate, and plugging these into Equation (2). For each test stimulus, we have also measured the antennal LFP. Plotting s versus the LFP response reveals a significant linear correlation between s and LFP ($p < 0.01$ for each glomerulus, Pearson's correlation, all data from Figs. 3 and 4 pooled together for each glomerulus). Moreover, this figure shows that the slope of this relation for the open symbols is well-described by the line fitted to the solid symbols. A sublinear function (exponential or hyperbolic) would not be a good description of this data.

In essence, this analysis explains why Equation (2) and Equation (4) together produce such good predictions of PN responses (Fig. 4): s is indeed a linear function of the LFP, and this can be used to predict how much lateral inhibition suppresses the PN's response to its direct ORN inputs.

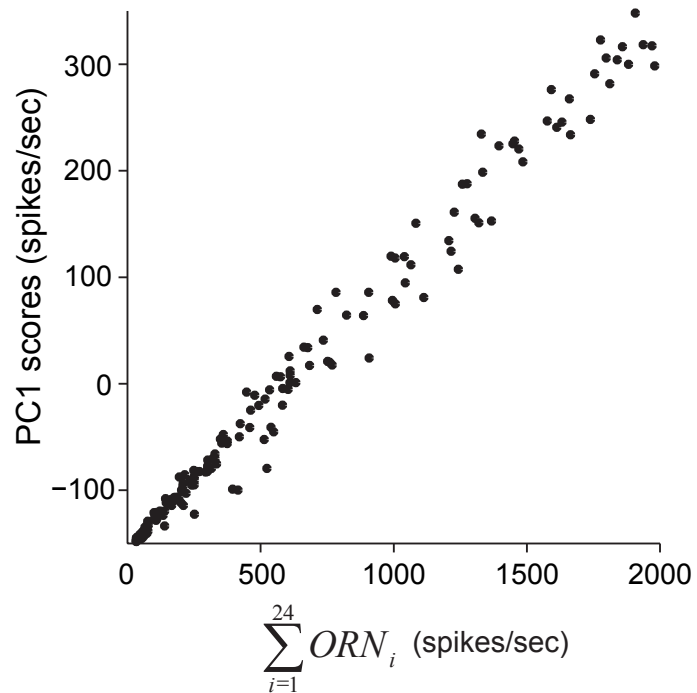


Figure S5: Odor intensity is the stimulus feature that accounts for the largest share of the variance in ORN responses.

We performed PCA on the 24-dimensional ORN response vectors obtained from Hallem & Carlson (2006). Here we plot scores on the first PC versus the total ORN spike rate evoked by that stimulus (each point is a different stimulus). This plot shows that scores on the first PC are highly correlated with the total number of ORN spikes evoked by that stimulus. This indicates that PC1 is essentially a proxy for the intensity of the stimulus from the perspective of the ORN population. This is why normalizing for intensity reduces correlations among PNs: stimulus intensity is a large source of covariance in this population.




Figure S6: Modeling noise in PN responses.

We used data collected for a previous study (Bhandawat et al., 2007) to derive a model of how trial-to-trial variability depends on mean firing rate. In collecting this data set, each odor was presented for a block of 5 trials per cell. We measured the firing rate of individual recorded PNs over the entire 500 msec odor presentation window for each trial. For each block of trials, we computed both the mean firing rate and the standard deviation (SD) of the firing rate. Each point corresponds to a different block of 5 trials and a different cell/odor combination ($n = 787$ blocks of trials). The blue line shows SD values binned by mean (bin width = 20 spikes/sec). We fit this data with an exponential function (red line):

$$SD = (9.5 \text{ spikes/sec}) - (7.2 \text{ spikes/sec}) \cdot e^{-\text{mean}/(76 \text{ spikes/sec})}$$

This function was used to simulate noise in PN responses (Figs. 6-7). Each entry in the matrix was treated as the trial-averaged mean firing rate for that stimulus-PN combination. For each mean firing rate, we used this equation to define a SD. We drew noise randomly from a Gaussian distribution with having this SD and a mean of zero.

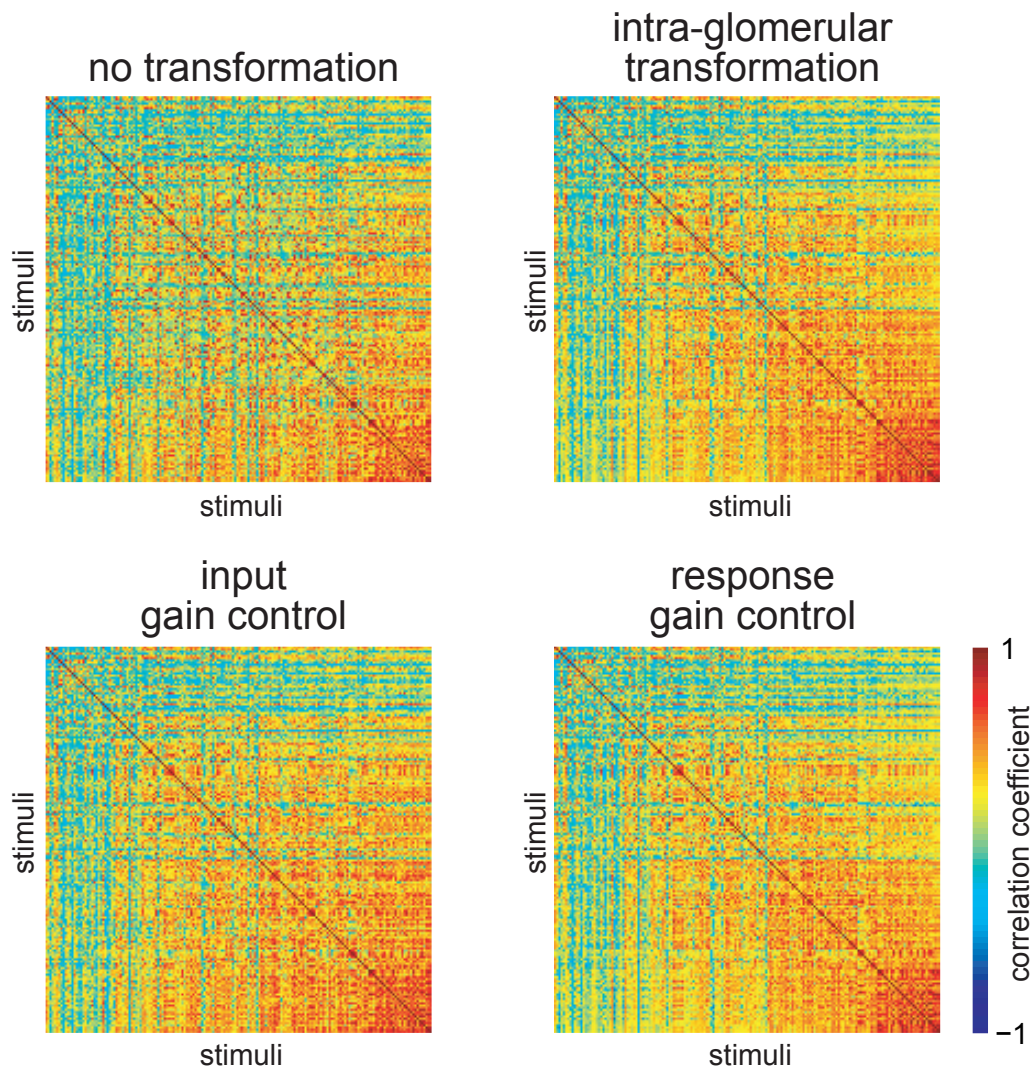


Figure S7: Correlations between stimuli are relatively unaltered in our model of antennal lobe transformations.

Each 176x176 matrix shows all pairwise correlation coefficients between stimuli. The first matrix shows correlations between stimuli in the ORN data (i.e., no transformation). Note that none of our simulated antennal lobe transformations has much effect on these correlations. The intra-glomerular transformation slightly increases correlations because it tends to push PNs toward saturation, but this effect is small. The two gain control models have little effect on correlations between odors because they affect all glomeruli uniformly in a manner that does not depend on the odor tuning of each glomerulus. Thus, each population odor response tends to be scaled down by some amount that depends on the odor, but the correlations between different odor-evoked glomerular patterns are not changed substantially. In geometric terms, the angles between the 176 (24-dimensional) odor response vectors are relatively unchanged. What *does* change is the norm of these vectors, such that all vectors now have more equal norms (Figure 5). This reduces correlations between glomeruli -- but not correlations between stimuli.

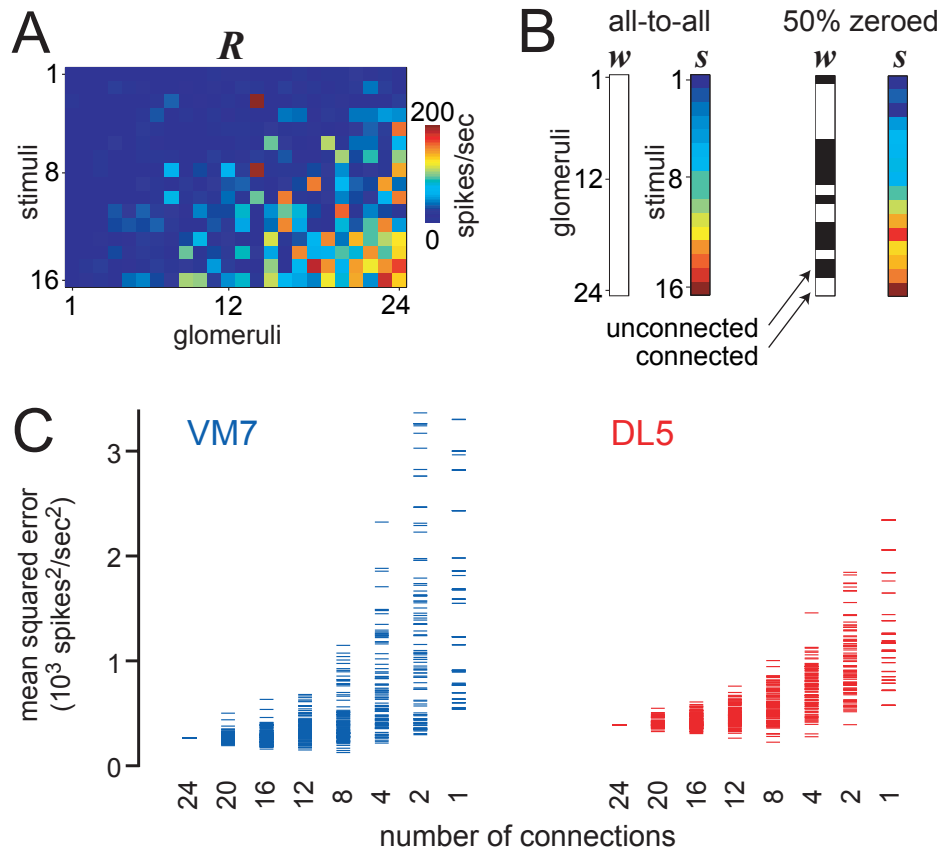


Figure S8: Impact of sparsening inter-glomerular connectivity on PN response predictions.

Throughout this study, we have assumed that all ORNs contribute equally to the normalization pool. Here we examine how prediction quality changes if we instead assume sparser connectivity. To investigate this, we allowed glomeruli to make differing contributions to the normalization pool by replacing the sum in Equation (6) with a weighted sum:

$$s = m(\sum_{i=1}^{24} w_i r_i) / 190 \text{ mV}\cdot\text{sec}^2/\text{spikes}$$

We then manipulated the vector w of all the glomerular weights (w_1, w_2, \dots, w_{24}) and examined how this affected the predictions of the model.

A. Responses to 16 odors in 24 ORN types (data from Hallem & Carlson, 2006). These 16 odors are the same as the set used to test predictions of the DL5 model (Fig. 4).

B. Inhibitory input to DL5 evoked by 16 odors is shown as a vector s with 16 elements. This vector represents the value of the suppression factor s for each odor (maximum inhibition in red, and minimum in blue). This vector is obtained by multiplying the ORN response matrix R by the 24-dimensional weight vector w and scaling by a constant. Two cases are shown: all-to-all connectivity (all weights =1) and sparser connectivity (50% of weights randomly zeroed). (In order to keep to the total level of inhibition constant when zeroing some weights, we scaled the non-zero weights by a factor of $24/N$, where N is the number of non-zero weights.) Because ORN responses are correlated with each other, the inhibition vector is similar for these two cases: weak stimuli (top) typically produce weak inhibition while intense stimuli (bottom) produce strong inhibition.

C. We used the equation above to determine the value of the suppression factor s for each odor. We then used Equation (2) to predict PN responses and computed the mean squared error of the prediction. For each connection number we ran 100 simulations. Each symbol represents the mean squared error for one iteration of the simulation.

Note that, on average, predictions degrade as more weights are randomly zeroed. Nevertheless, it is possible to find sets of weights that generate good predictions. This is not surprising because there are many degrees of freedom in this fit. Indeed, these sparse models overfit the data, because we can also obtain good predictions when we shuffle the odor labels on the ORN response vectors before fitting. In sum, the predictive power of our model does not place strong constraints on connectivity between glomeruli. What is notable is that good fits are obtained with the simplest model, where all glomeruli make an equal contribution.