Neuron
Article

Neural Encoding of Rapidly Fluctuating Odors

Maria N. Geffen,[1,2,5] Bede M. Broome,[3,5] Gilles Laurent,[3,*] and Markus Meister[1,4,*]
[1]Program in Biophysics, Harvard University, Cambridge, MA 02138, USA
[2]Center for Studies in Physics and Biology, Rockefeller University, New York, NY 10065, USA
[3]Division of Biology, California Institute of Technology, Pasadena, CA 91125, USA
[4]Department of Molecular and Cellular Biology and Center for Brain Science, Harvard University, Cambridge, MA 02138, USA
[5]These authors contributed equally to this work
*Correspondence: laurentg@caltech.edu (G.L.), meister@fas.harvard.edu (M.M.)
DOI 10.1016/j.neuron.2009.01.021

SUMMARY

Olfactory processing in the insect antennal lobe is a highly dynamic process, yet it has been studied primarily with static step stimuli. To approximate the rapid odor fluctuations encountered in nature, we presented flickering "white-noise" odor stimuli to the antenna of the locust and recorded spike trains from antennal lobe projection neurons (PNs). The responses varied greatly across PNs and across odors for the same PN. Surprisingly, this diversity across the population was highly constrained, and most responses were captured by a quantitative model with just 3 parameters. Individual PNs were found to communicate odor information at rates up to ~4 bits/s. A small group of PNs was sufficient to provide an accurate representation of the dynamic odor time course, whose quality was maximal for fluctuations of frequency ~0.8 Hz. We develop a simple model for the encoding of dynamic odor stimuli that accounts for many prior observations on the population response.

INTRODUCTION

Odors in the natural environment are transported by turbulent flow of air or water. In these chaotic eddies, odor filaments from multiple sources may be densely interleaved. An animal navigating through such an odor plume encounters rapid variations in odor content, exacerbated by its own intermittent movements (Koehl et al., 2001; Murlis et al., 1992; Carde, 1996) and breathing rhythms. These rapid fluctuations may pose some challenges to identifying an odor; on the other hand, the variation itself may provide important cues about the odor source (Hopfield, 1991; Murlis et al., 1992). One suspects that the neural circuits that process olfactory signals are adapted to this ecological reality. However, much of what we know about olfactory processing stems from studies using steady exposure to single odors (with a few exceptions: Stopfer and Laurent, 1999; Brown et al., 2005; Vickers et al., 2001). Here, we probe the function of the early olfactory system in locusts, with a particular focus on very dynamic odor stimuli, whose time course approximates that in turbulent plumes.

In the insect brain, the antennal lobe performs the first level of olfactory processing and encoding. Three populations of neurons interact in this circuit: olfactory receptor neurons (RNs) provide the sensory input signals, projection neurons (PNs) transform and transmit RN output to subsequent brain areas, and local neurons (LNs) mediate mostly inhibitory interactions through a densely connected local network. Broadly speaking, RNs excite PNs and LNs; PNs excite LNs but not other PNs (R.A. Jortner and G.L., unpublished data); LNs inhibit PNs, other LNs postsynaptically (Hansson and Anton, 2000) and possibly RNs presynaptically (Wachowiak et al., 2005; Olsen and Wilson, 2008). Both anatomy and function of this circuit are closely analogous to those of the vertebrate olfactory bulb (Friedrich and Laurent, 2001; Korsching, 2002; Wilson and Mainen, 2006).

By comparing the output signals of the antennal lobe—the firing of PNs—to the input—the odor concentration at the antenna—one finds that this circuit imposes a great deal of temporal patterning (Laurent, 1997; Laurent et al., 1996; Wehr and Laurent, 1996, 1999). Following an odor step of steady concentration, the firing rate of a PN typically waxes and wanes over several seconds. The dynamics of this response are reproducible but vary across PNs for a given odor and across odors for a given PN (Laurent et al., 1996). Given this intricate time-dependent response, one might worry that adding a strongly time-modulated stimulus would greatly complicate analysis and understanding. Fortunately, this is not the case. We found that a simple mathematical relationship captures the response of PNs to a time-varying odor stimulus: it includes a dynamic linear filter and a static nonlinear transformation. The dynamics of odor integration varied across PNs and across odors, reflecting the multitude of signaling pathways within the antennal lobe. Nevertheless, the range of dynamic behaviors in the population was very restricted and could be parametrized by a single number. These results lead to a "forward" model that predicts accurately the PN response to a rapidly varying stimulus and accounts for earlier observations on antennal lobe dynamics. We also considered the "reverse" problem of extracting the stimulus from the PN spike trains. Provided prior knowledge of the odor identity, we found that less than a handful of PNs suffice to reconstruct the time course of odor concentration to remarkable accuracy.

RESULTS

Projection Neurons Reliably Track Rapidly Varying Odors

We recorded the firing of 97 PNs in 8 locusts while an odorized air stream flowed over the antenna. The odor was mixed into carrier

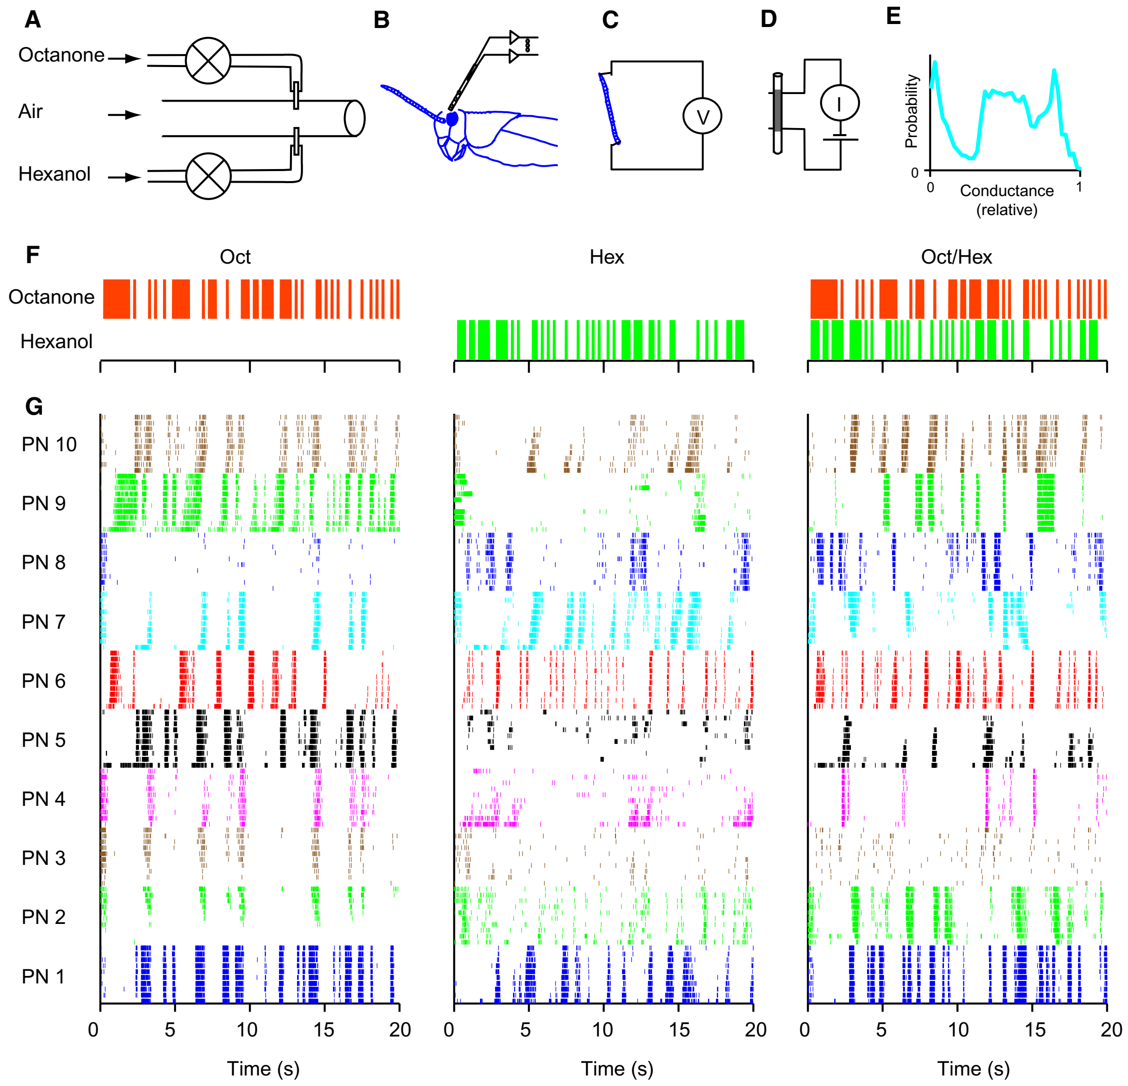

Figure 1. Projection Neurons Respond Reliably to Rapidly Flickering Odors

(A) Odor delivery. A large tubing delivered carrier air. Two needles inserted near the orifice injected dilute odorant vapors. Each odorant tube was switched by a solenoid valve. For single-odor experiments only one odorant tube was active.

(B) Locust preparation. One antenna was exposed to the odor delivery tube. Several silicon tetrodes were inserted in the antennal lobe for extracellular recording.

(C) Electroantennogram. A locust antenna was detached and thin wires inserted at either end. The voltage between the wires was recorded during odor exposure.

(D) Electric nose. A thin glass capillary tube was coated with a carbon-doped polymer film. A DC voltage was applied and the current measured during odor exposure.

(E) Histogram of the odor concentration during the experiment, as measured by conductance of the electric nose. The center 8 deciles of the distribution span a >10-fold range of concentration.

(F) State of the two odor valves in (A) during M sequence experiments using octanone alone (left), hexanol alone (center), and octanone and hexanol combined (right). These conditions will be called "Oct," "Hex," and "Oct/Hex." The full sequence lasted ~102 s.

(G) Response of 10 sample PNs under the flickering odor stimuli of panel F. For each neuron, the raster plot shows spikes during 10 identical stimulus repeats (sequence bottom to top).

air by a valve that opened and closed rapidly (see Experimental Procedures; Figures 1A–1E). A pseudorandom valve control signal was designed that effectively explored all odor fluctuations on the time scale of 0.2 s to 2 s. We also built a small odor sensor to confirm that changes in concentration at the antenna were indeed both rapid and reproducible across repeats. We used the odors hexanol and octanone, with concentrations adjusted to produce comparable activation of RNs (assessed with an electroantennogram) and of PNs. The two odors were presented individually, each according to a different temporal sequence. In a third experiment, the two odor sequences were mixed together in the air stream.

Many projection neurons responded reliably to one or both of these odors with firing locked to the stimulus sequence over many repeated trials (Figures 1F and 1G). Inspection of the raw raster plots reveals a great diversity among the responses. Some PNs are excited by the odor, others inhibited. Some PNs fire only sparsely at precise instants during the sequence; others are active half the time or more. Some PNs can track even the shortest odor pulse; others seem to respond more slowly. Finally, PNs that behave similarly under one odor can differ greatly under another. We will now show that this observed diversity can be mapped onto a small number of response parameters.

A Simple Model Predicts Most PN Responses

Direct inspection of raster plots under a complex stimulus (Figure 1G) can give a qualitative impression, but any deeper interpretation requires a quantitative analysis of each cell's behavior. For this purpose we sought a mathematical model of the response that would predict the time-varying firing rate of each PN from the preceding stimulus history. Here, we present this model and its performance. Subsequent sections will use this tool to inspect different aspects of PN behavior.

In the present experiments, a useful response model should capture whether a PN is excited or inhibited by the odor, the kinetics of its response, and whether it fires sparsely or densely. The simplest formalism that meets these demands is the linear-nonlinear (L-N) model (Figure 2A). In this scheme (Experimental Procedures; Equation 2), the stimulus waveform is first passed through a linear temporal filter with impulse response $F(t)$. Then the output of that filter is transformed by a nonlinear function $N(g)$ to give the predicted firing rate $r(t)$. The waveform $F(t)$ of the linear filter describes how the model neuron "weights" the odor stimulus at various times in the past. In turn, the shape of $N(g)$ accounts for all the instantaneous nonlinearities in the response. For now we merely seek a phenomenological model and thus refrain from a biophysical interpretation of the filter and nonlinearity. For each PN and each odor, we optimized the shapes of $F(t)$ and $N(g)$ such that the time-varying firing rate predicted by the model came closest to the observed firing rate (see Experimental Procedures).

We first discuss in detail how the model fits the responses from two sample neurons (Figures 2C–2H), followed in later sections by statistics about the population. One of the sample neurons (Figures 2C–2E) was inhibited by octanone, as revealed by the negative-going filter shape. In fact, the time course of the filter was biphasic: effectively, this neuron averaged the concentration in the preceding 0.4 s and subtracted it from the concentration in the 1 s prior to that. The nonlinearity resembled a half-wave rectifier: no firing for a negative filter output (increase in odor) and proportional firing for a positive output (decrease in odor). The model prediction comes very close to the real firing rate, with some systematic discrepancies that will be discussed below. The same PN was also inhibited by hexanol, with a similar filter shape and nonlinearity (Figure 2D). The other sample neuron (Figures 2F–2H) was excited by an increase in octanone, as revealed by a positive biphasic filter shape (Figure 2F). It was inhibited by hexanol (Figure 2G). The nonlinearity was rectifying, as for the preceding neuron, but with a higher threshold. Correspondingly, this neuron fired somewhat less often. Again, the L-N model accounted for many features of the neuron's time-varying firing rate, though better for one odor (Figure 2F) than the other (Figure 2G).

We also stimulated the antenna with the two different odor sequences superposed. This two-odor stimulus delivered essentially the arithmetic sum of the odor concentrations from the two single-odor experiments. However, the firing responses of the PNs were clearly different from the arithmetic sums of their single-odor responses (Figures 2E and 2H; see also Broome et al., 2006). Because the response to single odors is already nonlinear (Figures 2C–2D, 2F, and 2G), one should not expect plain summation across odors. A simple formalism that allows nonlinear summation is the two-dimensional L-N model (Figure 2B). In this scheme, the two odor sequences are each passed through a linear temporal filter, and the two filter outputs get combined through a nonlinear function to produce the predicted firing rate. As before, the shapes of the temporal filters represent the dynamics of stimulus integration, and the nonlinear function reflects how the two odor stimuli interact. For each PN, the two filters and the nonlinearity were optimized to produce the best fit to the measured firing response.

For the two sample neurons (Figures 2E and 2H), this two-dimensional L-N model correctly predicted most, though not all, of the firing episodes. For both neurons, the nonlinearity predicted a strong response only when both filter outputs are high. Consequently, the predicted firing rate was considerably sparser than the arithmetic sum of the two individual responses, consistent with the measured responses. An interesting observation emerged from inspecting the temporal filters in the two-odor condition. For one neuron (Figure 2E), these filter shapes were almost identical to the corresponding filters in the single-odor conditions. This implies that both odors had very similar effects on the firing rate as when they were applied individually. However, for the other neuron (Figure 2H), the filter for hexanol changed sign: this odor inhibited the PN when applied by itself, but excited it in the two-odor experiment. Such changes in behavior will receive more scrutiny below.

To assess how well these L-N models capture neural responses across the population, we compared several statistics of the fits (see Experimental Procedures). A simple indicator of the quality of the L-N fit is the difference between the predicted and the actual firing rate: we measured the root-mean-squared deviation between those two functions and call this the "residual" ($\sqrt{P_R}$ in Equation 1). To evaluate this, we compare it to a measure of reliability in the neural response, for which we used the standard deviation of individual trials from the average; this will be called the "noise" ($\sqrt{P_N}$ in Equation 1). Finally, both these quantities may be compared to the actual magnitude of the neuron's response. We measured this as the standard deviation of the firing rate and call it the "signal" ($\sqrt{P_S}$ in Equation 1).

Under all the stimulus conditions, the signal/noise ratio of the response varied a great deal across neurons (Figure 3); mostly this reflects the degree to which a neuron was responsive to the odor tested. In the two conditions with single odors, the L-N fits performed well (Figures 3A and 3B). For most neurons (92%, dots above diagonal in Figures 3A and 3B), the residual was smaller than the noise. This means that the L-N fit came closer to the average response than the typical single trial. Even for the PNs with the highest signal/noise ratio, the L-N model

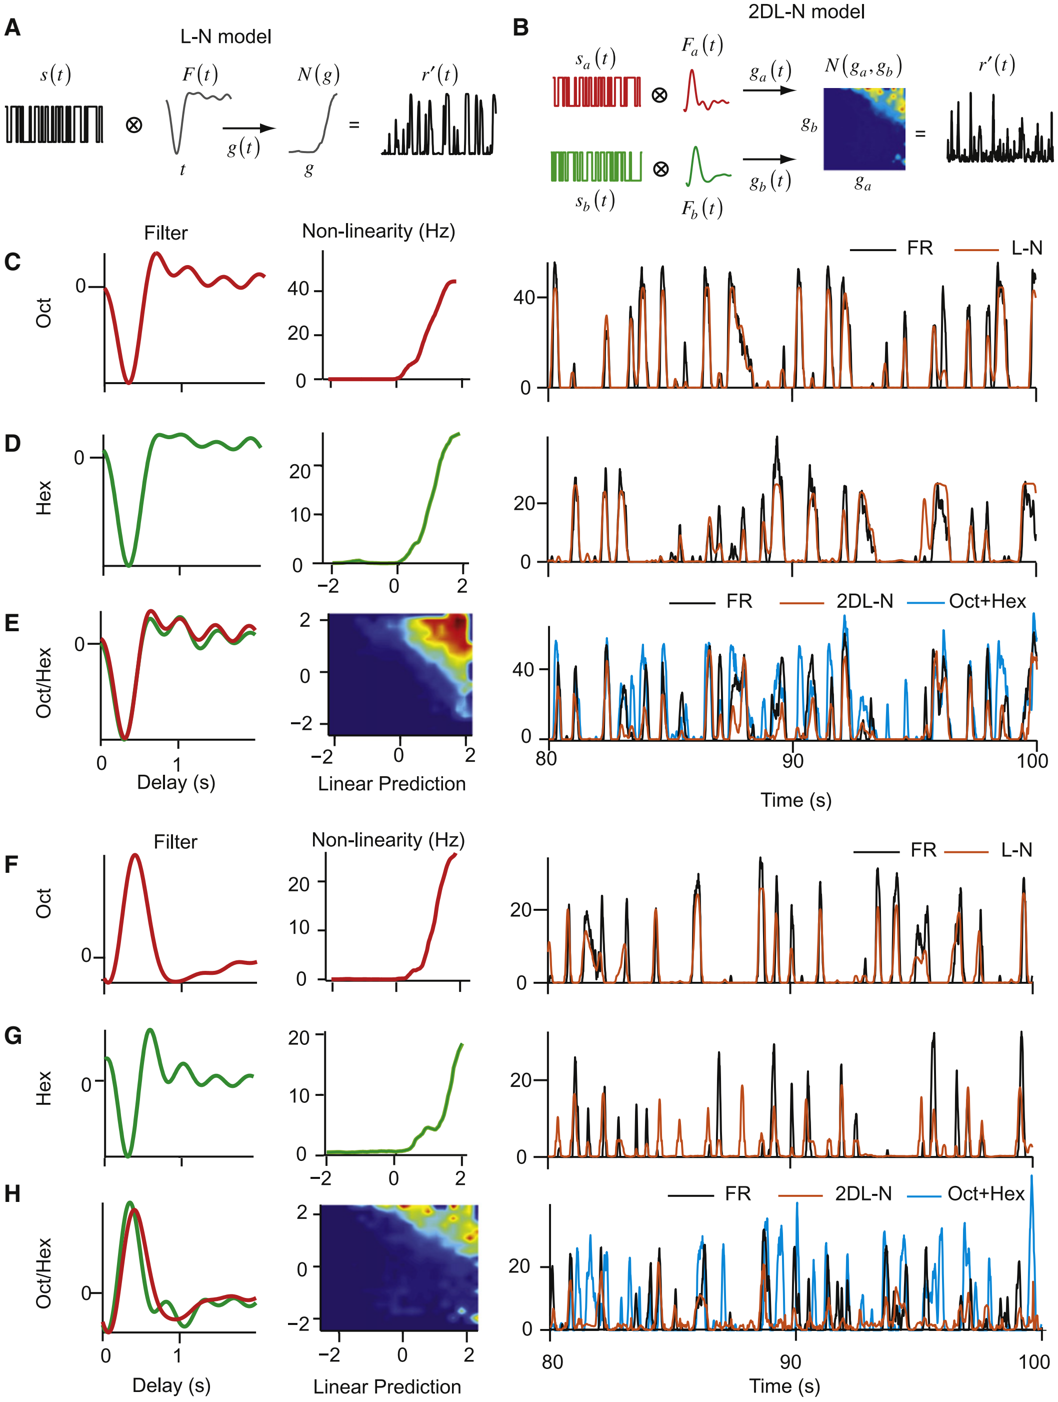

Figure 2. A Linear-Nonlinear Model Can Capture the Dynamic Odor Responses

(A) Diagram of the L-N model used to predict the PN response from the odor valve signal in the Oct and Hex single-odor experiments (see Equation 2).

(B) Diagram of the 2DL-N model used to fit the Oct/Hex two-odor experiment. Here each of the odor valve signals is passed through a separate filter, and the results combined in a nonlinear function of two variables to yield the predicted firing rate (see Equation 5).

(C–E) Responses and model fits of a sample neuron (PN1 in Figure 1G) in the three conditions: Oct (C), Hex (D), and Oct/Hex (E). Left: Filter functions used in the respective model. Middle: Nonlinearities that relate *g*(*t*) to the firing rate; note the function in (E) is two-dimensional (blue, low; red, high). Right: Actual firing rate averaged over 10 trials (FR), model fit (L-N and 2DL-N), and the sum of the two single-odor firing rates (Oct+Hex in [E]).

(F–H) Responses and model fits of another sample neuron (PN5 in Figure 1F), displayed as in (C)–(E).

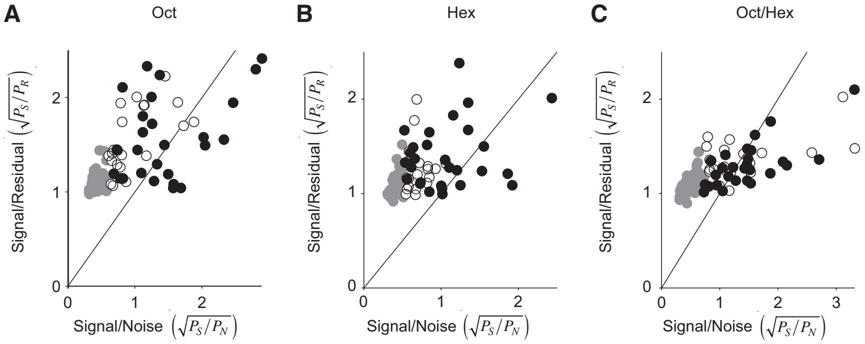

Figure 3. Linear-Nonlinear Models Provide a Good Fit for Many PN Responses

Statistics for the responses of all recorded PNs and their model fits in the three conditions: Oct (A), Hex (B), and Oct/Hex (C). Each panel shows one point per cell in a scatter plot of Signal/Residual ($\sqrt{P_S/P_R}$, see Equation 1) against Signal/Noise ($\sqrt{P_S/P_N}$). Diagonal line is the identity. For a cell with a point above this line, the model fit is more accurate than the trial-to-trial variation in the response. For further analysis of the model parameters, we used the top 50% of cells according to the reproducibility of the response, P_S/P_N. Gray symbols: cells not used for any analysis. Open symbols: cells used in analysis of one condition. Black symbols: cells used in analysis of both Oct and Hex (A and B), or in all three conditions (C). The numbers of cells selected this way were: Oct, 48; Hex, 48; Oct and Hex, 28; Oct/Hex, 48; Oct and Hex and Oct/Hex, 27.

approached this standard of accuracy. Because natural behavior is generally based on single trials of any given odor sequence, one can conclude that the L-N model is an adequate expression of how such a neuron responds. In the two-odor condition, the prediction error of the L-N model was somewhat larger (Figure 3C; 63% of neurons had smaller residual than noise). For the PNs with the highest signal/noise ratio, the fit's error was about twice the trial-to-trial noise.

For several PNs, we noticed substantial drifts in the firing rate over the 10 stimulus repeats (~20 min; Figure 1G). This likely reflects some adaptation process (Stopfer and Laurent, 1999) rather than general decline of the preparation, since neurons displaying increases and decreases of the firing rate could be observed simultaneously. In such cases we performed separate model fits to the early and late trials (see Figure S1 available online). The fits performed equally well at both ends. In the course of the drift, the nonlinearity changed, while the filter remained unchanged.

A PN's Response Is Characterized by Just Three Numbers

In the following sections, we analyze the repertoire of response properties in the PN population, based on the filter function and nonlinearity obtained from the L-N model. To be confident of the results, we restricted the analysis to sets of PN-odor pairs with large signal/noise ratios, specifically the best 50% by this criterion (Figure 3; Experimental Procedures). The excluded data were primarily from neurons that happened to respond only weakly to the odor in question, leading to an unreliable estimate of their response parameters. Some key analysis steps were repeated on the full data set.

The shape of the filter encompasses all the dynamics of a PN's response, and specifies how the response integrates odor concentration from various times in the past. We collected all the filter shapes from PN-odor combinations that passed the above quality test, including single-odor and two-odor experiments. The absolute amplitude of the filter function plays no role in the response dynamics, since it is redundant with the horizontal scale of the nonlinearity (Equation 2); therefore, we normalized all the filters to have unit variance. This family of filter shapes was subjected to a principal component analysis, which revealed two dominant components: one had an almost monophasic shape, the other was sharply biphasic (Figure 4A). These two principal components accounted for 74% of the variance in the filter set. When projected onto this two-dimensional subspace, the filter functions fell close to a unit circle (Figure 4B). Consequently, the filters could be parametrized almost entirely by a single variable, the angle ϕ around the unit circle (Figure 4B). Specific values of ϕ correspond to different response types: for example, the Off-type hexanol response of Figure 2D has $\phi = 0.9\pi$, and the On-type octanone response of Figure 2F has $\phi = .06\pi$. When all the original filter functions were sorted by this parameter ϕ, they did indeed demonstrate a systematic and smooth progression of shapes (Figure 4C).

Further evidence that ϕ is sufficient to characterize the odor response dynamics came from inspecting the raw firing responses themselves. When one plots the firing rate of each neuron during the flickering odor sequence (Figure 4D) and sorts these responses according to the parameter ϕ (Figure 4E), it is clear that responses of similar type get grouped together. One finds many events of high firing rate that are stereotyped across cells (Figure 4D) but shift gradually in time with increasing ϕ. This shift reflects the gradual change in the filter shape (Figure 4F). Therefore, what appeared as a broad and complex diversity of response types in the raw measurements (Figure 1G) can now be recognized as a continuous range of dynamics that vary only along a single scalar dimension. Remarkably, the same principal component shapes (Figure 4A) that define the parameter ϕ are appropriate for experiments with both single odors and also the two-odor condition (Figure 4B). Furthermore, the distribution of shapes is remarkably similar under the two odors (Figure 4E), even though each individual neuron may change filter shape (Figures 2F and 2G). This suggests that these shapes reflect global aspects of processing in the antennal lobe, rather than any odor-specific response properties.

To increase confidence in this analysis, we repeated it on the full set of PN-odor combinations without any quality selection (Figure S2). The first two principal components had the same

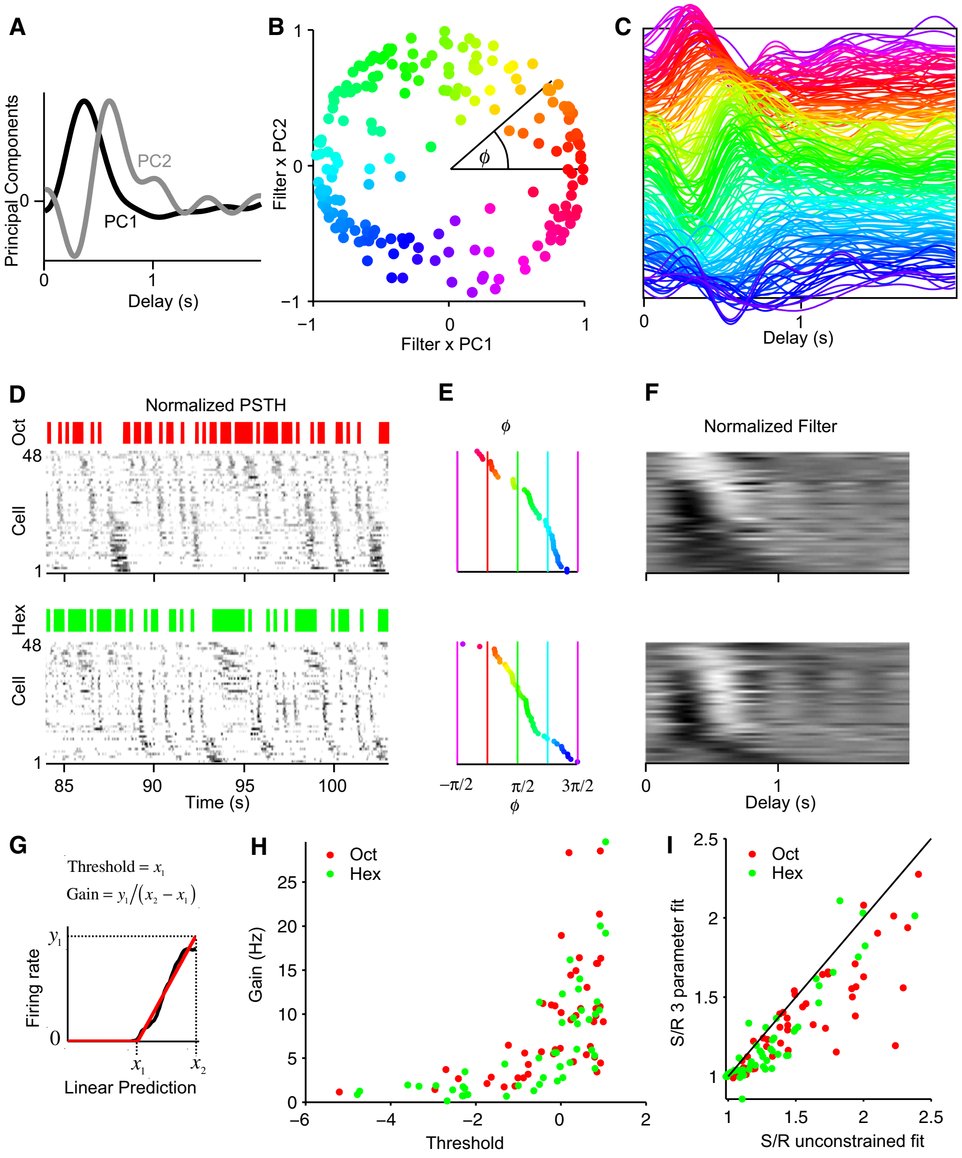

Figure 4. The Diversity in PN Response Dynamics Can Be Captured by One Parameter

A statistical analysis of 192 filter waveforms from the Oct, Hex, and Oct/Hex conditions, selected by the criteria of Figure 3.

(A) The first two principal components of the filter set. These account for 74% of the variance in the set of waveforms.

(B) Projection of all filter waveforms onto the two principal components in (A). Because each filter is normalized to unit variance and the two axes account for much of that variance, the points lie close to the unit circle. The points are color-coded by the angle ϕ around the origin.

(C) Plot of all filter waveforms, ordered top-to-bottom and color-coded by the angle ϕ from (B).

(D–F) Kinetics of PN responses, ordered by the angle ϕ. (D) Measured firing rate of each neuron, one row per cell. The rate is represented in grayscale, normalized to the average firing rate for each neuron. Here and in (E), (F), (H), and (I) the Oct and Hex plots include different sets of 48 PNs selected as in Figure 3. (E) The value of ϕ for each cell's filter. (F) The corresponding filter functions, represented in grayscale (white positive, black negative).

(G–I) Nonlinearity of PN responses. (G) Example of a threshold-linear fit to a neuron's nonlinearity, see Equation 7. (H) Relationship between gain and threshold for PN responses in the Oct and Hex conditions. (I) Quality of the 3-parameter response fit. For each cell, the signal/residual ratio of the 3-parameter L-N fit (ordinate, see Equation 9) is plotted against that of the unconstrained L-N fit (abscissa, Equation 2).

shape as before. The filters for the previously excluded responses were simply contaminated by greater amounts of noise, as expected from their weaker odor sensitivity.

The second component of the L-N model is the nonlinearity. This function generally had a sigmoid shape with a baseline of zero (Figure 2), reflecting the fact that the PN remains silent except for certain periods when the stimulus crosses some threshold (Figure 1G). Only rarely did we encounter significant saturation at the top of the sigmoid. Thus, we fitted this shape with a straight line, whose intercept and slope determine the threshold and the gain of the response (Figure 4G). These two response parameters were somewhat correlated across the

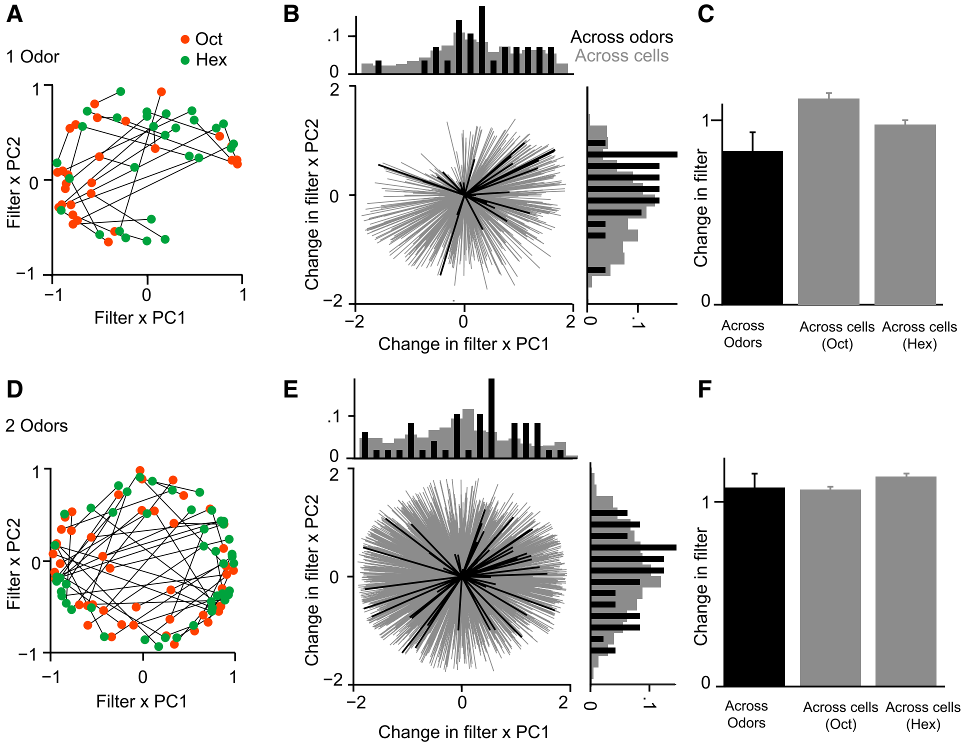

Figure 5. Response Dynamics Differ across Odors as Much as across Neurons

An analysis of the difference in a neuron's filter under the two odors.

(A) For each neuron, the filters in the conditions Oct and Hex (circles joined by a line) are plotted on the graph introduced in (B). Analysis based on 28 PNs with reliable responses under both Oct and Hex conditions (Figure 3).

(B) For each neuron from (A), the difference between the two filters for Oct and Hex is plotted as a vector from the origin (black). For comparison, the background (gray) shows the difference between the filters of two different neurons to the same odor (either Oct or Hex), plotted for all possible pairs. Marginal histograms show the distribution of differences along the two principal component axes.

(C) The average length of the difference vectors for the following filter comparisons: Left: same PN, Oct versus Hex. Middle: Oct, different PNs. Right: Hex, different PNs. Error bars: SEM. The change in the filter shape of the same PN across odors is as large as the spread of filters for the same odor among different PNs.

(D–F) Same analysis, but for pairs of filters that result from 2DL-N fits for the Oct/Hex condition, presented as in (A)–(C). Based on 48 neurons with reliable responses in the Oct/Hex condition (Figure 3).

population (Figure 4H): a combination of low threshold and low slope reflects a response that is only weakly driven by the odor. High threshold and high gain indicate a selective and strongly driven response. The two odors elicited the same range of nonlinearities (Figure 4H), as observed previously for the filter (Figure 4E).

It appears therefore that each PN's dynamic response to a given odor can be summarized by just 3 numbers: ϕ for the filter shape, and threshold and gain for the nonlinearity. To test this explicitly, we inspected the fits obtained from this 3-parameter L-N model (Equation 9) and found that they match the true firing rate almost as well as the unconstrained L-N fits (Equation 2), with very little compromise (Figures 4I and S3).

Diverse Response Dynamics across Neurons and Odors

An important issue in exploring the PN population is whether each neuron has a uniform response time course across odors. For example, one neuron might be excited by all odors, another inhibited. Contradicting this simple notion, we have already encountered a PN with filters of opposite sign for hexanol and octanone (Figures 2F and 2G) and indeed a diversity of response waveforms to different odors has been noted in previous work (Laurent et al., 1996). The characterization of response dynamics developed here (Figure 4) allows us to approach this issue more quantitatively and test for a neuron-specific bias.

For this purpose, the analysis focused on neurons with reliable filter measurements for both odors, presented either singly or together (Figure 3). For each PN, we measured the change in the filter shape between one odor and the other (Figure 5). These PN-specific changes across odors covered a wide range; in fact, they were as large as the differences between the responses to the same odor of two arbitrarily chosen PNs (Figures 5C and 5F). Thus, there was no detectable correlation between a given neuron's filters for different odors. In a statistical approximation, this suggests a picture of the PN population in which each neuron draws its response parameters at random from a limited set of possible filters, and draws them independently for each odor.

Once a neuron's response dynamics for a given odor are known, can they be altered by the context from other odors?

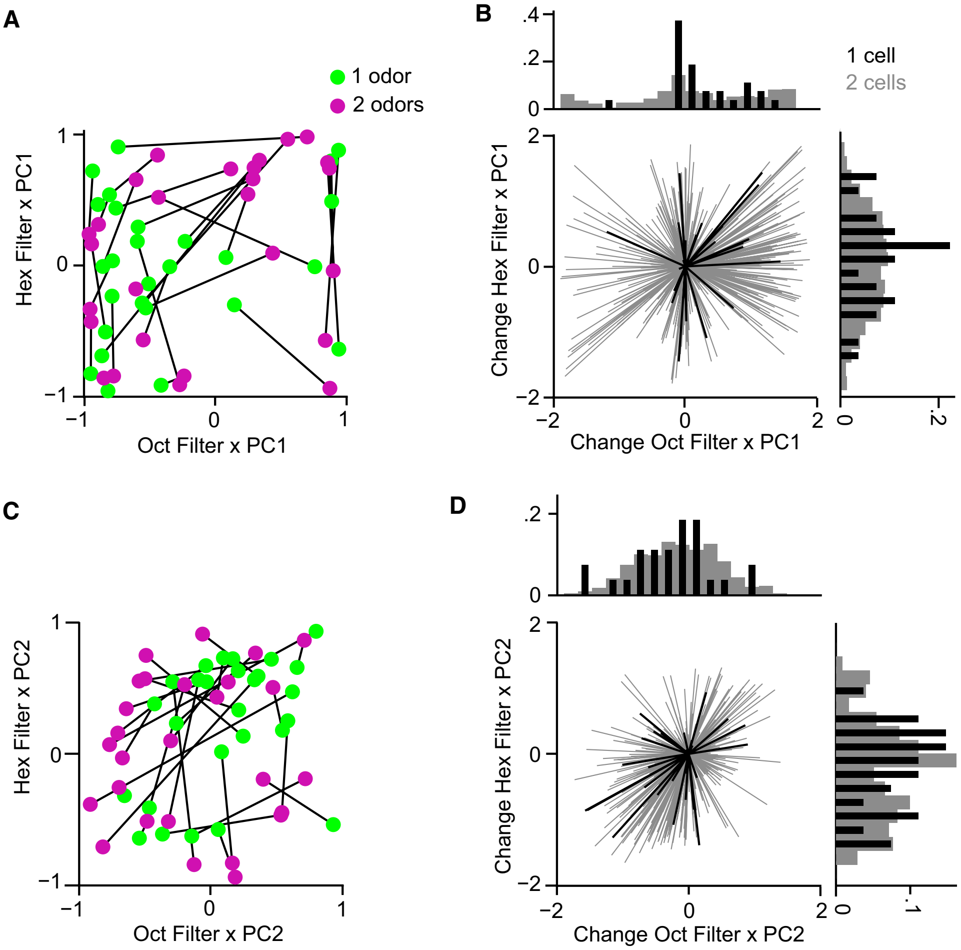

Figure 6. Response Dynamics to One Odor Can Change in Presence of a Second Odor

An analysis of the difference in a neuron's filters between the single-odor and two-odor conditions. Based on 27 neurons with reliable responses in the Oct, Hex, and Oct/Hex conditions (Figure 3).

(A) In this graph, each PN contributes two points, joined by a line. The "1 odor" point reflects the shape of the two filters in the single-odor conditions Oct and Hex, each projected onto the first principal axis (PC1 in Figure 4A). The "2 odors" point reflects the shape of the two filters in the two-odor condition Oct/Hex. The difference between the two points represents the change in filter shapes along the PC1 axis introduced by the two-odor condition.

(B) The filter change from (A), plotted for each PN as difference vectors from the origin (black). For comparison, the background (gray) shows the difference vectors obtained by comparing the single-odor filters (Oct and Hex) of two different neurons. Note the changes introduced by the two-odor condition are similar in magnitude to the diversity of filter shapes across cells. Marginal histograms show the distribution of changes for octanone and hexanol filters respectively. The distribution for the octanone filter is shifted significantly to positive values (t test, $p = 0.04$). This means that in the Oct/Hex condition, the octanone filter acquires a greater component along PC1 (Figure 4A), and thus becomes more On-like.

(C and D) Same analysis, but projecting the filters onto the second principal component PC2 (Figure 4A), presented as in (A) and (B). Again, the shape changes are large, but there is no systematic bias in one direction.

We already encountered an instance where the filter function for one odor changed substantially in the two-odor condition (Figure 2H). Again, the quantitative description of filter shapes (Figure 4) allows a more careful analysis across the population. This was restricted to the limited set of neurons for which high-quality filter measurements were available under both single odors and under the two-odor condition. In this group, many PNs exhibited large changes in response dynamics under the two-odor condition, even involving a change in sign of the temporal filter (Figure 6). Indeed, the changes in filter shapes were comparable in magnitude to the full range of filter shapes across the entire population.

These shifts in filter dynamics point to the relevance of nonlinear interplay between different signaling pathways within the antennal lobe. Interestingly, the changes observed with this stimulus set were not entirely random. For example, the filter for octanone very often became more "On-type" in the presence of hexanol (Figure 6B). The hexanol filter showed a similar trend, though the magnitude was not significant. Clearly, future study of these interactions will benefit from larger sets of PNs and odors.

Dynamics of the PN Population Vector

The above results lead to a simple but quantitative picture of neural coding by the population of PNs in the antennal lobe

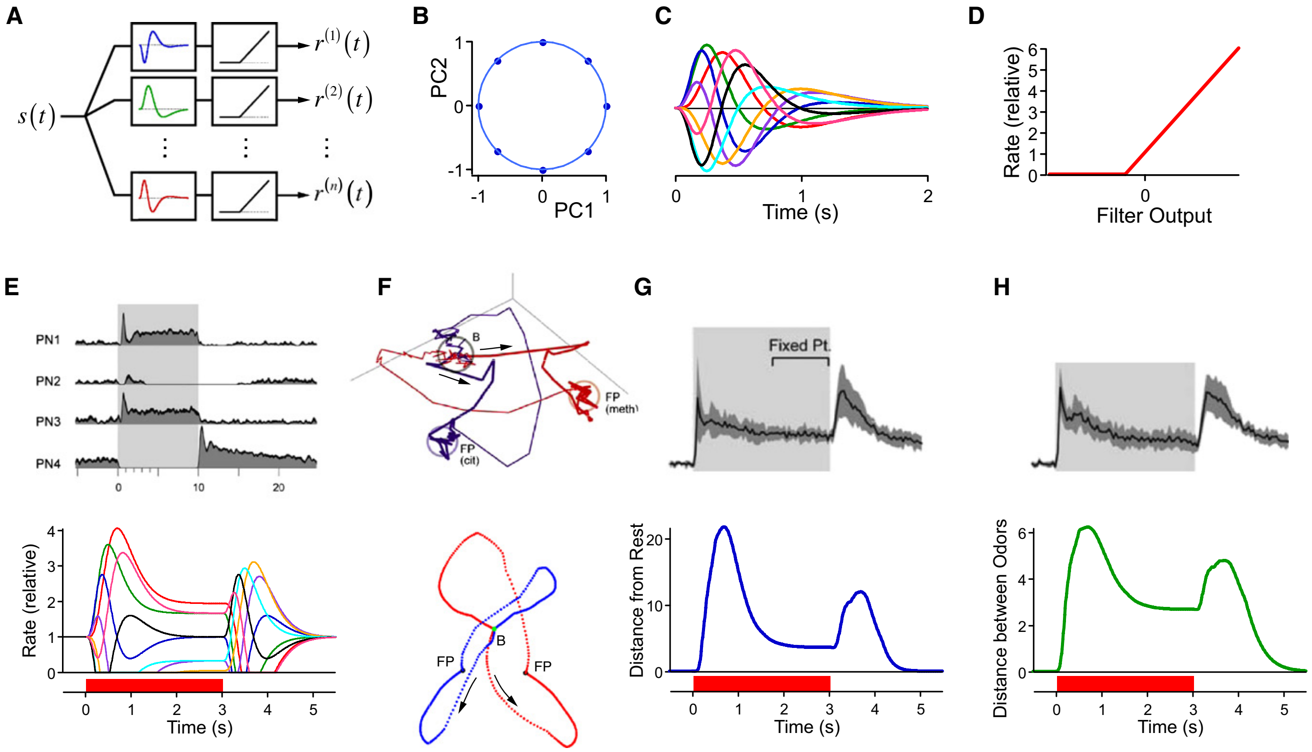

Figure 7. A Population of Linear-Nonlinear Model Cells Reproduces the Population Dynamics among PNs

(A) Population coding by a bank of L-N models. Each neuron is characterized by a filter and nonlinearity. In this way, the common stimulus $s(t)$ is converted to a population of firing rates $r^{(j)}(t)$.

(B–D) Parameters for a model of 8 L-N cells. (B) The linear filters were drawn from a circle in a shape space analogous to Figure 4B (Equation 11). (C) The resulting 8 filter waveforms (Equation 10). (D) The nonlinearity chosen for all 8 neurons: a rectifier with some maintained firing in absence of odor.

(E–H) Population activity of PNs under square odor pulses. Top: Experimental results from Mazor and Laurent (2005). Bottom: Results from the model in (A)–(D), see Equations 10–15. (E) Sample firing rate of several PNs under a square odor pulse. (F) Trajectory of the population vector in response to two different odors. The trajectories were projected into the subspace of the first three principal components and are viewed from a suitable angle. At the onset of the response (arrows), the vector rapidly leaves the baseline (B) and makes a large excursion before reaching the fixed point (FP). At the offset it rapidly leaves the fixed point then makes a large excursion back to baseline. The two odors make different trajectories (red versus blue). In computing the model response to the second odor, the filter functions in (B) were reassigned randomly to the 8 neurons. (G) Length of the population vector in response to a 3 s odor pulse. (H) Distance between the population vectors during responses to the two different odors.

(Figure 7A). A common dynamic odor stimulus is processed in parallel by a bank of L-N pathways, each corresponding to one PN. Across the PN population, the filters and nonlinearities follow a statistical distribution as defined in Figure 4. Note, of course, that this is a phenomenological model, not a mechanistic one; in particular, we do not suggest that each PN has a private set of neural connections to the stimulus. Nevertheless, such a dynamic model should be useful in describing or simulating odor coding in the antennal lobe and its consequences for downstream processing. As a test of such utility, we ask whether the model can account for some of the salient phenomena of population coding reported in prior work, under experimental conditions different from the ones used here.

Several studies have compiled recordings from many locust PNs that collectively provide a glimpse of the overall antennal lobe output (Mazor and Laurent, 2005; Stopfer et al., 2003; Broome et al., 2006; for mitral cells, see also Friedrich and Laurent, 2001, 2004; Bathellier et al., 2008). In these experiments, the stimulus was generally an odor pulse lasting between 0.5 s and several seconds. To generate predictions from our population model (Figure 7A), we simulated 8 L-N neurons and probed them with square pulse stimuli. Reflecting the diversity of the measured PN responses (Figure 4C), each neuron was assigned a different filter function chosen from a one-dimensional family (Figures 7B and 7C). The nonlinearity took a simple threshold-linear shape (Figure 7D, compare to Figure 4G) and was identical for all 8 model PNs.

The responses of these model neurons to a 3 s long odor pulse (Equation 12) bore good resemblance to the range of firing rate dynamics observed experimentally (Figure 7E): the neurons had some resting firing rate. Some cells were excited by the odor, others suppressed. Some had strong transients at the start, others at the end. The dynamics of this population response are sometimes summarized by the "population vector," which is simply the list of all firing rates from projection neurons (Equation 13), measured over successive time bins. The output of the antennal lobe is then fully specified by the trajectory that this vector executes in the course of an odor response.

Some interesting phenomena have been reported about these trajectories:

1. Before the pulse, the population vector hovers at a resting point given by the baseline firing rates of all the PNs (Figure 7F; Mazor and Laurent, 2005). At the onset of the odor pulse, the population vector rapidly moves from the resting point through an extended trajectory. If the pulse is long enough, the trajectory settles at an odor-specific fixed point. When the pulse terminates, the population vector returns from the fixed point back to baseline via a different trajectory. Different odors evoke different trajectories for the population vector. All these phenomena are reproduced faithfully by the population vector from the 8 model neurons (Figure 7F).
2. Following such a trajectory, the distance of the population vector from the baseline is typically much larger during the on- and off-transients than at the fixed point (Figure 7G; Mazor and Laurent, 2005). This implies that odor detection is most sensitive at the transients. Again, the L-N population model reproduces this effect faithfully (Figure 7G).
3. Comparing the population vectors for two different odors, their distance is zero at rest, nonzero at the fixed points, but considerably greater during the on- and off-transients. This implies that the two population vectors are more easily discriminated during the transients, and this enhancement of odor differences has been proposed as a principal function of antennal lobe dynamics (Friedrich and Stopfer, 2001; Laurent, 2002; Mazor and Laurent, 2005). Again, the L-N population model has this same property (Figure 7H).

In summary, much of the phenomenology reported previously for the dynamics of the population vector can be reproduced if one combines PNs described individually by a simple 3-number response model (Figure 7A). The important ingredients for getting the observed population effects are (1) the filters for odor integration are generally biphasic, (2) across neurons, the filters span a broad and more or less continuous range of shapes, (3) across odors, the filter shapes for each neuron change, spanning a similarly broad range of shapes. Note that these successful simulations required no custom-tweaking of the model parameters. Indeed, we obtained similar results by performing the simulation using the actual L-N fits from 20 neurons recorded in this study (not shown).

Small Groups of PNs Encode the Odor Wave Form Accurately

So far, the analysis has been concerned with a "forward" description of antennal lobe processing: given the stimulus, what can one say about the response? A complementary approach aims at a "reverse" description: given the spike trains of PN neurons, what can be said about the stimulus? In a sense, this formulation comes closer to the needs of the animal. To explore this, we tried to reconstruct the stimulus waveform using only the spike trains of projection neurons. Then we compared this reconstruction to the actual stimulus. Note we do not presume that the locust brain actually implements such an explicit stimulus reconstruction. Instead, the analysis serves to determine what information about the stimulus is transmitted by the spikes of projection neurons to the rest of the locust brain.

The simplest approach to stimulus reconstruction is "linear decoding" (Borst and Theunissen, 1999; Warland et al., 1997). Here, each action potential triggers a short waveform of stereotyped shape—the "kernel" (Figure 8A). The running estimate of the stimulus is simply the sum of all these waveforms (see Experimental Procedures, Equation 16). When multiple neurons are used for decoding, they each have a different kernel. Given the observed spike trains and a record of the true stimulus, one adjusts all the kernels to obtain the most faithful reconstruction. Then the quality of this decoding can be evaluated on a separate stretch of data.

For the odor octanone and the PN illustrated in Figures 8A and 8B, the optimal decoding kernel was a biphasic negative pulse extending back in time to ~0.8 s before the spike (Figure 8A): An action potential from this PN "votes" that the octanone valve was closed prior to the spike. Note this kernel differs substantially from the filter obtained during forward analysis (cf. Figure 2C for the same PN), a consequence of the nonlinear stimulus-response relationship (Borst and Theunissen, 1999). The convolution of this kernel with the spikes from a single trial yields a continuous estimate of the stimulus (Figure 8B). One can take the sign of this function to obtain a binary square wave function that estimates when the valve is open (Figure 8B). Remarkably, this estimate—obtained from a single recording of a single PN—was correct 84% of the time. This was the highest value observed among single PNs. The experiments with hexanol yielded a similar reconstruction quality of 82%, but here a different PN was most useful for the stimulus estimate.

As expected, combining spike trains from multiple PNs improved the reconstruction quality (Figures 8C and 8D), but only gradually: consulting the 5 most useful PNs, the state of the valve could be estimated correctly 95% of the time for both octanone and hexanol. Ten neurons provided little additional performance. Interestingly, stimulus reconstruction was quite successful even in the two-odor experiments, where both valve states need to be estimated simultaneously (Figure 8D). Here each PN gives separate "votes" as to the state of the octanone and the hexanol valve (Experimental Procedures). The best single neuron estimated the two valves correctly at 70% and 79%; a group of five neurons gave accuracies up to 86% and 85%. Ten neurons increased the reconstruction quality to 90% for both odors.

The binary truncation of the stimulus reconstruction misses some interesting aspects of neural coding. Note in particular that the stimulus estimate in Figure 8B gets closer to the real valve state during long closings than during short closings. Essentially, the sign of this function is the best guess about the state of the valve, but its magnitude reflects the level of certainty in that estimate. To assess how the decoder's information varies with the rapidity of the valve switches, we computed power spectra of the stimulus, the stimulus estimate, and the decoding error (Figure 8E). From these, one can obtain the spectrum of the information rate transmitted by the spike trains about the stimulus (Experimental Procedures; Figure 8F). This allows some interesting observations:

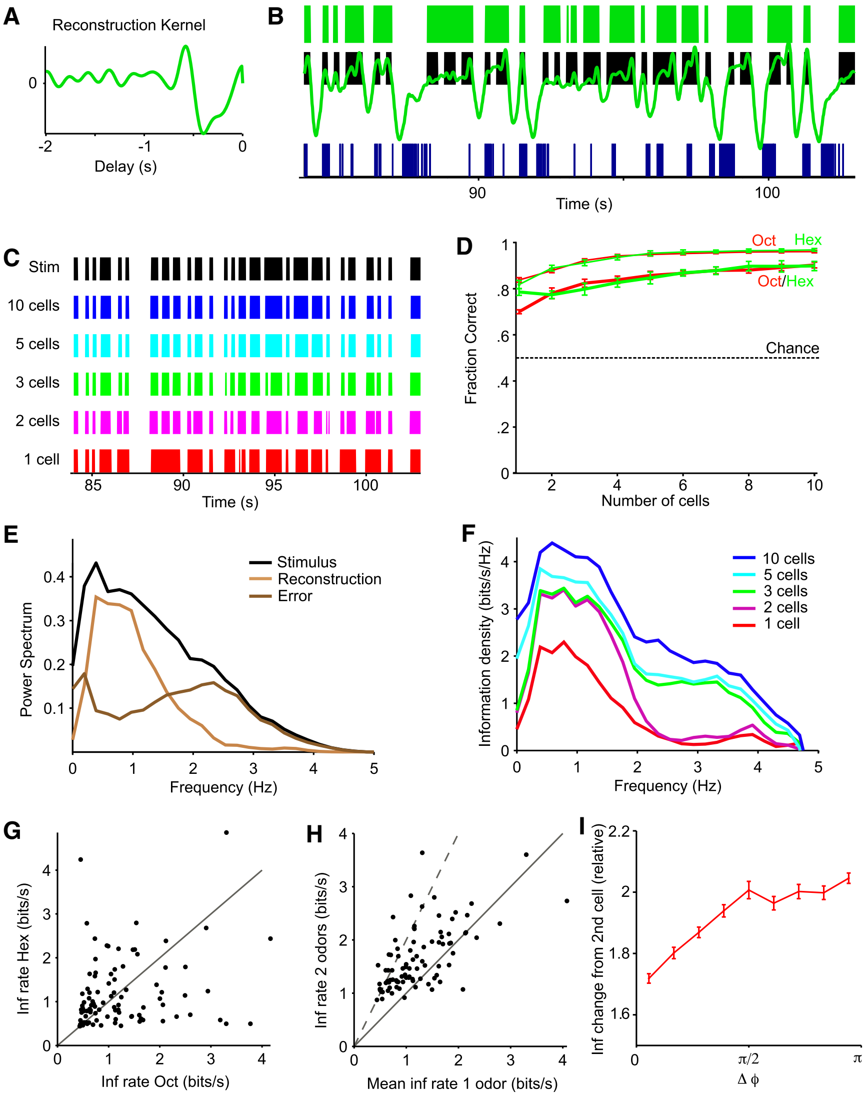

Figure 8. Small Groups of PNs Encode the Dynamic Odor Stimulus with High Accuracy

(A and B) Reconstruction of the odor stimulus from the spike train of a single PN on a single trial. (A) The reconstruction kernel $K(t)$ used for convolution with the spike train (Equation 16). (B) The spike train $r(t)$ (bottom), the convolution $s'(t)$ with the kernel from (A) (middle trace), the binary estimate of the valve state $s''(t)$ (middle bars, Equation 18), and the true valve state $s(t)$ (top bars). Condition: Oct.

(C) Reconstructions of the odor sequence from single trial responses of multiple PNs. For 1 cell, we chose the PN with the best reconstruction (A). Then we serially added to the group the PN that improved the reconstruction the most. Condition: Oct.

(D) Reconstruction quality, measured by the fraction of correct estimates of the odor valve, as a function of the number of contributing cells. In the Oct/Hex condition, results are given for each valve. As in (C), we started with the PN that gave the best one cell reconstruction, then serially added the PN that improved the reconstruction the most. Bars show mean ± SEM across 10 trials. Dashed line indicates chance performance of 50%.

(E) Power spectra of the stimulus, the reconstruction, and the reconstruction error (Equation 19), for the one cell decoding of (B), averaged over 10 trials. Note the reconstruction error is large at both low and high odor flicker frequencies.

(F) The information rate transmitted about the stimulus by small groups of PNs. The curves show the distribution across flicker frequencies (Equation 20), and the total information rate, measured in bits/s, is the area under the curve. Condition: Oct.

(G) Information rate obtained from single PNs in the Hex versus the Oct conditions. Each point represents one PN. All 97 PNs are included. Solid line: identity.

(H) Information rate obtained from single PNs in the Oct/Hex condition (abscissa) compared to the mean of the information obtained in the Oct and Hex conditions (ordinate). Note most PNs convey more information about the combination of both odors than they do on average about each odor alone (points above solid line, slope 1). Some PNs convey more information about two odors than in both single-odor conditions combined (points above dashed line, slope 2).

(I) The information rate obtained from combining two PNs plotted against the similarity of their response dynamics. The ordinate plots the two-neuron information rate relative to the average of the individual rates. The abscissa plots the difference of the two values for the angle ϕ (Figure 4B) that characterizes each neuron's response filter in the L-N model. Note the benefit of combining two PNs is smallest when their filters are most similar ($\phi \approx 0$). Condition: Oct. Analysis based on all pairs among 48 PNs with reliable responses (Figure 3A). Bars are SEM across pairs of cells.

First, single neurons transmitted information about either odor at rates up to ~4 bits/s (Figure 8G). In two-odor experiments, where a neuron can convey information about the state of both odor valves, the information rates were comparable. Remarkably, some neurons conveyed more information about the stimulus in the two-odor conditions than in both one-odor conditions combined. The information conveyed by each spike ranged up to 1.5 bits, with an average across cells of 0.5 bits in the one-odor and 0.7 bits in the two-odor condition. These values are remarkably similar to the information per spike encountered in visual, auditory, mechanosensory, and electrosensory systems (Borst and Theunissen, 1999; Warland et al., 1997), despite the dramatic differences in stimulus dynamics. Second, the encoded information was maximal for odor fluctuations around 0.8 Hz, with a substantial falloff at both lower and higher frequencies, even though the stimulus covered a much broader spectrum (Figure 8E). Third, the information increased when multiple neurons were included, and again there was a marked saturation after the best five cells (Figure 8F).

Finally, the benefit from combining multiple neurons depended considerably on which neurons were chosen. For example, in Figure 8F, the second neuron added information about low flicker frequencies, the third neuron about high frequencies. One might expect this from the analysis of Figure 4, which showed that the "preferred features" of PNs vary a great deal, as summarized by the parameter ϕ. Presumably two neurons that encode the same stimulus feature would be redundant, whereas those with different preferred features would provide complementary information. Indeed, we confirmed that combining a pair of PNs with maximally different filter shapes ($\Delta\phi = \pi/2$) provided a great enhancement of the information rate, approximately the sum of the individual rates (Figure 8H), whereas similar filter shapes ($\Delta\phi \approx 0$) provided less benefit. Interestingly, combining PNs with opposite filter shapes ($\Delta\phi \approx \pi$) was also beneficial. This can be understood based on the rectifying nonlinearity in the response (Figures 4G and 4H): two neurons with filters of opposite sign report different excursions of the flickering stimulus (compare Figures 2C and 2F). Such complementarity between On- and Off-type responses also occurs in the visual system (Warland et al., 1997).

We conclude that a small number of PNs provides sufficient information to decode the time-varying concentration of an odor stimulus with remarkable accuracy, provided one knows what PNs to select. Because the range of dynamic features reported by the PNs is rather limited (Figure 4), a few neurons can cover this range effectively. At the same time, these limitations explain why certain features of the stimulus are poorly represented, specifically components at high and low frequencies.

DISCUSSION

We explored the processing of rapidly varying odor stimuli in the locust antennal lobe. The projection neurons produced complex specific responses, varying in sparseness, timing, and intertrial variability (Figure 1). For many PNs, responses were highly reproducible across trials. The firing rate of individual PNs could be fitted effectively by a simple mathematical model operating on the odor stimulus (Figures 2 and 3). Just 3 numbers were sufficient to specify the dynamics, threshold, and gain of the response (Figure 4). Interestingly, individual PNs responded with different dynamics to different odors: these filters varied equally across PNs for each odor as across odors for each PN (Figures 5 and 6). From this emerged a quantitative picture of dynamic odor coding in the antennal lobe, which generalizes and explains phenomena observed in previous studies (Figure 7). Finally, we considered the task of downstream circuits to extract the odor time course from PN spike trains and found that this can be accomplished with high accuracy relying on just a few well-chosen PNs (Figure 8).

Functional Modeling

The application of complex stimuli necessitates some mathematical formalism to extract features of interest from the ensuing spike trains. In this study, the linear-nonlinear model served this purpose, as it has on many previous occasions (Chichilnisky, 2001; Nagel and Doupe, 2006; Hunter and Korenberg, 1986; Mancini et al., 1990; Stockbridge et al., 1991; Poliakov et al., 1997). We were somewhat surprised by the good performance of this model, especially after the structure was simplified and reduced to just 3 numbers for any given PN-odor combination (Figure 4). This led us to a quantitative and predictive picture of the entire PN population response (Figure 7A). Such a working model of the antennal lobe output should be helpful in thinking about odor coding, for example to evaluate future stimulus designs, to predict population response patterns, to simulate PN spike trains, or to model downstream processing. To demonstrate the model's utility, we showed that it extends beyond the current study and that it accounts—without modification—for many aspects of population dynamics reported previously. The simple mathematical structure (Figure 7A) is an attractive feature: all the calculations for Figure 7 were performed analytically, without numerical simulation.

That said, this working model does not account for every detail of PN activity, and there is room for further developments. In particular, a few neuron-odor combinations were poorly described by the L-N fits, and this was more severe in the two-odor condition (Figure 3C). Note that a single odor generally drives multiple odor receptors. At high concentrations of the ligand, high-affinity receptors saturate and low-affinity receptors are recruited, which can alter the time course of the PN response (Figure 9C). Use of multiple odors exacerbates this diversity of receptor activations at the very input to the antennal lobe. An improved treatment might be obtained from including the binding nonlinearities directly, leading to an N-L-N model structure. Such an approach might also serve to explain the change in filter shape caused by presence of another odor (Figure 6), which falls outside the range of the current model. Another limitation relates to the detailed time course of PN firing: The L-N fit generally predicts the moments during a complex odor sequence when the neuron fires, but the shape of recorded firing events is sometimes sharper and higher than predicted (Figure 2). This same flaw has been noted in other sensory systems, where it has been fixed by adding an explicit spike-generating mechanism (Keat et al., 2001) or a gain control that depends on stimulus history (Victor, 1987; Berry et al., 1999).

Neural Mechanisms of Response Dynamics

While the phenomenological model serves to summarize the response properties of projection neurons, one would also like to understand what neural mechanisms shape those responses. In particular, the dynamics of a PN's odor response—encapsulated by the filter of its L-N fit—must reflect the entire chain of events from odor binding through sensory transduction, synaptic transmission, network processing within the antennal lobe, to postsynaptic events and spike generation in the PN. Here, we consider to what extent the known components and pathways of the antennal lobe can explain the observed response dynamics.

To gauge the role of purely sensory transduction, we recorded the electroantennogram (EAG) from an isolated antenna: this is a field potential derived from the summed transduction currents of many receptor neurons (Vickers et al., 2001; Boeckh et al., 1965). The EAG was subjected to the same L-N analysis as the PN responses recorded under the identical stimulus. The resulting filter function for the EAG was broad and monophasic (Figure 9A). Among the PN filters, the fastest ones had a rising phase and peak time very similar to the EAG. For these PNs, antennal lobe processing seems to entail little delay beyond sensory transduction, and they may well receive direct excitation from the RNs. However, the falling phase of these PN filters was considerably faster than for the EAG sensory response, leading to biphasic shapes with noticeable undershoots. Other PNs in the same experiment had filters of opposite sign or with

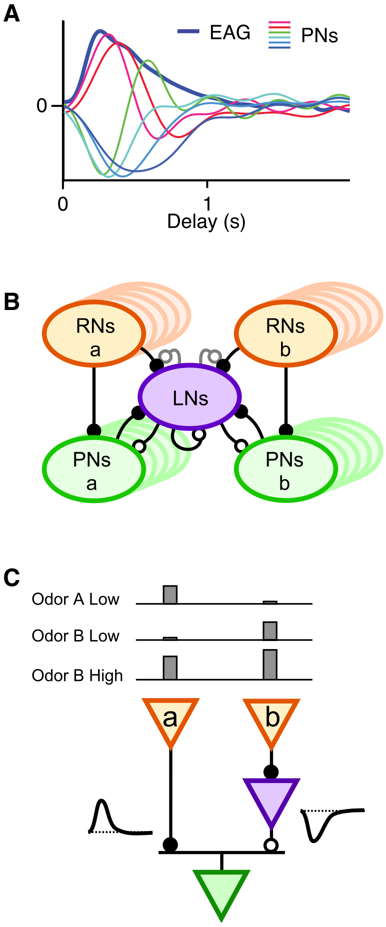

Figure 9. Response Dynamics and Neural Circuits in the Antennal Lobe

(A) The role of sensory transduction. Filter functions obtained for octanone responses of the electroantennogram (EAG, bold line) and of several projection neurons (PNs, thin lines) measured in the same M sequence experiment.

(B) Sketch of the principal circuits among different populations (ovals) of antennal lobe neurons. Receptor neurons (RNs) come in many types, of which only two (a and b) are represented; each type excites a distinct population of projection neurons (PNs). Local neurons (LNs) receive broad excitation from RNs and PNs. They deliver broad inhibition to PNs and to LNs, and possibly to RNs. Synapses are excitatory (closed circles) or inhibitory (open circles). Unconfirmed connections are shaded.

(C) Convergence of pathways with different dynamics. Two of the many pathways by which the odor signal can travel from the RNs to a PN: direct excitation, or inhibition via a local neuron. Next to each pathway is a plausible filter shape for its kinetics. The PN pools these and other inputs, and its overall response kinetics depend on the relative synaptic weight of each pathway. If the RNs are of different types (e.g., RNa prefers odor A, RNb odor B) then they will respond differently to the presented odors (bars). Depending on the odor type (A versus B) and concentration (low versus high) the direct and indirect pathways will be activated in different proportion, altering the response dynamics of the PN.

substantial delays of 0.3–0.4 s relative to the EAG. We conclude from these measurements that the latency and overall time scale of the PN response are determined by sensory transduction. Within that range, the diversity of filter functions among PNs reflects in part a diversity among receptors (Spors et al., 2006; B. Raman et al., 2008, Soc. Neurosci., abstract) as well as the internal dynamics generated by the circuits of the antennal lobe.

Most of this diversity among PN response dynamics was captured by a superposition of just two waveforms (Figures 4A and 4B). This could be explained by the confluence of two major pathways through the antennal lobe. In the fruit fly (Komiyama and Luo, 2006) and probably in other insects (Hansson and Anton, 2000; Jortner et al., 2007), each PN receives direct synaptic input from just one type of RN. Indirect input to the PN, on the other hand, arrives from the pool of local neurons (LNs), which connect promiscuously to many types of RNs and PNs (Figure 9B). A direct path from RN to PN would likely produce an On-type filter (Figure 9C), because RNs are generally excited by odors (reviewed in Buck, 1996; though see Hallem and Carlson, 2006) and are themselves excitatory. An indirect path from RN to PN via a sign-reversing LN would produce an Off-type filter (Figure 9C). Each PN's response will then be a weighted mixture of these two major pathways (Figure 4), depending on the strength of its direct input from the RNs activated by the odor (Figure 9C).

We also found that a given PN can exhibit very different dynamics under different odors (Figure 5; see also Perez-Orive et al., 2002; Broome et al., 2006). Again, this can be understood by the convergence of direct and indirect pathways. A new odor stimulates a different set of RNs (Figure 9C), altering the proportion of direct and indirect input and thus the filter shape. Indeed, this argument explains why the family of filter shapes encountered on changing odors should be the same as the shapes encountered across the population for a given odor (Figures 4 and 5).

Though the basic features of response dynamics in the PN population can be understood from the known neural pathways in the antennal lobe, the details of that circuitry may be rather intricate. For example, in addition to the straight excitatory and indirect inhibitory paths (Figure 9C), one can draw many elaborate signal paths through the diagram of Figure 9B that involve multiple LNs and PNs and feedback loops. These various paths may well interact in very nonlinear fashion, for example via presynaptic inhibition. That could account for some of the more intricate functional observations, specifically that the dynamics of a PN's response to one odor can change in the presence of another odor (Figures 2F–2H and 6). Such effects cannot be explained by weighted summation over parallel pathways (Figure 9C). Since mixtures of odor plumes are common in the environment, it will be useful to pursue the effects of dynamic odor mixing further, and they will likely provide qualitatively different insights about antennal lobe processing. As we have seen, the basic L-N model can serve as a useful working framework that highlights unexpected phenomena.

Decoding of Odor Signals

It seems remarkable that the time course of such a rapidly varying and complex odor waveform can be estimated quite accurately from the spikes of just a few PNs, or even a single one (Figure 8C). This is probably related to our finding that the response dynamics in the PN population are well described by a two-component set of curves (Figure 4). Accordingly two well-driven neurons that represent one component each should cover most of the accessible range of odor fluctuations

(Figure 8F). Of course, the most useful subset of PNs for any such estimation depends on the odor identity and can therefore be determined only once the odor has been identified. For a given olfactory task, downstream circuits might be able to learn the identity of this subset from experience. Such learning would require inspecting many possible subsets of PNs, and assessing their utility based on subsequently assessed value. The mushroom body might play such a role.

The decoding analysis also revealed the limits of what the locust brain can extract from the firing of PNs. These are best appreciated in the frequency domain: Stimulus reconstruction was optimal for odor fluctuations around 0.8 Hz and degraded substantially at both lower and higher frequencies (Figure 8E). The failure at low frequencies means that the brain has little access to the DC level of odor concentration. This is consistent with the biphasic waveform of most PN response filters (Figure 4). The roll-off at high frequencies reflects the temporal averaging of the odor concentration, which is likely dominated by the slow process of sensory transduction. As a result, we predict that locusts should be most sensitive to fluctuations around 0.8 Hz. It could be rewarding to measure these dynamics of odor coding in the antennal lobe for insects with different lifestyles or flight patterns, and explore whether their sensitivity spectrum is adapted to the environmental spectrum of fluctuations, as has been observed in the insect visual system (O'Carroll et al., 1996).

EXPERIMENTAL PROCEDURES

Stimulation

The experiments required reproducible delivery of a rapidly varying odor stimulus to the locust's antenna (Figures 1A–1D). A stream of desiccated and filtered carrier air was blown through a delivery tube aimed at the antenna (Figure 1A). Odorized air was injected into this tube a short distance upstream from the orifice to minimize the mixing time. The flow through the injector line was switched by a fast solenoid valve, and compensated by a change in the carrier flow to keep the overall flow rate constant. To superpose two odors with independent modulation, we added a second injector line and valve. The valves were controlled by a custom computer interface (LabView). To produce the odors, desiccated filtered air was passed through a vial containing 1-hexanol or 2-octanone (Sigma), diluted 1:100 in mineral oil (J.T. Baker). Given the various dilution steps, we estimate the final concentrations of the two odors at 10^{-4}–10^{-3} of the saturated vapor. The concentrations were chosen to yield about equivalent and subsaturating responses in the electroantennogram (EAG). The dynamic variations of odor concentration were measured with a custom electric sensor (see Supplemental Data). A large vacuum hose was placed behind the animal to rapidly remove odors from the antenna.

Recording

Experiments were performed on 8 young (<14 days post-fifth-instar) male locusts (*Schistocerca americana*) that were raised in a breeding colony. The locusts were immobilized in wax cups atop Plexiglas holders with their antennae protruding from the cup through small-diameter Teflon tubing. The brain was exposed by removing the cuticle and sheath while the top half of the head was submerged in locust saline, as previously described (Laurent and Davidowitz, 1994; Stopfer and Laurent, 1999). Silicon tetrodes obtained from the Center for Neural Communication Technology were used to record PN activity (Drake et al., 1988). Silicon probes were placed into the antennal lobe at different sites for each experiment, chosen to maximize detectable activity. Spikes were sorted from the tetrode signals using custom algorithms (Pouzat et al., 2002). All recorded spikes were assigned to PNs, because interneurons of the locust antennal lobe do not produce fast action potentials (Laurent and Davidowitz, 1994).

Stimulus Design

In early experiments, we mapped the general time scale of the PN response. Some PNs responded reliably to valve openings of just 100 ms, but most were somewhat slower. The response typically started several 100 ms after the valve opening and ended after 2 s or less. With these parameters in mind, we designed an appropriate broadband flicker stimulus. The odor valve was switched on and off according to a 9-bit M sequence (Golomb, 1967; Schuckel and French, 2008). This is a pseudorandom sequence of zeros and ones that presents all possible binary words of length 9 exactly once. Each "1" (odor on) or "0" (odor off) lasted 0.2 s, and a 9-bit word spanned 1.8 s. Thus, the stimulus produced all possible odor patterns of length 1.8 s—comparable to the integration time of a PN's response—at a resolution of 200 ms—comparable to the temporal resolution of a PN. These stimulus statistics also approximate the structure of natural odor plumes measured in wind tunnels, which involve the random overlay of concentration transients of about 0.2 s duration (Justus et al., 2002). The experiment involved stimulation first with odor A alone, then odor B alone, then with both odors together but modulated independently. In each case, delivery of the full M sequence required $\sim 2^9 \cdot 0.2\ \mathrm{s} \approx 102\mathrm{s}$. This was repeated for 10 trials to assess reproducibility of the neural responses.

Analysis

Preprocessing

For each cell and each trial, spike times were histogrammed in 0.01 s bins. The result was smoothed by a Gaussian with standard deviation of 3 bins. The firing rates were computed as the average of smoothed spike trains from 10 trials. The waveforms of the valve state and of the electronic nose were also sampled at 0.01 s resolution. In all the analyses reported here, the stimulus variable is the binary state of the solenoid valve; this proved more reliable than the electronic nose signal (see Supplemental Data).

Notation

We will use the following notation:

$$\begin{aligned}
r_i(t) &= \text{firing rate during trial } i\text{, sampled every 0.01 s}\\
r(t) &= \langle r_i(t)\rangle_i = \text{firing rate averaged over all trials}\\
\bar{r} &= \langle r(t)\rangle_t = \text{mean firing rate over all trials and time}\\
s(t) &= \text{valve state} = \begin{cases} -1, \text{off} \\ +1, \text{on} \end{cases}\\
r'(t) &= \text{firing rate predicted from the stimulus by some model}\\
P_S &= \text{``signal power in the response''}\\
&= \text{variance of firing rate over time } = \left\langle (r(t)-\bar{r})^2 \right\rangle_t \qquad (1)\\
P_N &= \text{``noise power in the response''}\\
&= \text{variance of firing rate over trials, averaged over time}\\
&= \left\langle \left\langle (r_i(t)-r(t))^2 \right\rangle_i \right\rangle_t\\
P_R &= \text{``power in the residual of a fit''}\\
&= \text{squared error of the fit, averaged over time}\\
&= \left\langle (r'(t)-r(t))^2 \right\rangle_t
\end{aligned}$$

L-N Model

In the L-N model of the odor response (Figure 2A), the stimulus waveform $s(t)$ is first passed through a linear filter $F(t)$; the output of that filter is then transformed by a nonlinear function $N(g)$ to produce the neuron's estimated time-varying firing rate,

$$r'(t) = N\left(\int_{-\infty}^{t} s(\tau)F(t-\tau)d\tau \right). \qquad (2)$$

Fitting the model for a specific PN involves optimizing the functions $F(t)$ and $N(g)$ so as to minimize the difference between the predicted firing rate $r'(t)$ and the measured rate $r(t)$. Following established practice (Chichilnisky, 2001) we first performed a linear regression to compute the filter $F(t)$ that produced the best linear fit between the stimulus $s(t)$ and the observed response $r(t)$, minimizing the squared error,

$$\min_{F(t)} \int (g(t)-r(t))^2 dt, \text{ where } g(t) = \int s(\tau)F(t-\tau)d\tau. \qquad (3)$$

This deconvolution operation is sensitive to high-frequency noise; we therefore smoothed the filter using a singular value decomposition restricted to the first 20 components. The resulting waveform was truncated to a length of 2 s and normalized to have unit variance,

$$\int_{0}^{2s} (F(t))^2 dt = 1. \tag{4}$$

Then we plotted the observed firing rate $r(t)$ against the filter output $g(t)$ for all time points and averaged both quantities over bins of g containing equal numbers of points. This resulted in a piecewise linear function $N(g)$.

In general, the model parameters were fit to the data over the first 80 s of the odor sequence. The remaining 22 s were used to test the quality of model predictions.

2DL-N Model

We expanded the L-N model to two dimensions (2DL-N model) to fit the firing rate of the neurons under the Oct/Hex condition (Figure 2B). The estimated firing rate of the neuron $r'(t)$ was modeled as a 2-dimensional nonlinear transformation $N(g_a, g_b)$ of the two odor-pulse sequences, $s_a(t)$ and $s_b(t)$, convolved with two distinct linear filters, $F_a(t)$ and $F_b(t)$,

$$r'(t) = N\left(\int s_a(\tau)F_a(t-\tau)d\tau, \int s_b(\tau)F_b(t-\tau)d\tau\right). \tag{5}$$

The filters $F_a(t)$ and $F_b(t)$ were again fitted through reverse-correlation of the response $r(t)$ with the stimuli, $s_a(t)$ and $s_b(t)$, respectively,

$$\begin{aligned} &\min_{F_a(t)} \int (g_a(t) - r(t))^2 dt, \quad \text{where } g_a(t) = \int s_a(\tau)F_a(t-\tau)d\tau \\ &\min_{F_b(t)} \int (g_b(t) - r(t))^2 dt, \quad \text{where } g_b(t) = \int s_b(\tau)F_b(t-\tau)d\tau. \end{aligned} \tag{6}$$

The two-dimensional nonlinearity $N(g_a, g_b)$ was constructed by plotting the observed firing rate $r(t)$ against the filter outputs $g_a(t)$ and $g_b(t)$ for all time points and averaging all three quantities over bins of g_a and g_b containing equal numbers of points.

Linear Filter Analysis

The ensemble of filter functions obtained from the model fits was subjected to a principal component analysis. We computed the eigenvectors and eigenvalues of the crosscorrelation matrix of the filters and picked the two eigenvectors with the highest eigenvalue, $P_1(t)$ and $P_2(t)$ (Figure 4A). In Figures 4B, 5A, and 5D, each dot is the projection of a filter function $F(t)$ onto the two principal components, with coordinates $x = \int F(t)P_1(t)dt$ and $y = \int F(t)P_2(t)dt$. Accordingly, the angle ϕ assigned to a filter is $\phi = \arctan(y/x)$. In Figure 4D, we plot the firing rates $r(t)$ of many PNs, with each cell's rate normalized by its mean value.

In Figure 6A, we plot the projections of the filters $F_a(t)$ and $F_b(t)$ for the same neuron onto $P_1(t)$. Each dot has coordinates $x = \int F_a(t)P_1(t)dt$ and $y = \int F_b(t)P_1(t)dt$. In Figure 6C, we plot the projections onto $P_2(t)$, namely $x = \int F_a(t)P_2(t)dt$ and $y = \int F_b(t)P_2(t)dt$.

3-Parameter Fits

To capture the response model with just 3 parameters, the measured nonlinearity $N(g)$ was approximated as a threshold rectifier determined by two numbers, namely

$$N'(g) = \begin{cases} 0, \text{if } g < a \\ b(g-a), \text{if } g > a \end{cases}, \quad \text{where } a = \text{threshold and } b = \text{gain}. \tag{7}$$

The filter was approximated by a linear combination of the two principal components, determined by the angle ϕ:

$$F'(t) = \cos\phi \cdot P_1(t) + \sin\phi \cdot P_2(t). \tag{8}$$

Then the predicted response was given by

$$r'(t) = N'\left(\int s(\tau)F'(t-\tau)d\tau\right). \tag{9}$$

Figure 4H illustrates the relationship between the threshold a and the gain b (Equation 7). Figure 4I compares the quality of the 3-parameter fit (Equation 9) to that of the unconstrained L-N fit (Equation 2).

Population Model of Pulse Responses

We simulated a model network of 8 PNs whose odor responses $r^{(j)}(t)$, $j = 1, \ldots, 8$ followed the L-N model. Each neuron's linear filter $F^{(j)}(t)$ was chosen as a weighted sum of two principal components

$$F^{(j)}(t) = x^{(j)}P_1(t) + y^{(j)}P_2(t). \tag{10}$$

The shapes of $P_1(t)$ and $P_2(t)$ were inspired by the measured components (Figure 4A), in particular

$$\begin{aligned} &P_1(t) = (t/\tau)^n e^{-t/\tau} - \tfrac{a}{2}(t/2\tau)^n e^{-t/2\tau} \\ &P_2(t) = \frac{d}{dt}P_1(t) \\ &\text{with } \tau = 0.08 \text{ s}, \; n = 5, \; \text{and } a = 0.8 \end{aligned} \quad . \tag{11}$$

The coefficients $(x^{(j)}, y^{(j)})$ were chosen from a unit circle (Figure 7B), yielding the 8 filter shapes $F^{(j)}(t)$ in Figure 7C. The nonlinearity $N(g)$ was identical for all 8 cells: a linear half-wave rectifier with negative threshold (Figure 7D). The odor input was modeled as a square pulse of amplitude A and duration $T = 3$ s, yielding a model firing rate (Figure 7E) of

$$r^{(j)}(t) = N\left(\int s(\tau)F^{(j)}(t-\tau)d\tau\right) = N\left(\int_{\tau=0}^{T} A \cdot F^{(j)}(t-\tau)d\tau\right). \tag{12}$$

The population vector has as its components the firing rates of the individual PNs:

$$\vec{v}(t) = \left[r^{(1)}(t), \ldots, r^{(n)}(t)\right]. \tag{13}$$

The distance of the population vector from baseline (Figure 7F) was computed as

$$|\vec{v}(t) - \vec{v}(0)| = \sqrt{\sum_{j=1}^{n} (r^{(j)}(t) - r^{(j)}(0))^2}. \tag{14}$$

To model the response to a second odor (Figure 7G), the 8 filter shapes $F^{(j)}(t)$ were randomly reassigned to each of the 8 neurons. The distance between two population vectors under the two different odors (Figure 7G) was computed as

$$|\vec{v}_a(t) - \vec{v}_b(t)| = \sqrt{\sum_{j=1}^{n} \left(r_a^{(j)}(t) - r_b^{(j)}(t)\right)^2} \tag{15}$$

where $r_a^{(j)}(t)$ and $r_b^{(j)}(t)$ are the responses of neuron j under the two odors. This distance was further averaged over all possible assignments of the 8 filters to the 8 neurons.

Decoding

For reconstruction of the stimulus and information estimates from spike trains (Figure 8), we followed established procedures (Borst and Theunissen, 1999; Warland et al., 1997), summarized here in brief. To obtain a linear estimate of the stimulus, each spike train was convolved with a kernel function (Figure 8A), and the results were summed to produce the stimulus estimate:

$$s'(t) = \sum_j \int r^{(j)}(\tau)K^{(j)}(t-\tau)d\tau \tag{16}$$

where

$s'(t) = $ stimulus estimate
$r^{(j)}(t) = $ firing rate of cell j
$K^{(j)}(t) = $ decoding kernel of cell j.

The kernels were optimized by linear regression to minimize the squared deviation between the actual and the estimated stimulus, $\int (s'(t) - s(t))^2 dt$. In the two-odor experiments, each cell has two decoding kernels, one for each of the two stimulus estimates:

$$\begin{aligned} s'_a(t) &= \sum_j \int r^{(j)}(\tau)K_a^{(j)}(t-\tau)d\tau \\ s'_b(t) &= \sum_j \int r^{(j)}(\tau)K_b^{(j)}(t-\tau)d\tau. \end{aligned} \tag{17}$$

In computing the decoding filters, we averaged the firing rate $r^{(j)}(t)$ over all repeated trials using 80% of the stimulus sequence. In performing the stimulus reconstruction, we used the spike train from a single trial on the remaining 20% of the sequence. To make a binary guess for the state of the odor valve (Figures 8B and 8C), we took the sign of the stimulus estimate,

$$s''(t) = \mathrm{sign}(s'(t)). \tag{18}$$

The quality of this estimate was measured by the fraction of time it agreed with the actual valve state $s(t)$.

The frequency-dependence of the reconstruction quality was assessed by comparing various power spectra,

$$\begin{aligned} &P_s(f) = \text{power spectrum of the stimulus } s(t) \\ &P_{s'}(f) = \text{power spectrum of the stimulus estimate } s'(t) \\ &P_e(f) = \text{power spectrum of the error } e(t) = s'(t) - s(t). \end{aligned} \tag{19}$$

All spectra were computed by the Welch method with 1024 or 512 point segments.

Given an estimate of the stimulus obtained from the responses, a lower bound on the information rate transmitted by the spike trains (Borst and Theunissen, 1999) is given by

$$I_{\mathrm{LB}} = \int_0^\infty \log_2 \frac{P_s(f)}{P_e(f)}\, df. \tag{20}$$

If the estimate is obtained from linear decoding, then this bound is equal to

$$I_{\mathrm{LB}} = -\int_0^\infty \log_2\left(1 - |C_{rs}(f)|^2\right) df \tag{21}$$

where $C_{rs}(f)$ is the coherence spectrum between the response $r(t)$ and the stimulus $s(t)$. When there are two uncorrelated stimulus variables, as in the two-odor condition, one can show by the same methods that the information rate is greater than

$$I_{\mathrm{LB}} = -\int_0^\infty \log_2\left(1 - |C_{ra}(f)|^2 - |C_{rb}(f)|^2\right) df \tag{22}$$

where $C_{ra}(f)$ is the coherence spectrum between the response $r(t)$ and the stimulus $s_a(t)$, and $C_{rb}(f)$ is the coherence spectrum between the response $r(t)$ and the stimulus $s_b(t)$.

For each cell, the information rate was computed for each individual trial, and then averaged over trials. The information per spike was computed by dividing the information rate by the average firing rate. This was analyzed for all neurons that fired at least 1 spike/s.

SUPPLEMENTAL DATA

The Supplemental Data include three figures and supplemental text and can be found with this article online at http://www.neuron.org/supplemental/S0896-6273(09)00088-9.

ACKNOWLEDGMENTS

The authors thank members of the Laurent and Marcelo Magnasco groups for lively discussion and the group of Nate Lewis for guidance on electronic noses. This work was supported by a HHMI pre-doctoral fellowship (M.N.G.), Burroughs Wellcome Career at the Scientific Interface Award (M.N.G.), a pre-doctoral NRSA (B.M.B.), the National Institute of Deafness and Communication Disorders (G.L.), and the Gordon and Betty Moore Foundation (M.M.).

Accepted: January 21, 2009
Published: February 25, 2009

REFERENCES

Bathellier, B., Buhl, D.L., Accolla, R., and Carleton, A. (2008). Dynamic ensemble odor coding in the mammalian olfactory bulb: sensory information at different timescales. Neuron *57*, 586–598.

Berry, M.J., Brivanlou, I.H., Jordan, T.A., and Meister, M. (1999). Anticipation of moving stimuli by the retina. Nature *398*, 334–338.

Boeckh, J., Kaissling, K.E., and Schneider, D. (1965). Insect olfactory receptors. Cold Spring Harb. Symp. Quant. Biol. *30*, 263–280.

Borst, A., and Theunissen, F.E. (1999). Information theory and neural coding. Nat. Neurosci. *2*, 947–957.

Broome, B.M., Jayaraman, V., and Laurent, G. (2006). Encoding and decoding of overlapping odor sequences. Neuron *51*, 467–482.

Brown, S.L., Joseph, J., and Stopfer, M. (2005). Encoding a temporally structured stimulus with a temporally structured neural representation. Nat. Neurosci. *8*, 1568–1576.

Buck, L.B. (1996). Information coding in the vertebrate olfactory system. Annu. Rev. Neurosci. *19*, 517–544.

Carde, R.T. (1996). Odour plumes and odour-mediated flight in insects. Ciba Found. Symp. *200*, 54–70.

Chichilnisky, E.J. (2001). A simple white noise analysis of neuronal light responses. Network *12*, 199–213.

Drake, K.L., Wise, K.D., Farraye, J., Anderson, D.J., and BeMent, S.L. (1988). Performance of planar multisite microprobes in recording extracellular single-unit intracortical activity. IEEE Trans. Biomed. Eng. *35*, 719–732.

Friedrich, R.W., and Laurent, G. (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. Science *291*, 889–894.

Friedrich, R.W., and Stopfer, M. (2001). Recent dynamics in olfactory population coding. Curr. Opin. Neurobiol. *11*, 468–474.

Friedrich, R.W., and Laurent, G. (2004). Dynamics of olfactory bulb input and output activity during odor stimulation in zebrafish. J. Neurophysiol. *91*, 2658–2669.

Golomb, S.W. (1967). Shift Register Sequences (San Francisco: Holden-Day).

Hallem, E.A., and Carlson, J.R. (2006). Coding of odors by a receptor repertoire. Cell *125*, 143–160.

Hansson, B.S., and Anton, S. (2000). Function and morphology of the antennal lobe: new developments. Annu. Rev. Entomol. *45*, 203–231.

Hopfield, J.J. (1991). Olfactory computation and object perception. Proc. Natl. Acad. Sci. USA *88*, 6462–6466.

Hunter, I.W., and Korenberg, M.J. (1986). The identification of nonlinear biological systems: Wiener and Hammerstein cascade models. Biol. Cybern. *55*, 135–144.

Jortner, R.A., Farivar, S.S., and Laurent, G. (2007). A simple connectivity scheme for sparse coding in an olfactory system. J. Neurosci. *27*, 1659–1669.

Justus, K.A., Murlis, J., Jones, C., and Carde, R.T. (2002). Measurement of odor-plume structure in a wind tunnel using a photoionization detector and a tracer gas. Environ. Fluid Mech. *2*, 115–142.

Keat, J., Reinagel, P., Reid, R.C., and Meister, M. (2001). Predicting every spike: a model for the responses of visual neurons. Neuron *30*, 803–817.

Koehl, M.A., Koseff, J.R., Crimaldi, J.P., McCay, M.G., Cooper, T., Wiley, M.B., and Moore, P.A. (2001). Lobster sniffing: antennule design and hydrodynamic filtering of information in an odor plume. Science *294*, 1948–1951.

Komiyama, T., and Luo, L. (2006). Development of wiring specificity in the olfactory system. Curr. Opin. Neurobiol. *16*, 67–73.

Korsching, S. (2002). Olfactory maps and odor images. Curr. Opin. Neurobiol. *12*, 387–392.

Laurent, G. (1997). Olfactory processing: maps, time and codes. Curr. Opin. Neurobiol. *7*, 547–553.

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

Laurent, G., and Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. Science *265*, 1872–1875.

Laurent, G., Wehr, M., and Davidowitz, H. (1996). Temporal representations of odors in an olfactory network. J. Neurosci. *16*, 3837–3847.

Mancini, M., Madden, B.C., and Emerson, R.C. (1990). White noise analysis of temporal properties in simple receptive fields of cat cortex. Biol. Cybern. *63*, 209–219.

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

Murlis, J., Elkinton, J.S., and Carde, R.T. (1992). Odor plumes and how insects use them. Annu. Rev. Entomol. *37*, 505–532.

Nagel, K.I., and Doupe, A.J. (2006). Temporal processing and adaptation in the songbird auditory forebrain. Neuron *51*, 845–859.

O'Carroll, D.C., Bidwell, N.J., Laughlin, S.B., and Warrant, E.J. (1996). Insect motion detectors matched to visual ecology. Nature *382*, 63–66.

Olsen, S.R., and Wilson, R.I. (2008). Lateral presynaptic inhibition mediates gain control in an olfactory circuit. Nature *452*, 956–960.

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

Poliakov, A.V., Powers, R.K., and Binder, M.D. (1997). Functional identification of the input-output transforms of motoneurones in the rat and cat. J. Physiol. *504*, 401–424.

Pouzat, C., Mazor, O., and Laurent, G. (2002). Using noise signature to optimize spike-sorting and to assess neuronal classification quality. J. Neurosci. Methods *122*, 43–57.

Schuckel, J., and French, A.S. (2008). A digital sequence method of dynamic olfactory characterization. J. Neurosci. Methods *171*, 98–103.

Spors, H., Wachowiak, M., Cohen, L.B., and Friedrich, R.W. (2006). Temporal dynamics and latency patterns of receptor neuron input to the olfactory bulb. J. Neurosci. *26*, 1247–1259.

Stockbridge, L.L., Torkkeli, P.H., and French, A.S. (1991). Intracellular nonlinear frequency response measurements in the cockroach tactile spine neuron. Biol. Cybern. *65*, 181–187.

Stopfer, M., and Laurent, G. (1999). Short-term memory in olfactory network dynamics. Nature *402*, 664–668.

Stopfer, M., Jayaraman, V., and Laurent, G. (2003). Intensity versus identity coding in an olfactory system. Neuron *39*, 991–1004.

Vickers, N.J., Christensen, T.A., Baker, T.C., and Hildebrand, J.G. (2001). Odour-plume dynamics influence the brain's olfactory code. Nature *410*, 466–470.

Victor, J.D. (1987). The dynamics of the cat retinal X cell centre. J. Physiol. *386*, 219–246.

Wachowiak, M., McGann, J.P., Heyward, P.M., Shao, Z., Puche, A.C., and Shipley, M.T. (2005). Inhibition of olfactory receptor neuron input to olfactory bulb glomeruli mediated by suppression of presynaptic calcium influx. J. Neurophysiol. *94*, 2700–2712.

Warland, D.K., Reinagel, P., and Meister, M. (1997). Decoding visual information from a population of retinal ganglion cells. J. Neurophysiol. *78*, 2336–2350.

Wehr, M., and Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. Nature *384*, 162–166.

Wehr, M., and Laurent, G. (1999). Relationship between afferent and central temporal patterns in the locust olfactory system. J. Neurosci. *19*, 381–390.

Wilson, R.I., and Mainen, Z.F. (2006). Early events in olfactory processing. Annu. Rev. Neurosci. *29*, 163–201.

Neuron, Volume *61*

Supplemental Data

Neural Encoding of Rapidly Fluctuating Odors

Maria N. Geffen, Bede M. Broome, Gilles Laurent, and Markus Meister

Electric Nose Measurements

To measure the dynamics of odor concentration delivered by our apparatus, we constructed a sensor that mimicked the slender shape of the locust antenna. A thin glass capillary (Warner Instrument Corp.) was coated with a film of doped polymer (mixture of carbon black, polyethylene vinyl di-acetate and resin) using an airbrush and established protocols (Matzger, A. J., Lawrence, C. E., Grubbs, R. H., and Lewis, N. S. 2000. Combinatorial approaches to the synthesis of vapor detector arrays for use in an electronic nose. J Comb Chem 2, 301-304). Electrical contacts were made on either side of a 3 mm exposed patch of polymer by attaching copper wires to the glass with silver paint. At baseline, the electrical resistance ranged between 50 kOhms and 500 kOhms. Exposure to vapor increased the resistance. Recordings with this sensor showed that the odor concentration can be rapidly switched on and off every 100 ms. The timing of successive odor pulses was highly reproducible: the delay between valve opening and half-maximal rise of the odor pulse varied by only 0.7 ms (SD). This indicates that the odor can indeed be modulated rapidly and reliably, and over a wide range of concentration (Fig 1E).

In all our reported analyses, the stimulus is the binary state of the solenoid valve in the odor injector line. We considered using the signal from the electric odor sensor, but rejected it for the following reasons. First, the detailed waveform reported by these electric sensors varied considerably among different devices. Second, when we did perform some L-N model fits of neural responses starting from the sensor output, they were less successful than L-N fits starting from the valve signal; they required more convoluted and acausal filter waveforms; and they depended on the device used. These observations suggest that the polymer sensor introduced significant dynamics and distortion in the reported signal. While the sensors served to confirm that odor modulation is rapid and reproducible across trials, they need to be engineered further to provide a quantitative record of odor concentration. In evaluating the reported filter shapes $F(t)$, one therefore needs to recall that these include a brief delay of <0.1 s from valve opening to odor onset. This air-flow delay is short compared to the neural response delays, and it is identical for all neurons in the sample.

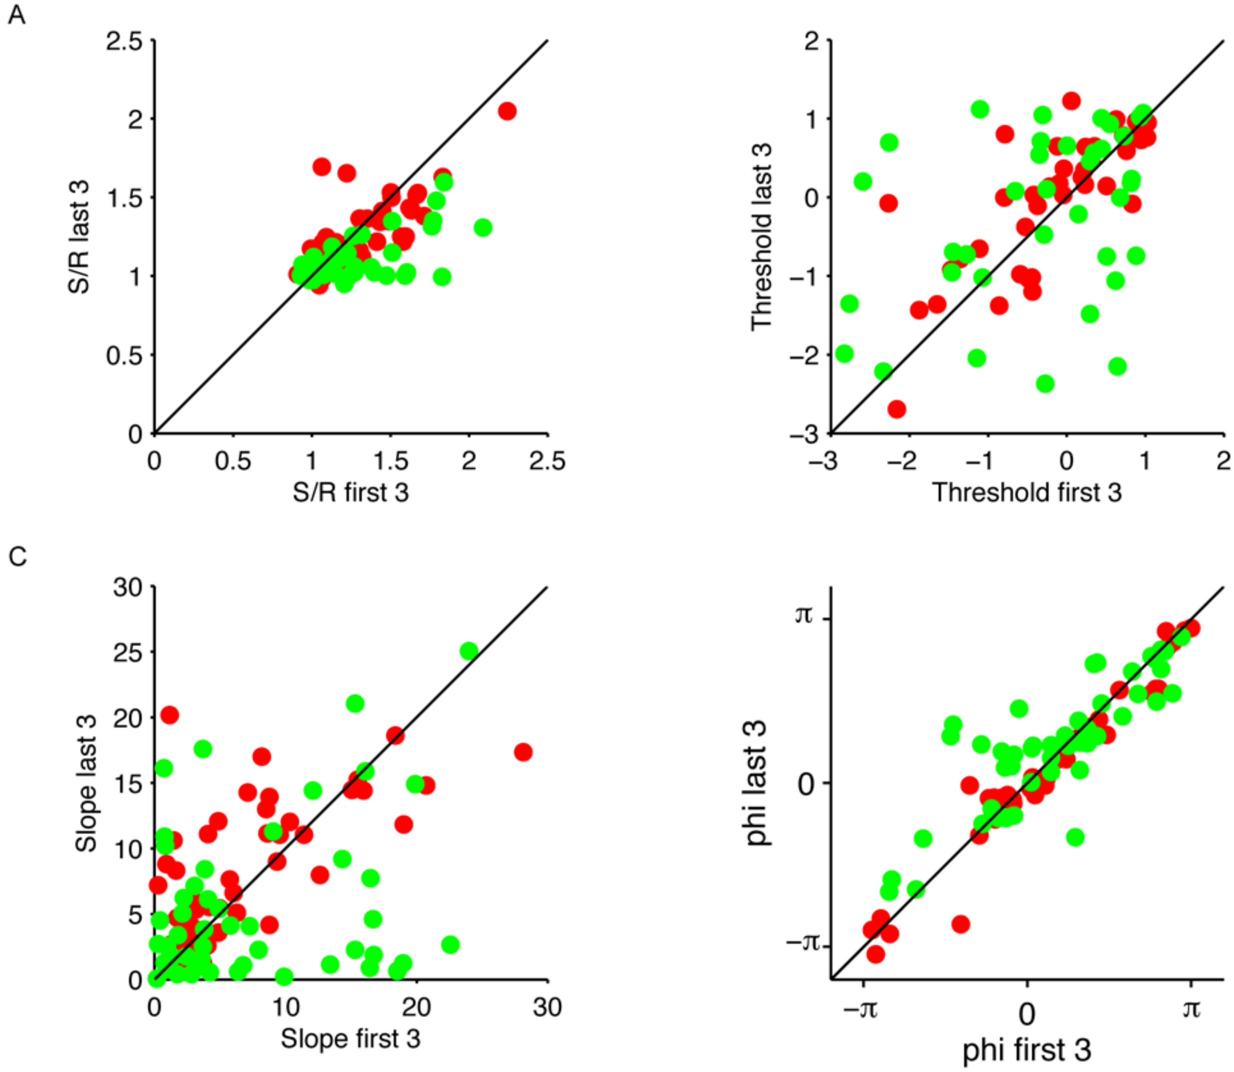

Figure S1. Change in L-N Fits across Trials

Some PNs showed a systematic change in firing rate over the ten trials (Fig 1G). Here we evaluate what aspects of the odor response change during this period. The 3-parameter L-N model (Eqns 7-9) was fit to the first 3 trials and again to the last 3 trials. Analysis included 48 PNs each with reliable responses in the Oct, Hex, and Oct/Hex conditions (Fig 3).

A. The ratio Signal/Residual ($\sqrt{P_S/P_R}$, see Experimental Procedures) compared in early and late trials. Line is the identity. By this measure, the goodness-of-fit is comparable in early and late trials.

B. The threshold of the nonlinearity (Eqn 7) compared in early and late trials. On average, thresholds increase somewhat in late trials.

C. The slope of the nonlinearity (Eqn 7) compared in early and late trials. This parameter reflects the gain of the response and changes considerably between early and late trials, increasing for some PNs and decreasing for others.

D. The shape of the filter compared in early and late trials, as summarized by the shape parameter ϕ (Eqn 8). Note that the filter shape remains essentially the same throughout all trials.

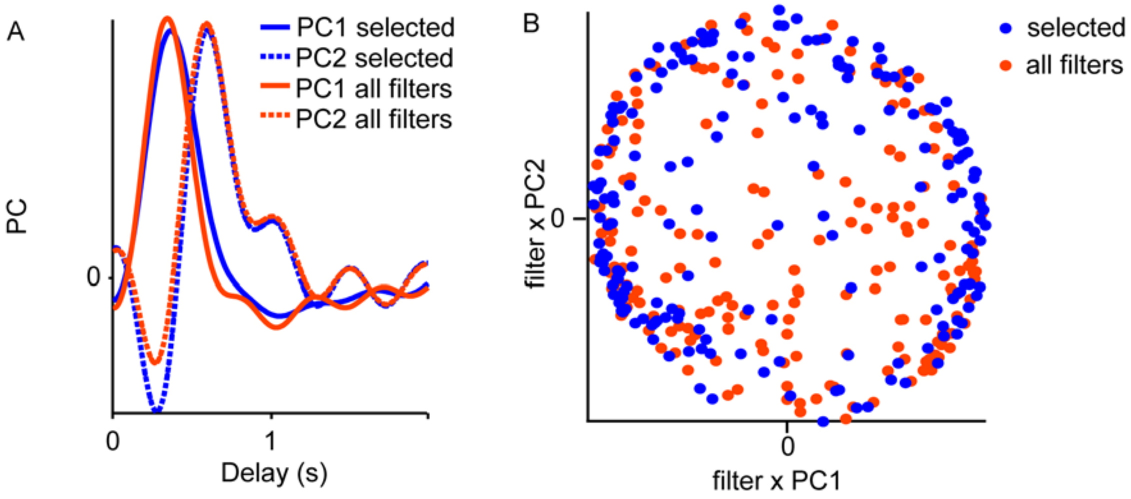

Figure S2. Principal Component Analysis of All Filter Waveforms

The analysis of Figure 4 was restricted to the PNs with the most reliable odor responses, selected as the top 50% in terms of Signal/Noise ($\sqrt{P_S/P_N}$, see Fig 3). Here we test whether the excluded cell/odor combinations with weaker responses had some systematically different kinetics, by repeating the analysis of filter shapes over all neurons. We found no evidence for additional kinetic components.

A. The first two principal components of the filter shapes, computed either for the selected 48 strong responses (same as Fig 4A) or for all 97 responses. No systematic difference is discernable.

B. The filter shapes projected onto the first two principal components, PC1 and PC2 from panel A. The selected filters for strong responses lie close to the unit circle (same as Fig 4B), because these two components account for most of their variance. The excluded filters for weak responses fill more of the space within the circle, because they are corrupted by noise with random waveform other than PC1 and PC2. The fact that PC1 and PC2 are virtually identical to those of the strong responses (panel A) confirms that this noise has no systematic component.

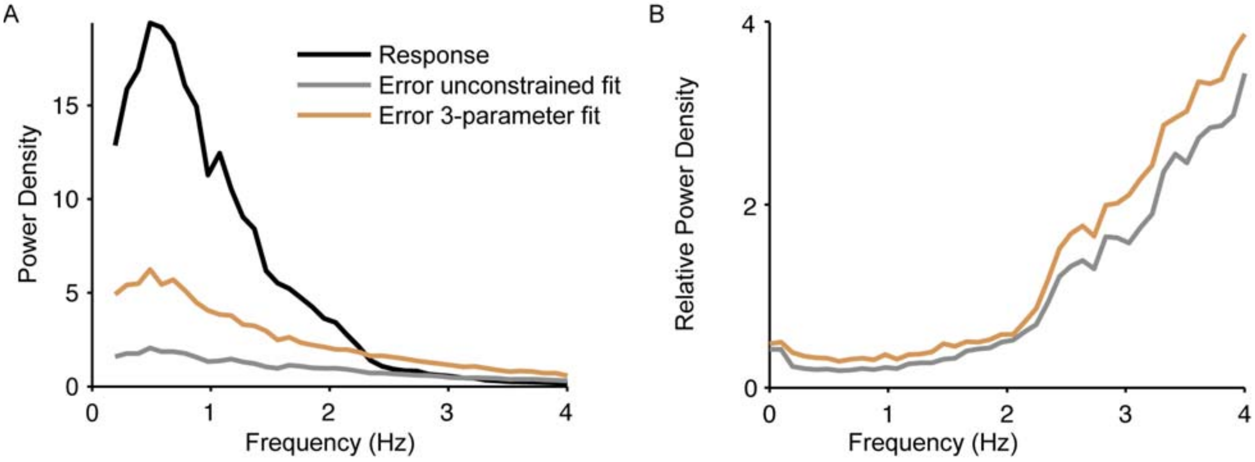

Figure S3. Frequency-Dependence of the 3-Parameter L-N Model Fits

For most cell-odor combinations, the 3-parameter L-N model fits almost as well as the unconstrained L-N model (Fig 4I). On average there is a small loss in fit quality, more so for some cells in the Oct condition. Here we explore the nature of these losses. In particular, one might worry that the 3-parameter model fails at high frequencies, because the waveform of the filter function was approximated by just 2 principal components with rather smooth waveform (Fig 4A). This concern was not validated.

A. Power spectrum of the response $r(t)$ (see Eqn 1), the residual $r'(t) - r(t)$ from the unconstrained L-N fit, and the residual from the 3-parameter fit. Results were averaged over all 48 selected cells (Fig 3) in the Oct condition.

B. Power of the two residuals plotted as a fraction of the power in the response. Note the two residuals have very similar frequency dependence.