

Functional Roles for Synaptic-Depression within a Model of the Fly Antennal Lobe

Aaditya V. Rangan*

Courant Institute of Mathematical Sciences, New York University, New York, New York, United States of America

Abstract

Several experiments indicate that there exists substantial synaptic-depression at the synapses between olfactory receptor neurons (ORNs) and neurons within the drosophila antenna lobe (AL). This synaptic-depression may be partly caused by vesicle-depletion, and partly caused by presynaptic-inhibition due to the activity of inhibitory local neurons within the AL. While it has been proposed that this synaptic-depression contributes to the nonlinear relationship between ORN and projection neuron (PN) firing-rates, the precise functional role of synaptic-depression at the ORN synapses is not yet fully understood. In this paper we propose two hypotheses linking the information-coding properties of the fly AL with the network mechanisms responsible for ORN→AL synaptic-depression. Our first hypothesis is related to variance coding of ORN firing-rate information — once stimulation to the ORNs is sufficiently high to saturate glomerular responses, further stimulation of the ORNs increases the regularity of PN spiking activity while maintaining PN firing-rates. The second hypothesis proposes a tradeoff between spike-time reliability and coding-capacity governed by the relative contribution of vesicle-depletion and presynaptic-inhibition to ORN→AL synaptic-depression. Synaptic-depression caused primarily by vesicle-depletion will give rise to a very reliable system, whereas an equivalent amount of synaptic-depression caused primarily by presynaptic-inhibition will give rise to a less reliable system that is more sensitive to small shifts in odor stimulation. These two hypotheses are substantiated by several small analyzable toy models of the fly AL, as well as a more physiologically realistic large-scale computational model of the fly AL involving 5 glomerular channels.

Citation: Rangan AV (2012) Functional Roles for Synaptic-Depression within a Model of the Fly Antennal Lobe. *PLoS Comput Biol* 8(8): e1002622. doi:10.1371/journal.pcbi.1002622

Editor: Peter E. Latham, Gatsby Computational Neuroscience Unit, University College London, United Kingdom

Received: January 6, 2012; **Accepted:** June 11, 2012; **Published:** August 23, 2012

Copyright: © 2012 Aaditya V. Rangan. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This study was funded by the NSF. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The author has declared that no competing interests exist.

* E-mail: rangan@cims.nyu.edu

Introduction

The early stages of the drosophila olfactory system include a primary sensory structure called the antenna lobe (AL). The AL receives input from olfactory sensory neurons (ORNs) at the sensory periphery, and is organized into glomerular clusters, with each cluster corresponding to a specific olfactory receptor class [1–5]. Each glomerulus within the AL contains dendrites of local neurons (LNs) whose projections are limited to the AL, as well as projection neurons (PNs) whose axons extend beyond the AL deeper into the fly brain [6]. The PNs are excitatory, whereas there is evidence that both excitatory local neurons (LNEs) and inhibitory local neurons (LNIs) exist [7–9]. The LNs associated with each glomerulus have local projections, which connect to that glomerulus, as well as lateral projections which connect to other glomeruli [10].

Various experiments indicate that there exists substantial synaptic-depression at the synapses between olfactory receptor neurons (ORNs) and neurons within the drosophila antenna lobe (AL); by ‘synaptic-depression’, we refer to any mechanism which gives rise to short-term depression of the ORN-induced EPSCs within the AL following an increase in ORN activity. While it has been proposed that this synaptic-depression contributes to the nonlinear relationship between ORN and PN firing-rates, the precise functional role of synaptic-depression at the ORN synapses is not yet fully understood. To investigate the relationship between

synaptic-depression and the coding properties of the fly AL, we created and analyzed the dynamics of several models of the fly AL. We have been able to distill two hypotheses linking the information-coding properties of the fly AL with the network mechanisms responsible for ORN→AL synaptic-depression.

Our first hypothesis is related to the variance coding of ORN firing-rate information — once stimulation to the ORNs is sufficiently high to saturate PN responses within any particular glomerular channel, further stimulation of the ORNs can reduce the amount of fluctuation of the ORN→PN input within that channel, thus increasing the regularity of PN spiking activity while maintaining PN firing-rates. Thus, given two different stimuli which saturate the responses of a given glomerulus, it may still be possible to distinguish between these two stimuli solely by using this saturated glomerulus’ activity. In order to distinguish these saturated responses, a readout mechanism must be sensitive to higher-order statistics (such as variance) in the saturated glomerulus’ activity.

Our second hypothesis proposes a tradeoff between trial-to-trial reliability and sensitivity governed by the mechanisms responsible for ORN→AL synaptic-depression. Within the fly, synaptic-depression may be partly caused by vesicle-depletion, and partly caused by presynaptic-inhibition due to the activity of inhibitory local neurons within the AL [11,12]. Our second hypothesis is that synaptic-depression caused primarily by vesicle-depletion will give rise to a very reliable system, whereas an equivalent amount of

Author Summary

Understanding the intricacies of sensory processing is a major scientific challenge. In this paper we examine the early stages of the olfactory system of the fruit-fly. Many experiments have revealed a great deal regarding the architecture of this system, including the types of neurons within it, as well as the connections those neurons make amongst one another. In this paper we examine the potential dynamics produced by this neuronal network. Specifically, we construct a computational model of this early olfactory system and study the effects of synaptic-depression within this system. We find that the dynamics and coding properties of this system depend strongly on the strength, and sources of, synaptic-depression. This work has ramifications for understanding the coding properties of other insect olfactory systems, and perhaps even other sensory modalities in other animals.

synaptic-depression caused primarily by presynaptic-inhibition will give rise to a less reliable system that is more sensitive to small shifts in odor stimulation. Using this second hypothesis, one can further postulate that a balance of vesicle-depletion and presynaptic-inhibition within the AL is required in order to optimize the discriminability of the network over short observation-times.

Results

The relationship between the architecture of the fly AL and its odor-coding properties largely remain a mystery. Specifically, the precise functional role of synaptic-depression at the ORN synapses is still unclear. In order to investigate the possible function associated with these network mechanisms, we have designed and built a scaled down computational network model of the fly AL. By analyzing the dynamics of this model we have been able to distill two hypotheses linking the information-coding properties of the fly AL with the network mechanisms responsible for ORN→AL synaptic-depression. We will discuss these hypotheses

later in the sections below, after first introducing a few pertinent details regarding our computational model.

Sketch of computational network model

In brief, our computational network model incorporates 5 glomerular channels, each with 6 PNs, 6 LNEs, 6 LNIs and 60 ORNs, in rough accordance with the experimentally observed ratio of ORNs to PNs and LNIs [13]. As the real fly AL has ~ 50 glomerular compartments, each of roughly this size [10], this model is $\sim 10\%$ the size of the full AL. Each neuron in this network model is modeled using Hodgkin-Huxley-type equations. The synaptic currents in this network allow neurons to affect other neurons in the same glomerulus, as well as neurons in other glomeruli. The input to this network takes the form of noisy stimulus current to the ORNs, with different ‘odors’ corresponding to different levels of stimulus current to different ORN input channels. Importantly, the model is built to accommodate synaptic-depression of the ORN synapses, allowing for both the mechanisms of presynaptic-inhibition as well as vesicle-depletion. An illustration of the network’s connectivity, as well as an abridged list of network parameters, is given in Fig. 1. We have built this network to respect physiological constraints, and we have tuned this model using several experiments as benchmarks. Here we provide a brief summary of these results. A more detailed description of the model as well as the details regarding the benchmarking are contained in the Methods section.

Our goal while benchmarking this model was to ensure that our model produced reasonable statistical features of AL activity during the $\sim 500ms$ following odor onset. The reason we focused on matching the statistics of this transient period is that evidence indicates that this period is likely critical for many basic olfactory discrimination and classification tasks [14,15]. One of the simplifications we have made in our model is that the input to the ORNs following odor onset is assumed to be a Poisson process with a time-varying rate that is roughly stereotyped across ORN classes (see Methods). While natural odor stimuli are likely temporally complex [16] and even static stimuli generate odor-specific temporal fluctuations at the level of the fly ORNs after

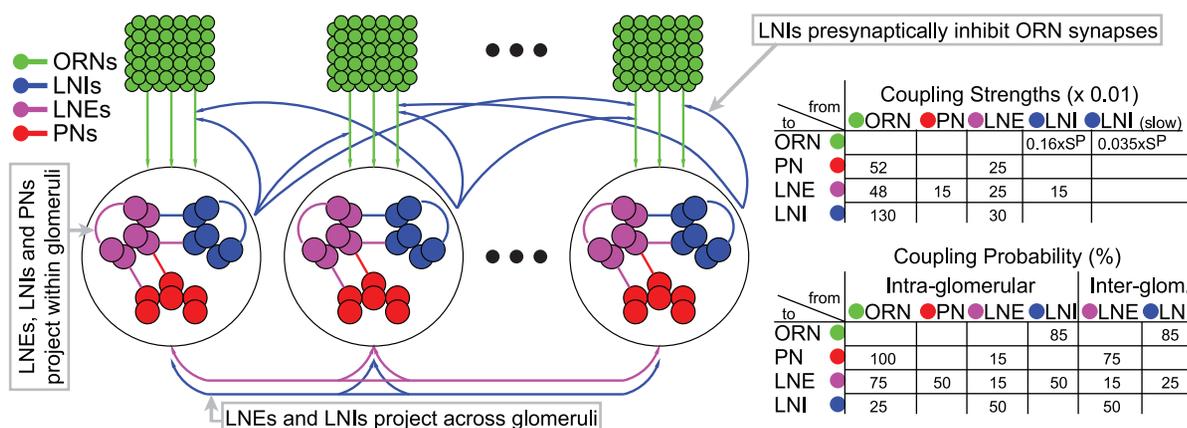


Figure 1. A schematic of the large-scale network model. [Left]: The network consists of 5 glomerular channels, each incorporating 60 olfactory receptor neurons (ORNs in green) which stimulate a ‘glomerulus’ consisting of 6 projection neurons (PNs in red), 6 excitatory local neurons (LNEs in magenta) and 6 inhibitory local neurons (LNIs in blue). The PNs, LNEs and LNIs are connected to one another randomly within each glomerulus, and the LNEs and LNIs also affect the neurons in other glomeruli. The LNIs affect the ORN→AL synapses via presynaptic-inhibition. [Right]: The non-negligible connection strengths are listed on top, with the slow-inhibitory connection strengths listed separately from the fast-inhibition strengths. The relevant connection probabilities are listed on the bottom. The parameter S^P refers to S^{presyn} , which characterizes the overall strength of presynaptic-inhibition. See Methods for full details. doi:10.1371/journal.pcbi.1002622.g001

several hundred ms [17], the dynamics of the ORN responses during the first 500ms following odor onset seems to be relatively stereotypical, involving either a sharp increase in activity or, more rarely, an inhibitory phase [18,19]. Thus, the idealized input to the ORNs we employ in our model is intended to capture these simple features of ORN activity which drive the AL during the first $\sim 500ms$ following odor onset.

The experimental phenomena we used to benchmark our model ultimately provided three constraints on the connectivity of our model network. First, the convergence ratio of ORNs to PNs must be high, otherwise the PNs do not receive sufficient convergent input to fire quickly after odor onset. Second, the synaptic-depression at the ORN synapses must be sufficient to ensure that PN firing-rates peak earlier than ORN firing-rates (in response to odor stimulus), and that ORN \rightarrow PN input is strong and relatively stable during the first 500ms after odor stimulus onset. Finally, the inter-AL connectivity (governed by the LN \rightarrow LN, PN \rightarrow PN, PN \rightarrow LN, and LN \rightarrow PN connection matrix) must be sufficiently strong to create PNs which are more broadly responsive than their ORN inputs, yet sufficiently sparse to place the network in a dynamic regime which does not develop spontaneous oscillations (which are not observed experimentally during the initial transient following odor onset – [17]).

In addition, to further understand the network mechanisms underlying the two proposed hypotheses, we have designed simpler neuronal network models which distill the relevant phenomena, while allowing for a more comprehensive analysis. The analytical tools we use include the analysis of return-maps for simple network models, as well as the analysis of population-dynamics equations for more complicated network models (see the sections to follow for more details).

Hypothesis 1: a monotonically decreasing map between ORN activity and PN input variance

As evidenced in [20,21], the relationship between ORN firing-rate (m_{ORN}) and PN firing-rate (m_{PN}) for a given glomerulus is often nonlinear, with the PN firing-rate saturating rather quickly as a function of ORN firing-rate. One consequence of this nonlinearity is that, for low m_{ORN} , the gain in m_{PN} is high — as m_{ORN} varies from 0–50 Hz, m_{PN} can vary from 0–150 Hz or more. Another consequence of the nonlinear relationship is that, for high m_{ORN} , the gain in m_{PN} is low — as m_{ORN} varies from 100–200 Hz, m_{PN} may remain almost constant. Many have noted that the region of high gain allows for ‘odor separation’ — namely, odors which give rise to similar m_{ORN} profiles for a given glomerulus may in turn produce very different m_{PN} profiles within that glomerulus [20]. However, this ‘odor separation’ only works when the odors in question generate m_{ORN} which are sufficiently low as to lie in the region of high m_{PN} gain. It is tempting to conclude that if two odors generate m_{ORN} which are sufficiently high (such that the induced m_{PN} lie in the region of low gain), then the m_{PN} generated by these odors would be similar, and the odors would not be ‘separated’.

The first hypothesis we propose is that, even if two odors generate m_{ORN} which correspond to similar m_{PN} , the dynamics of the glomerulus may still serve to separate these odors. However, in this case the odor separation takes place not in terms of PN firing-rates (as, indeed, the m_{PN} generated by these two odors may be very similar or identical), but rather in terms of higher-order statistics of PN activity. In other words, even though the set of PN firing-rates produced at the plateau of the $m_{PN}(m_{ORN})$ relationship are similar, we hypothesize that there is in fact a systematic difference in the PN dynamics underlying these similar PN firing-rates.

To be more specific, we claim that for values of m_{ORN} along the plateau of the $m_{PN}(m_{ORN})$ relationship, as m_{ORN} increases (and m_{PN} stays roughly the same), the synaptic-depression at the ORN synapses continues to increase. One consequence of this increase in synaptic-depression is that, as m_{ORN} increases along the plateau of $m_{PN}(m_{ORN})$, the number of ORN firing-events increases, but the effect of each ORN firing-event on postsynaptic PNs decreases. Thus, the postsynaptic conductance induced within any PN by the ORNs (i.e., the ORN input to the PN) maintains roughly the same mean, but decreases in variance. When discussing a reduction in the variance of ORN input, we refer specifically to a reduction in the variance across short time-windows of the PN excitatory-conductance due to ORN activity.

If the m_{ORN} is not very high, then each ORN generates relatively few spikes, each resulting in a large EPSC in the postsynaptic PN. Thus, the ORN input to the PNs will have large fluctuations (i.e., the PNs will be ‘fluctuation-driven’). On the other hand, if m_{ORN} is very high, then each ORN generates very many spikes, each resulting in a small EPSC within the postsynaptic PN. In this case the PN conductance due to the ORNs will be nearly constant (and the PNs will be ‘mean-driven’). We further hypothesize that, as m_{ORN} increases along the plateau of $m_{PN}(m_{ORN})$, the decrease in variance of ORN input to the PNs will correspond to a decrease in the variance of PN spiking activity. Because (i) the ORN activity is not deterministic, but rather driven by many independent stochastic molecular binding events [18], and (ii) many ORNs are presynaptic to each PN, the accumulation of ORN firing-events observed by any given PN during any trial of odor presentation is well-approximated by a Poisson process with time-varying rate. Thus, a decrease in the ORN input variance across short time-windows will be associated with a decrease in the ORN input variance across multiple trials (for the same time-window). Thus, one would expect the variance in PN spiking activity mentioned above to decrease both across short time-windows and across multiple trials (for the same time-window). This reduction in variance of PN spiking activity is equivalent to an increase in the regularity of PN spiking activity, which is equivalent to a reduction in the variance of the inter-spike-interval distribution associated with a PN within the given glomerulus.

Thus, in summary, our first hypothesis is that the dynamics of a glomerulus can serve to separate ORN inputs in two ways. Not only can similar ORN inputs within the high-gain region of $m_{PN}(m_{ORN})$ be mapped to significantly different PN firing-rates (see [20]), but ORN inputs within the low-gain region of $m_{PN}(m_{ORN})$ can give rise to PN activity with differing degrees of regularity, even when the PN firing-rates associated with those ORN inputs are not significantly different. This hypothesis may have significance for odor discrimination, as the variance in PN activity may encode features of the odor even in situations where the ORN input is sufficiently high that PN firing-rates have saturated (see Discussion).

A simple cartoon of variance coding. As a simple cartoon which illustrates this hypothesis, we have simulated a single conductance-based integrate-and-fire PN, driven by a set of 4 ORNs, each endowed with a simple model of synaptic-depression. This simple model exhibits the following dynamical features: (i) the $m_{PN}(m_{ORN})$ relationship exhibits high gain and saturation, and (ii) for different values of m_{ORN} on the plateau of the $m_{PN}(m_{ORN})$ relationship, the variance in PN activity decreases as m_{ORN} increases, even though m_{PN} remains roughly constant.

Within this simple model, we describe each ORN as a Poisson process with fixed rate η ($0Hz \leq \eta \leq 180Hz$). The coupling strength $S^{PN \leftarrow ORN}$ between the ORNs and the PN is modulated by a term $\mu(t)$ ($0 < \mu < 1$), which is intended to model vesicle-depletion at the

ORN synapses. As each ORN fires, this μ term will give rise to synaptic-depression between the ORNs and the PN. If $\mu=1$, the synapses between the ORNs and the PN are 100% exhausted. If $\mu=0$, the synapses between the ORNs and the PN are completely refreshed. The model details are given in a section entitled “An idealized model used to illustrate variance coding” in Methods.

With this simple model, it can be seen that the PN firing-rate $m_{PN}(\eta)$ is a nonlinear function of the ORN firing-rate η , and that m_{PN} saturates (plateaus) at values of $\eta \sim 30\text{Hz}$ (See Fig. 2A). The time-averaged mean total excitatory conductance $\langle G_E \rangle_t$ of the PN enjoys a similar nonlinear relationship (Fig. 2B). Notably, for values of $\eta > 30\text{Hz}$, the time-averaged mean vesicle-depletion parameter $\langle \mu \rangle_{t,ORN}$ increases as a function of η , and the standard deviation in the total PN conductance $\sigma_{G_E} = \sqrt{\langle G_E^2 \rangle_t - \langle G_E \rangle_t^2}$ decreases as a function of η (Fig. 2C and Fig. 2D). This decrease in standard deviation is associated with a decrease in coefficient-of-variation for the total PN conductance. Qualitatively speaking, the PN is more ‘mean driven’ when $\eta \sim 180\text{Hz}$, and the PN is more ‘fluctuation driven’ when $\eta \sim 30\text{Hz}$, even though the firing-rate of the PN is similar in both cases (Fig. 2E and Fig. 2F). This can be quantified by measuring, for example, the autocorrelation of the PN. In the case $\eta = 180\text{Hz}$, the PN autocorrelation shows several significant peaks, the first of which is at $\sim 13\text{ms}$, indicating periodic-firing at $\sim 70\text{Hz}$ (Fig. 2E). In the case $\eta = 30\text{Hz}$, the PN autocorrelation does not indicate a strong periodicity to the PN firing-patterns (Fig. 2E).

The simple cartoon described above only considers synaptic-depression resulting from vesicle-depletion. The real AL displays evidence of presynaptic-inhibition as well. Nevertheless, the same general principle still holds regardless of the source of synaptic-depression at the ORN synapses, as long as the PNs become more mean driven as ORN firing-rates increase. In fact, it is possible to show analytically that similar results hold across a wide range of parameters for an idealized system similar to this one (see the section entitled “A simple analyzable cartoon of variance coding” in Methods).

If this picture is accurate in the real AL, then the PN dynamics within any given glomerulus in the AL will change as a function of ORN input to that glomerulus, even when the mean PN firing-rates have saturated for that glomerulus. These dynamical changes will only be observable through measurements of statistics that are ‘higher-order’ than mean firing-rate. We note that synaptic-depression of the ORN synapses is not the only mechanism via which the PNs may become more mean-driven as ORN firing-rates increase — other mechanisms, such as spike-frequency adaptation, could also contribute to this effect. As long as the postsynaptic influence of each ORN spike decreases as m_{ORN} increases, the PN activity will become more mean-driven as m_{ORN} increases. As the PN activity becomes more mean-driven, we expect the firing-sequences produced by that PN to become more regular [22].

An illustration of variance coding within a large-scale model. We also observe this phenomenon within our large-scale model (described in Methods), which contains both presynaptic-inhibition and vesicle-depletion. To illustrate this phenomenon at work, we created a panel of 16 odors, all of which saturated the PN firing-rates (i.e., produced average PN firing-rates at the ‘plateau’ of the $m_{PN}(m_{ORN})$ curve for the model). We presented each of these odors to the model network 64 times.

For each of the 64 trials of each stimulus we measured the 30-component vector of PN firing-counts collected over the 256ms following odor onset. Each component of this vector represents the number of spikes fired by one of the 5×6 PNs during this time.

We then used this vector to perform each possible 2-way and 3-way stimulus discrimination task (see the section entitled “Odor Discrimination” in the Methods). Each of these 2-way and 3-way discrimination tasks results in a discriminability rate (i.e., the fraction of correctly categorized trials — note that chance performance for a 2-way task is 50%, and chance performance for a 3-way task is 33%).

We construct a histogram of the discriminability rates for the $\binom{16}{2}$ 2-way discrimination tasks, and as expected (see Fig. 3A), the typical discriminability rate for the system is not particularly high (recall that each odor saturated the PN firing-rates). Similarly, the $\binom{16}{3}$ 3-way discrimination tasks performed using PN firing-rate vectors also do not yield high discriminability rates (Fig. 3B). However, if instead of merely using PN firing-rate information we also use information regarding PN-PN correlations within the system, then the typical discriminability rates for the 2-way and 3-way tasks increase (see Fig. 3C,D). To produce the discriminability rates shown in Fig. 3C,D, we measured not only the 30-component vector of PN firing-counts for each odor trial, but also the 30^2 -component vector of PN-PN correlations (with correlation time 32ms). As expected, these higher-order statistics contain enough information to discriminate odors significantly more reliably than mere firing-rates.

The difference between the performance of these low-order and high-order readouts is more noticeable when the synaptic-depression in the system is strong. Conversely, in a network with no vesicle-depletion and reduced presynaptic-inhibition, the low- and high-order readouts yield more similar discriminability-rates (see Fig. 3E,F). Thus, the presence of strong synaptic-depression within our system is one factor which allows the network’s dynamics to encode input-specific information within the PN-PN correlations.

For the example shown in Fig. 3, the difference between the typical 2-way discriminability rates observed when using high-order versus low-order readouts is maximized when the synaptic-depression is strongest; the effect of variance coding is seen quite clearly. However, for the 3-way discriminability rates, the difference between the high- and low-order readouts is greatest when the presynaptic-inhibition is not too strong. A natural question is: why does the performance for the 3-way discrimination task not parallel that for the 2-way task? Why is the difference in performance between high- and low-order readouts not maximized when both presynaptic-inhibition and vesicle-depletion are at their strongest?

This effect arises in part because the 3-way task is quite difficult and the observation time T_{obs} over which the task is carried out is rather short — $T_{obs} = 256\text{ms}$ in this case. As we will argue below, one consequence of strong presynaptic-inhibition is that the network’s ability to perform fine discrimination will be compromised when T_{obs} is small. In order to perform very well on fine discrimination tasks when T_{obs} is small, the network should have only moderate amounts of presynaptic-inhibition (consistent with Fig. 3F).

Hypothesis 2: a tradeoff between reliability and sensitivity

It has been hypothesized that one functional role for the AL is to separate similar odors and that the nonlinear gain curve $m_{PN}(m_{ORN})$ is instrumental in this process. As shown in [11], the nonlinearity of $m_{PN}(m_{ORN})$ is influenced strongly by substantial synaptic-depression at the ORN synapses. Thus, it is

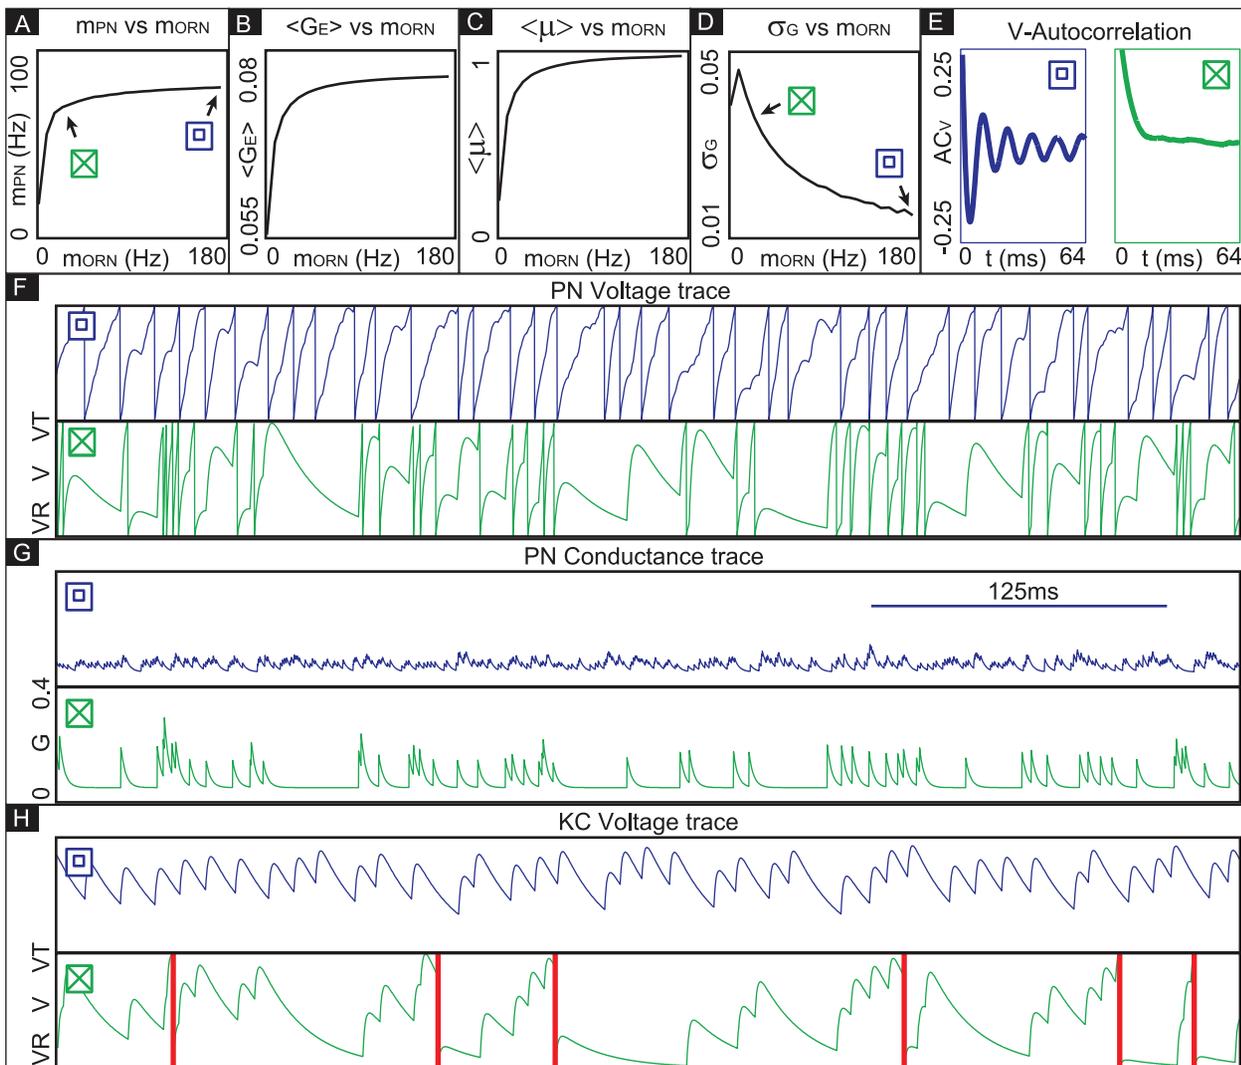


Figure 2. A simple illustration of variance coding. Here we presume the simple model described in the section entitled “An idealized model used to illustrate variance coding”. [A] There is a nonlinear relationship between the ORN firing-rate and the PN firing-rate. [B] There is also a nonlinear relationship between the ORN firing-rate and the time-averaged conductance of the PN. [C] As the ORN firing-rate increases, the time-averaged vesicle-depletion parameter increases and saturates. [D] Since the average vesicle-depletion parameter increases as the ORN firing-rate increases, the variance in the PN conductance is a decreasing function of ORN firing-rate, for sufficiently high ORN firing-rates. Two different points along this curve are indicated, corresponding to two different PN dynamical regimes with similar PN firing-rates. The ‘ \times ’ and ‘ \square ’ symbols indicate, respectively, an irregularly firing-regime and a regularly firing-regime. [E] As a result of the fact that the PN conductance has a low variance when the ORN firing-rates are high, the PN activity is very regular when the ORN firing-rate is high. In contrast, the PN activity is less regular when the ORN firing-rate is not as high. This is reflected in the normalized PN autocorrelation, which shows several significant peaks when the variance in the PN conductance is low (‘ \square ’-regime, left). In contrast, when the variance in the PN conductance is high the autocorrelation does not show significant peaks (‘ \times ’-regime, right). [F] The regularity in the PN spiking activity is seen in PN voltage trace, as shown for the ‘ \square ’-regime (top) and ‘ \times ’-regime (bottom). [G] The variance in the PN conductance is seen in PN conductance trace, as shown for the ‘ \square ’-regime (top) and ‘ \times ’-regime (bottom). [H] In this panel we show the voltage-trace of a putative Kenyon cell, a conductance-based integrate-and-fire-neuron, driven by either the PN from the ‘ \square ’-regime (top) or the PN from the ‘ \times ’-regime (bottom). Thick vertical lines indicate firing-events for this putative KC. When driven by the regular activity of the ‘ \square ’-PN, the KC maintains an elevated subthreshold voltage, but does not fire often. On the other hand, when driven by the irregular activity of the ‘ \times ’-PN, the KC does not maintain an elevated subthreshold voltage but fires after each burst in ‘ \times ’-PN-activity. This provides a simple illustration of one possible way in a variance-code could be ‘read-out’ by downstream neurons.
doi:10.1371/journal.pcbi.1002622.g002

reasonable to conclude that one functional role of synaptic-depression at the ORN synapses is to enhance the odor separation capabilities of the AL.

Within the fly AL there are multiple sources of synaptic-depression at the ORN synapses. Two major mechanisms which contribute to this synaptic-depression are vesicle-depletion and presynaptic-inhibition. While either one of these mechanisms

could, in principle, be the major contributing factor to the synaptic-depression observed within the fly AL, it seems as though both of these mechanisms play a substantial role in producing synaptic-depression [11,12]. Thus, one is faced with the following natural question: What purpose do these two distinct mechanisms serve within the fly AL? How would the odor-coding properties of the fly AL change if, say, only one of

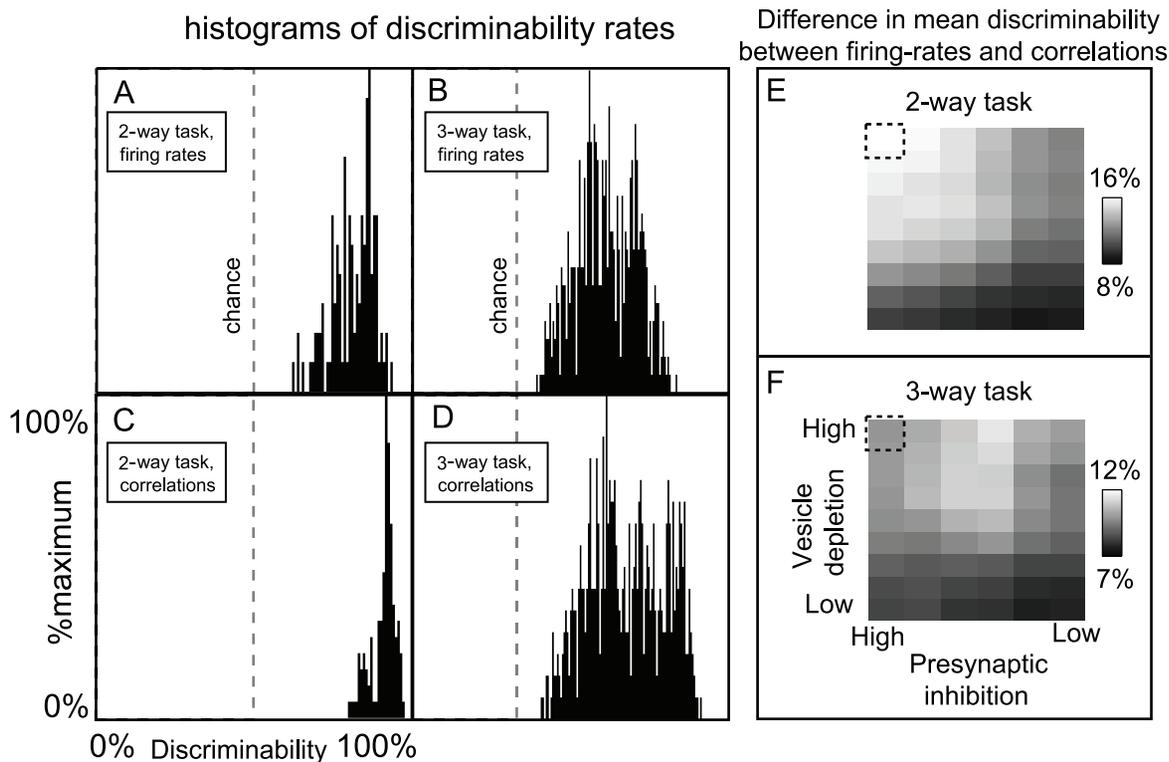


Figure 3. A manifestation of variance coding within the large-scale model. The large scale model (described in Methods) exhibits a phenomenon similar to the variance coding shown in Fig. 2. We constructed a panel of 16 odors, all of which only directly stimulated the same 3 glomeruli (although to differing degrees). Moreover, we chose every odor within this panel such that the ORN firing-rates of the 3 directly stimulated glomeruli were sufficient to saturate the firing-rates of the associated PNs (i.e., the directly stimulated ORN firing-rates were ≥ 12 Hz, see Fig. 10). Given this panel of odors, we presented each odor multiple times, and used the collection of 30-component PN firing-rate vectors (measured over the 256ms period immediately following odor onset) to perform a variety of odor discrimination tasks (see Results for details). [A] The histogram of discriminability rates associated with 2-way discrimination tasks when only firing-rate data is used. Note that 50% is chance level for these tasks (chance level is also shown in panels B,C,D). [B] The histogram of discriminability rates associated with the 3-way discrimination tasks when only firing-rate data is used (note that 33% is chance level for these tasks). [C] The histogram of discriminability rates associated with 2-way discrimination tasks when firing-rate data and 2-point correlations (correlation time 32ms) are used. [D] The histogram of discriminability rates associated with 3-way discrimination tasks when firing-rate data and 2-point correlations (correlation time 32ms) are used. Note that the typical discriminability rate is higher when correlations are used. [E] Here we plot the difference in mean discriminability for the 2-way discrimination task between the cases (i) when firing-rate data and 2-point correlations are used, and (ii) only firing-rate data is used. We plot this difference as a function of the parameters S^{vesdep} and S^{presyn} used in our large-scale model. The vesicle-depletion parameter S^{vesdep} ranges from 0 to 0.95 across the vertical axis, and the presynaptic-inhibition parameter S^{presyn} ranges from 1/3 to 2.25 across the horizontal axis. The data shown in panels A–D is taken from the simulation indicated by the dashed square. Note that, as the total amount of synaptic-depression decreases, the discriminability computed using only firing-rates is closer to the discriminability computed using both firing-rates and 2-point correlations. [F] Similar to panel-E, except for the 3-way discrimination task, rather than the 2-way discrimination task. doi:10.1371/journal.pcbi.1002622.g003

these mechanisms were responsible for the observed levels of synaptic-depression at the ORN synapses? Is there some functional advantage gained by having both of these mechanisms at play?

In what follows we introduce a hypothesis which links the underlying nature of synaptic-depression at the ORN synapses to information-coding properties of the AL, such as reliability, sensitivity and discriminability. First we will define these terms, and then we will explain our hypothesis in more detail throughout the rest of this section.

sources of noise: There are two sources of ‘noise’ in our network which influence the reliability (or unreliability) of the AL’s activity across trials. The first is the initial condition of the system (i.e., the state of the system at odor onset). Different initial conditions will give rise to different dynamic trajectories. The second source of noise is the odor-driven Poisson input to the ORNs in the model. Different trials will give rise to different sequences of ORN spikes.

reliability: We define the reliability of the AL as the inverse of the coefficient-of-variation in spike-counts of AL neurons, as measured across trials over a given stimulus-driven time-window. Reliability is high if the spike-counts of the AL neurons are similar from trial-to-trial. Reliability is low if the spike-counts vary significantly from trial to trial. In our analysis we will consider a family of networks with the same mean firing-rate, hence the notion of reliability can be constructed using standard-deviation in spike-counts across trials, rather than coefficient-of-variation.

sensitivity: Given two similar stimuli, we can measure the time-averaged firing-rates of the various neurons in the AL, collected over a long time (e.g., 1024ms). If the firing-rates induced by these two similar stimuli are nearly identical, we say that the AL is ‘not sensitive’ to the difference between these two stimuli. On the other hand, if the firing-rates induced by these two stimuli are quite different, then we would describe the AL as ‘sensitive’ to the stimulus difference. More specifically, we define sensitivity to be the magnitude of the derivative of the vector of steady-state AL-

firing-rates, when considered as a function of the odor input. In this sense, our notion of sensitivity is built around firing-rates, and does not explicitly consider higher order dynamical structure.

discriminability: Given an unknown odor from amongst a set of possible known candidates, we can use the AL as a discriminator: by presenting this mystery odor to the AL and measuring PN firing-counts over a time-period T_{obs} , we can attempt to classify the input as one of the possible candidate odors. We define the discriminability of the AL as the accuracy (i.e., correct-classification rate) of this procedure. The discriminability depends strongly on T_{obs} . If T_{obs} is sufficiently long, the discriminability of the AL is related directly to its sensitivity. If T_{obs} is short, then unreliability may come into play and reduce discriminability. As with our definition of sensitivity, our definition of discriminability is built around measurements of firing-rates, and does not take into account higher order dynamic structure.

The main thrust of our second hypothesis is that the combination of the mechanisms of vesicle-depletion and presynaptic-inhibition allows the fly AL to balance sensitivity and reliability in such a manner as to maximize the discriminability of AL activity (with respect to similar ORN inputs) over short observation times.

An illustration of the tradeoff between reliability and sensitivity within a large-scale model. In this subsection we will show how the hypothesis introduced above manifests within our large scale model. First we will discuss some features of this model which are pertinent to this hypothesis, then we will discuss our hypothesis in more detail.

We used simulations to investigate and benchmark our large-scale model (see the sections regarding benchmarking in the Methods). By analyzing these simulations we determined that, even after benchmarking, there were still a handful of free parameters that were left unconstrained. Two parameters in particular were not fully constrained by our benchmarking: (i) the strength of vesicle-depletion as characterized by S^{vesdep} , and (ii) the strength of presynaptic-inhibition as characterized by S^{presyn} . Within our large-scale model the combination of these two parameters produced synaptic-depression of the ORN synapses. While the total amount of synaptic-depression was constrained by our benchmarking, the relative strengths of S^{vesdep} versus S^{presyn} were not constrained.

As an example of this lack of constraint, consider the following benchmark: assume that we expect the average PN firing-rate within the AL to saturate at a certain level m_{PN} when stimulated sufficiently by ORN input. What we found was that there is a spectrum of possible AL architectures which could produce this desired firing-rate m_{PN} : (A) on one end of the spectrum is an AL in which there is hardly any vesicle-depletion of the ORN synapses, but for which the LNIs give rise to substantial presynaptic-inhibition at these synapses. This type-A AL would be characterized by a large value of S^{presyn} and a small value of S^{vesdep} . (B) on the other end of the spectrum is an AL in which vesicle-depletion is primarily responsible for synaptic-depression, and the presynaptic-inhibition of the ORN synapses due to LNIs is negligible. For this type-B AL S^{presyn} would be small and S^{vesdep} would be large.

An example of this spectrum is given in Fig. 4. Given a fixed value \bar{m}_{PN} for the saturated firing-rates of PNs in a strongly driven glomerulus, there exists a 1-parameter family of values ($S^{\text{presyn}}, S^{\text{vesdep}}$) which corresponds to networks exhibiting saturated firing-rates equal to \bar{m}_{PN} . This 1-parameter family of values ranges from networks with high S^{presyn} and low S^{vesdep} (i.e., type-A networks) to networks with high S^{vesdep} and low S^{presyn} (i.e., type-B networks).

Shown in Fig. 4A are the mean stimulus-driven PN spike-counts for several networks with varying values of S^{presyn} and S^{vesdep} . To construct this example we performed a systematic scan of parameter space for our large-scale network model. We selected a 2-dimensional array of parameter values for $S^{\text{presyn}}, S^{\text{vesdep}}$, ranging from $S^{\text{presyn}} \in \left[(4/3)^{-4}, (4/3)^3 \right]$ and from $S^{\text{vesdep}} \in [0.0, 0.9]$. For each fixed $S^{\text{presyn}}, S^{\text{vesdep}}$ within this array, we ran a large-scale simulation using a panel of 9 odors, and we ran 64 trials per odor. The first eight of the odors used stimulated three glomerular channels — the first glomerular channel was stimulated strongly, and an odor-specific subset of two other glomerular channels was stimulated weakly. The ninth odor only stimulated the first glomerular channel strongly. We remark that the simulations used to construct this array differ only in their values of S^{presyn} and S^{vesdep} . The architecture and connectivity of the rest of the model network were fixed.

In Fig. 4A we show the mean spike-count of PNs in the first glomerulus, for each pair of parameter-values $S^{\text{presyn}}, S^{\text{vesdep}}$. The mean spike-count is calculated as the mean of the number of spikes/48ms time-bin averaged across all 64 trials, and further averaged over the 128–512ms period following odor onset, and further averaged across all 9 odors. Overlaid on top of the mean spike-counts are contour lines for the spike-count. Each of these contours represents a 1-parameter family of networks with a different constant mean stimulus-induced spike-count. Note that, as indicated in Fig. 4B, these contours extend from regions of high S^{presyn} and low S^{vesdep} to regions of low S^{presyn} and high S^{vesdep} . In this example Type-A networks correspond to the lower-left corner of the array, and Type-B networks correspond to the upper-right corner of the array. Thus, in Fig. 4A it can be seen that m_{PN} is constant along contours extending from type-A networks (lower left) to type-B networks (upper right).

We observed two important systematic differences between the candidate networks along these 1-parameter families. First, type-B networks are more reliable than type-A networks. This can be understood as follows. First consider the ORN inputs to PNs in a type-B network (for which synaptic-depression is dominated by vesicle-depletion). A typical odor stimulates many ORNs to fire at a high rate. Each of the ORN synapses likely has a high quantal release rate [11], implying that the fraction of active vesicles remaining after several rapid ORN spikes is likely to be small. Moreover, there are ~ 30 such ORNs which converge onto each PN within their target glomerulus [23]. Thus each PN within a strongly stimulated glomerulus receives a large number of input spikes from a large number of presynaptic ORNs, each firing with a high rate, each synapse of which is likely to experience profound vesicle-depletion. Moreover, the vesicle-depletion experienced by the ORN synapses is only dependent on the ORN activity, and is independent of the activity of the AL. Thus, we expect the ‘feed-forward’ synaptic-depression observed within a type-B network to always exhibit very similar dynamic transients from trial to trial, with the only differences due to the variation in ORN spike-sequences induced by the trial-to-trial variability of the Poisson input to the ORNs [19]. Now, on the other hand, let us consider the ORN inputs to PNs within a type-A network. In such a network, synaptic-depression is primarily governed by ‘feedback’ from the AL in the form of presynaptic-inhibition. ORNs in a type-A network rely on the odor-specific firing patterns of LNIs in order to exhibit synaptic-depression, and therefore may receive different amounts of presynaptic-inhibition from trial to trial (or over disjoint time-windows within a single trial). Moreover, there are only a few LNIs per glomerulus, and a given stimulus may not cause all these LNIs to fire at high rates. A few extra LNI spikes

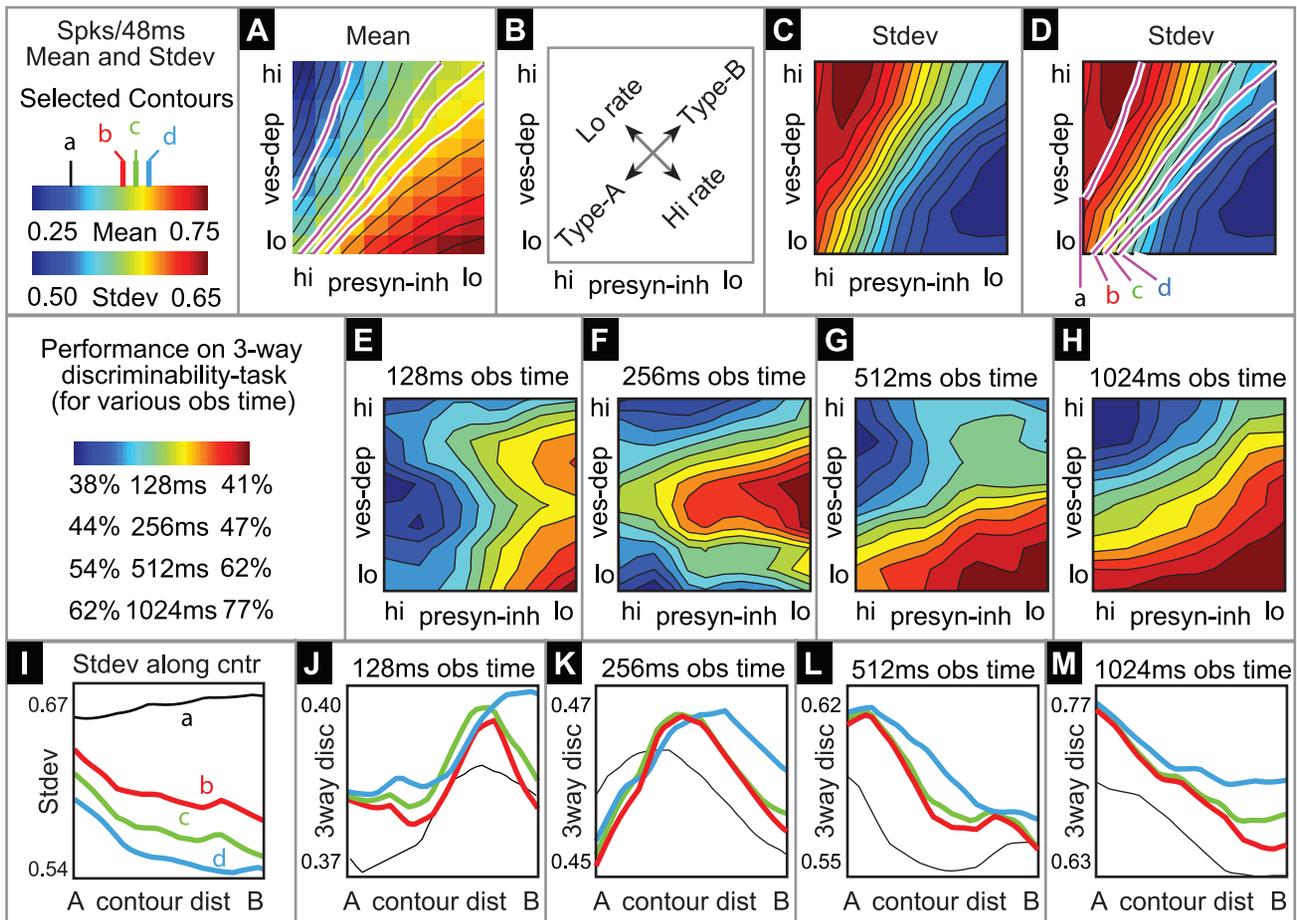


Figure 4. A tradeoff between reliability and sensitivity within our large-scale model. We performed a systematic scan of our large-scale network model, varying S^{presyn} and S^{vesdep} (see the section entitled “An illustration of the tradeoff between reliability and sensitivity within a large-scale model” in the main text for details). For each point in this parameter array we measured various features of the network dynamics (such as mean PN spike-counts and reliability), as well as the performance of each of these networks on a 3-way odor discrimination task. [A] Shown is the mean PN spike-count of PNs in the first glomerulus, for each pair of parameter-values $S^{\text{presyn}}, S^{\text{vesdep}}$. Overlaid on top of the mean spike-counts are contour lines for the spike-count. Four of these contours are highlighted in magenta, and will be referenced later. [B] Indications of the type-A and type-B network regimes. [C] Shown are the standard deviation in PN spike-counts of PNs in the first glomerulus (see colorbar on far left). [D] Reproduction of panel-C, along with the contours of panel-A. [E–H] Shown are contour plots associated with $R^{0.4}(S^{\text{presyn}}, S^{\text{vesdep}})$ for various values of T_{obs} . These panels use the colorbar shown to the far left. [I] Here we plot the standard-deviation in spike-count (taken from panel-D) as a function of the distance along each of the contours indicated in panel-D, with values bi-linearly interpolated as necessary. [J] Here we plot the discriminability values $R^{0.4}(S^{\text{presyn}}, S^{\text{vesdep}})$ indicated in panel-E as a function of the distance along each of the contours shown in panel-D. The contours are indicated using the colorcode from panel-I. [K–M] Similar to panel-J, except for $T_{\text{obs}} = 256\text{ms}, 512\text{ms},$ and 1024ms respectively.
doi:10.1371/journal.pcbi.1002622.g004

induced on any one trial may substantially change the footprint of synaptic-depression across the ORN synapses, thus leading to even more extra LNI spikes later on, and so forth. This ‘feed-back’ mechanism allows the synaptic-depression observed within type-A networks to exhibit quite different dynamic transients from trial to trial. Put another way, the ‘feed-back’ structure within type-A networks allows the trial-to-trial variability in LNI activity to affect and magnify the trial-to-trial variability in ORN input to the AL. In conclusion, we expect that ORN inputs to PNs in type-A networks will be less reliable than the corresponding ORN input to PNs for type-B networks when measured either (a) over multiple trials, or (b) over different time-windows within a single trial.

The second systematic difference between networks along such a 1-parameter family is that type-A network-dynamics is more sensitive than type-B network-dynamics to subtle changes in ORN input. To see why this might be true, let’s revisit the argument

used above. Consider a subtle change in ORN input which is only large enough at first to shift PN and LN firing rates slightly. This subtle change in ORN input will not create a large shift in the PN input for type-B networks, yet the same subtle change in ORN input may give rise to a few different LNI firing-events in the type-A network, which may then presynaptically inhibit different ORNs, giving rise to even more different type-A-network-activity, and so forth. In other words, due to the feedback between the type-A LNIs and the type-A ORNs, we expect the type-A system’s dynamics to be more sensitive than the type-B network’s dynamics to certain perturbations in input.

These systematic differences (i.e., type-A networks are less reliable, but more sensitive to perturbations in input than type-B networks) manifest within our large-scale model.

To quantify reliability for each network along such a 1-parameter family, we measured the trial-to-trial standard deviation

in PN spike-counts of PNs in the first glomerulus, for each pair of parameter values S^{presyn} , S^{vesdep} . The standard-deviation is calculated as the standard-deviation of the number of spikes/48ms time-bin across all 64 trials, averaged over the 128–512ms period following odor onset, then further averaged across all odors. The coefficient-of-variation in spike-counts is equal to the standard-deviation in spike-count divided by the mean. Thus, along contours of m_{PN} (where the mean is constant) the coefficient-of-variation in spike-count will be proportional to the standard-deviation in spike-count. Shown in Fig. 4D are the standard deviation in PN spike-counts along with the 4 contours highlighted in Fig. 4A. These four contours (labelled $\alpha, \beta, \gamma, \delta$) each correspond to a 1-parameter family of networks exhibiting a fixed mean spike-count, and are each associated with a different color (black, red, green, cyan, respectively) on the colorbar to the far left. In Fig. 4I we plot the standard-deviation evaluated along these contours. Note that, since contour γ is longer than contour β , the graphs shown in Fig. 4I are not directly comparable. However, there is a clear trend amongst all these graphs: As one moves along the 1-parameter family of networks with constant mean stimulus-induced spike-count from type-A networks to type-B networks the standard-deviation in spike-count decreases as long as the mean spike-count is sufficiently high (i.e., contours β, γ, δ). This is equivalent to the statement that, along contours β, γ, δ , type-B networks are more reliable than type-A networks.

Recall that, for each network (i.e., for each fixed value of $S^{\text{presyn}}, S^{\text{vesdep}}$), we ran 64 trials for each of 9 different odors. Using this data, we can quantify the sensitivity of each of these networks to input perturbations. For each odor trial we measure the 30-component PN firing-rate vector averaged over the $T_{\text{obs}} = 1024\text{ms}$ time-window including and immediately following a 512ms odor presentation. We use these time-averaged firing-rate vectors to perform each of the $84 = \binom{9}{3}$ 3-way odor discrimination tasks, and thus we obtained a distribution of discriminability rates for each 3-way odor task (see the section entitled “Odor Discrimination” in the Methods). For each network we then record the 40th-percentile of the distribution of discriminability rates (across odors), denoted by $R^{0.4}$. We chose to display $R^{0.4}$, as this 40th-percentile discriminability rate illustrates our conclusions most clearly. However, our main results do not change if we choose another percentile in the range [10,60]. Higher percentiles, such as the 70th-percentile, are usually all near 100% correct-classification, since the set of odors used contain several rather distinct odors. Note that $R^{0.4}$ will depend on S^{presyn} and S^{vesdep} . Shown in Fig. 4H are the contour plots associated with $R^{0.4}(S^{\text{presyn}}, S^{\text{vesdep}})$ for $T_{\text{obs}} = 1024\text{ms}$. In Fig. 4M we plot these discriminability rates along the contours shown in panel-A. For each of these contours the maximum discriminability (when $T_{\text{obs}} = 1024\text{ms}$) occurs at the type-A end of the spectrum. This indicates that the discriminability of type-A networks (using firing-rates measured over long observation times) is superior to that of the type-B networks. This is a reflection of the fact that type-A networks are more sensitive than type-B networks to subtle changes in input.

A combination of vesicle-depletion and presynaptic-inhibition is required to optimize discriminability over short observation-times within a large-scale model. Within our model network we have observed a further functional consequence associated with the tradeoff between reliability and sensitivity described above. Type-A networks are indeed more sensitive than type-B networks to shifts in input, and this sensitivity is reflected in the long time (or trial averaged) PN firing-rate vector associated with any given input. As a result, type-A networks

outperform type-B networks in odor-discrimination tasks when the discriminability rate is calculated using a long time observation (such as $T_{\text{obs}} = 1024$, as shown in Fig. 4H,M). However, type-A networks are less reliable than type-B networks, and thus, if the observation-time of any given odor stimulus is sufficiently short, the higher variability associated with the single-trial short-time responses of type-A networks will pollute the performance of any discrimination task which uses only these short observations. On the other hand, since type-B networks are rather reliable, shortening the observation-time associated with a discrimination task will not affect the discriminability rate associated with that task for a type-B network as much. As demonstrated in our model network, if the observation-time of any given odor trial is shortened from 1024ms (as shown in Fig. 4H,M) to merely 128ms after odor onset, the decreased reliability associated with type-A networks will drastically lower the discriminability rate of the odor-discrimination tasks which use only these short observations (see Fig. 4E,J). Moreover, since the type-B networks are more reliable than type-A networks, the decrease in discriminability associated with reducing the observation-time of the discrimination task is lower for type-B networks than it is for type-A networks (compare Fig. 4J,M). Most intriguingly, there is a midpoint in the spectrum — a balance between vesicle-depletion and presynaptic-inhibition — which gives rise to the maximum discriminability rates using only short-time observations. This optimal point depends on the length of the observation-time associated with the discrimination-task. With long observation-times type-A networks are optimal. With very short observation-times type-B networks are optimal.

This feature is shown in more detail for our large scale model in Fig. 4E,F,G,J,K,L, which illustrate the discriminability capabilities of our model for a variety of observation times $T_{\text{obs}} = 128, 256, 512$. Note that, for any particular contour $\alpha, \beta, \gamma, \delta$, The point of maximum performance occurs closer to the type-B extreme when T_{obs} is small, and this maximum occurs closer to the type-A extreme when T_{obs} is large. In other words, when T_{obs} is low, type-B networks outperform type-A networks, whereas when T_{obs} is large type-A networks outperform type-B networks.

In conclusion, we have demonstrated that for a particular set of discrimination tasks the network which performs optimally lies in between the type-A and type-B extremes. Moreover, as the observation-time associated with this task increases (or decreases) the optimal point shifts towards the type-A (or, respectively, type-B) end of the spectrum. Although the details of Fig. 4 only pertain to a particular discrimination task, we mention now that this systematic dependence of the optimal point on observation-time is actually a natural consequence of the fact that type-A networks are more sensitive, and type-B networks are more reliable. Indeed, as we will argue below (in a section entitled “A simple cartoon of optimizing discriminability over short observation-times”), this feature is to be expected for a rather general class of discrimination tasks in which estimates of the mean firing-rates (sampled over an observation-time) are used to classify the input.

A simple analyzable cartoon of the tradeoff between reliability and sensitivity. In this section we introduce a simple deterministic 2-neuron model network which will allow us to discuss various aspects of hypothesis-2. This simple network has the property that the sequence of neuronal firing-events is a sensitive function of the network’s initial conditions as well as the input to the network and the source of synaptic-depression within the network. The model itself consists of 2 LNIs, each driven by a single ORN. Each LNI (labelled *A* and *B*) is modeled by a simple phase-oscillator (similar to a current-based integrate-and-fire neuron), and each ORN is modeled by a fixed input-current (i.e., η_A and η_B). This input current indicates the rate at which

each neuron would fire if there were no presynaptic-inhibition or vesicle-depletion.

In this system the strength of presynaptic-inhibition is modeled by a constant parameter ξ . As ξ increases, the firing-events of each neuron have a greater inhibitory effect on the input to the other neuron. Specifically, whenever LNI A fires, the ORN input to LNI B is shut off for ξ -time. Similarly, whenever LNI B fires, the ORN input to LNI A is shut off for ξ -time. To ensure that both neurons fire, we assume $\xi < \min(1/\eta_A, 1/\eta_B)$. As the amplitudes η_A, η_B of the ORN processes are constant, the vesicle-depletion in this system is assumed to attain a steady state, and is modeled via a single constant parameter $\mu \geq 1$, which reduces the ORN input to both A and B .

In keeping with the description above, the membrane potentials for LNI A and B obey the differential equations

$$\begin{aligned} \frac{dV_A(t)}{dt} &= \eta_A/\mu \Theta_A(t) = \eta_A/\mu \left[1 - \int_{t-\xi}^t \sum_k \delta(s - T_k^B) ds \right] \\ \frac{dV_B(t)}{dt} &= \eta_B/\mu \Theta_B(t) = \eta_B/\mu \left[1 - \int_{t-\xi}^t \sum_k \delta(s - T_k^A) ds \right], \end{aligned} \quad (1)$$

and whenever the potential V_A reaches $V_T = 1$, we say that LNI A fires, and reset V_A to $V_R = 0$. The k^{th} spiketime of neuron A is recorded as T_k^A . Similarly, whenever V_B reaches 1, we say that LNI B fires, and record the k^{th} spiketime of neuron B as T_k^B . The term $\Theta_A = 1 - \int_{t-\xi}^t \sum_k \delta(s - T_k^B) ds$ is equal to 1, unless neuron B has fired within ξ of the current time, in which case $\Theta_A = 0$. Similarly, the term $\Theta_B = 1 - \int_{t-\xi}^t \sum_k \delta(s - T_k^A) ds$ is equal to 1, unless A has fired within ξ of the current time, in which case $\Theta_B = 0$. Note that, since $\xi < \min(1/\eta_A, 1/\eta_B)$, the terms Θ_A and Θ_B are each either 0 or 1 at each time.

This simple network is easy to analyze, and the firing-rates m_A, m_B of each LNI in the network, as well as the interspike-interval-distributions (ISI_A , and ISI_B) can be directly calculated in terms of the inputs to the network η_A, η_B and the sources of synaptic-depression ξ , and μ . See the section entitled ‘‘A simple model illustrating the tradeoff between reliability and sensitivity’’ in the Methods for more details.

An example of such calculations is shown in Fig. 5. If we fix $\mu = 1$, then the firing-rates of the two neurons is a decreasing function of ξ for small ξ (see Fig. 5A). As expected, this decrease in firing-rate corresponds to the two neurons interfering with and slowing down one another. However, this interference causes the firing-events of each neuron to occur at irregular intervals, and hence the variance in the ISI-distributions of these neurons is a monotonic *increasing* function of ξ for small ξ (see Fig. 5B). Thus, as ξ increases from 0, the neurons fire less, and have a lower trial-to-trial reliability.

If we fix ξ , then the firing-rates of the two neurons are decreasing functions of μ . Thus, there is clearly a 1-dimensional family of synaptic-depression parameters which gives rise to networks exhibiting the same mean firing-rates for any fixed set of inputs. This 1-parameter family ranges from type-A networks (with high ξ and low μ) to type-B networks (with low ξ and high μ) – see Fig. 5C. We can index networks along this 1-parameter family using ξ , assuming that $\mu(\xi)$ is chosen so that the average firing-rate \bar{m} is maintained (see Fig. 5D).

Intuitively, one expects that for an extreme type-B network (i.e., $\xi = 0$) the activity should be perfectly regular: each neuron fires independently of the other neuron. On the other hand, for an extreme type-A network (i.e., $\xi > 0$), each neuron fires in spurts,

constantly disrupting the periodicity of the other neuron’s activity. This interplay between the neurons (resulting from presynaptic-inhibition) gives rise to a greater variability in the ISI-distributions of the neurons within the type-A networks. This increased variability implies that the neurons in the type-A networks have a lower trial-to-trial reliability than the analogous neurons within the type-B networks (assuming that different trials have different initial conditions).

This same intuition can be extended to see that, in the type-A network, there is a ‘rich-get-richer’ phenomenon: the neuron which gets more input will slow the other neuron down more than it is slowed down by the other neuron. Thus, the presynaptic-inhibition ξ magnifies the sensitivity of the type-A networks, increasing the difference in firing-rates between the two neurons when the input to these neurons is similar. In other words, the firing-rates produced by the type-A network should be more sensitive than those produced by the type-B network to small differences in inputs.

This intuition is borne out by analysis. When constraining μ, ξ so that the mean firing-rate \bar{m} is constant, we see that as ξ increases from 0 the network becomes more sensitive (i.e., the difference in firing-rates $m_A - m_B$ increases) and less reliable (i.e., the variance of the ISI-distributions of the two neurons increases). These functions are plotted in Fig. 5D,E. Thus, for this simple system, we can show analytically that hypothesis-2 holds: type-A networks are more sensitive, and type-B networks are more reliable. See the section entitled ‘‘A simple model illustrating the tradeoff between reliability and sensitivity’’ in Methods for more details.

A simple cartoon of optimizing discriminability over short observation-times. As postulated above, and illustrated for both a large-scale and idealized network architecture, we expect there to be a 1-parameter family of networks with the same mean firing-rate for any fixed set of inputs. This 1-parameter family ranges from type-A networks (with significant presynaptic-inhibition, lower reliability and higher sensitivity) to type-B networks (with significant vesicle-depletion, higher reliability and lower sensitivity). It turns out that, under rather general conditions, the networks which perform best on discriminability tasks with finite observation-times are the networks in the middle of this spectrum (i.e., networks with a combination of presynaptic-inhibition and vesicle-depletion).

To illustrate this principle, we will use the network architecture discussed in the section entitled ‘‘A simple analyzable cartoon of the tradeoff between reliability and sensitivity’’. It should be noted however, that the argument we will present here is not specific to the 2-neuron architecture discussed above, and a modified version of this argument will hold for any 1-parameter family of networks ranging from the type-A to the type-B extremes discussed above.

To begin, let us consider the following discriminability task. Assume that a simple 2-neuron network (of the type described in the section entitled ‘‘A simple analyzable cartoon of the tradeoff between reliability and sensitivity’’) is driven by one of two inputs — either (\mathcal{O}_1) neuron A is driven at rate $\bar{\eta} + \Delta\eta/2$ and neuron B is driven at rate $\bar{\eta} - \Delta\eta/2$, or (\mathcal{O}_2) A is driven at $\bar{\eta} - \Delta\eta/2$ and B is driven at $\bar{\eta} + \Delta\eta/2$. The steady-state ISI distributions $ISI_A(\mathcal{O}, \mu, \xi)$ and $ISI_B(\mathcal{O}, \mu, \xi)$ will depend on the unknown input \mathcal{O} , as well as the known system parameters μ, ξ . By performing a measurement of the system, is it possible to tell which input (i.e., either \mathcal{O}_1 or \mathcal{O}_2) is driving the system? Let us assume that our measurement process consists of 2 steps. First, we estimate the mean of $ISI_A(\mathcal{O}, \mu, \xi)$ by drawing T_{obs} samples from $ISI_A(\mathcal{O}, \mu, \xi)$. For example, we can either measure A within a single system for a long time, or we can measure multiple systems from an ensemble. The longer a single

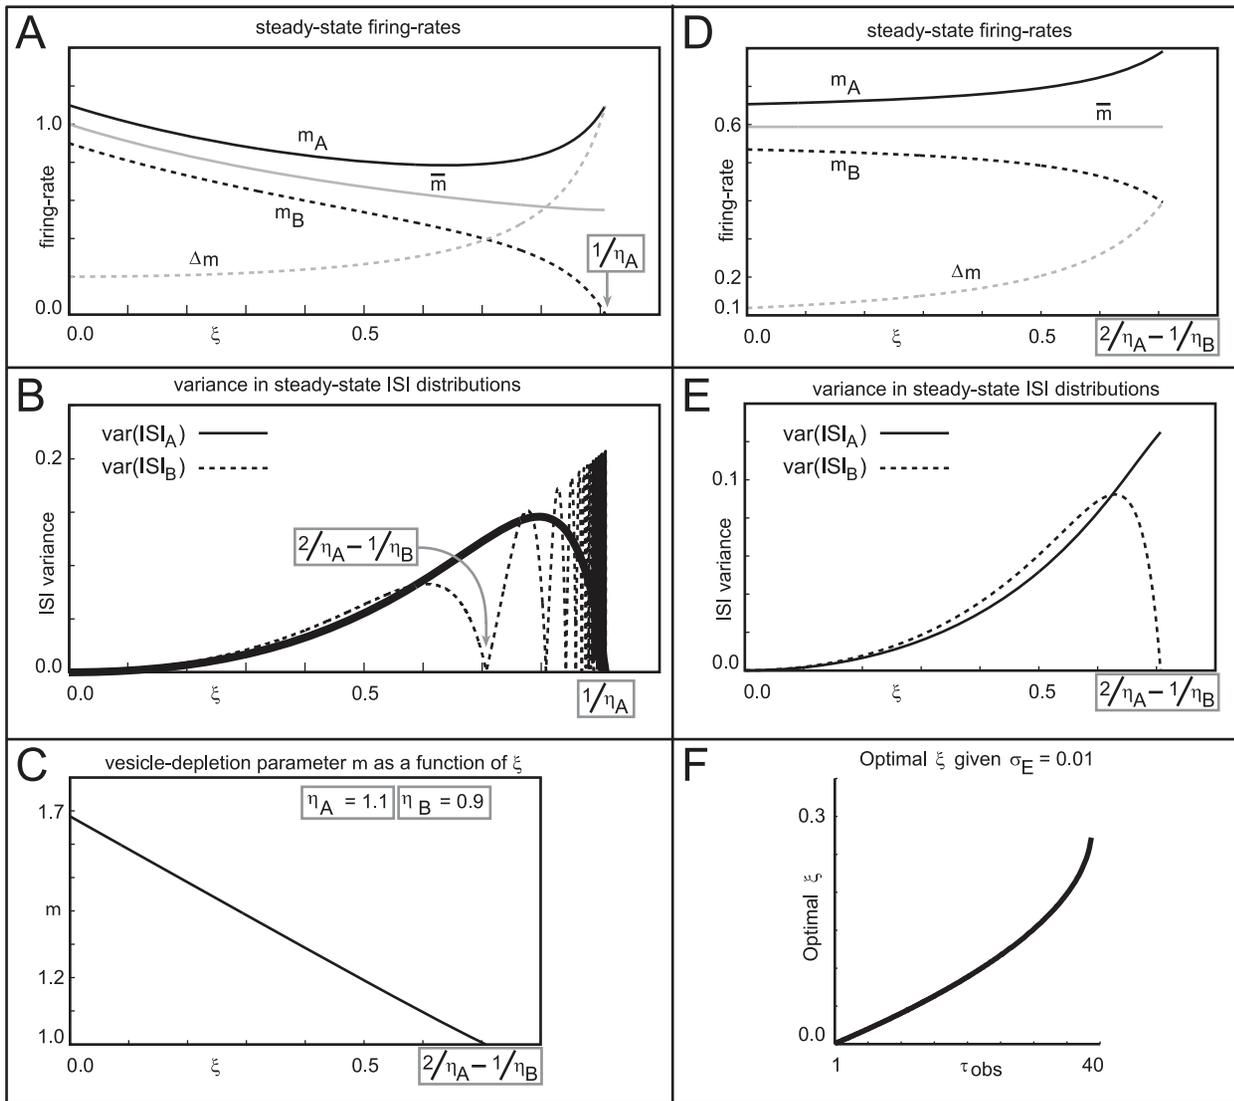


Figure 5. A simple analyzable cartoon of the tradeoff between reliability and sensitivity. In this example $\eta_A = 1.1$, and $\eta_B = 0.9$. In panels A and B the vesicle-depletion parameter $\mu \equiv 1$. In panels C,D,E and F, the vesicle-depletion parameter $\mu = \mu(\xi)$, such that the mean firing rate $\bar{m} = (m_A + m_B)/2 \approx 0.594$ is held constant. [A] Graphs of m_A (solid), m_B (dashed), $\bar{m} = (m_A + m_B)/2$ (gray), and $\Delta m = m_A - m_B$ (gray dashed), as functions of ξ , for the case $\mu \equiv 1$. [B] Graphs of $\text{var}(ISI_A)$ (solid) and $\text{var}(ISI_B)$ (dashed) as functions of ξ , for the case $\mu \equiv 1$. [C] Graph of μ as a function of ξ , subject to the constraint that \bar{m} remain constant. The constant value of \bar{m} chosen (essentially arbitrarily) in this case is the value of $\bar{m} \approx 0.594$ shown in panel A for $\xi = 2/\eta_A - 1/\eta_B$. Other choices of \bar{m} yield similar results. Note that this graph is monotonically decreasing, implying the existence of a 1-parameter family of networks possessing the same \bar{m} — ranging from type-A networks with low μ and high ξ , to type-B networks with high μ and low ξ . [D] Graphs of $m_A(\xi)$ (solid), $m_B(\xi)$ (dashed), $\bar{m}(\xi)$ (gray), and $\Delta m(\xi)$ (gray dashed), for the case $\mu = \mu(\xi)$. [E] Graphs of $\text{var}(ISI_A)$ (solid) and $\text{var}(ISI_B)$ (dashed) as functions of ξ , for the case $\mu = \mu(\xi)$. [F] Graph of the optimal choice of ξ (implying a vesicle-depletion parameter of $\mu(\xi)$) for which discriminability is maximized, as a function of the sample number T_{obs} . The notion of discriminability is described in the section entitled “A simple cartoon of optimizing discriminability over short observation-times”. In this case the observation error σ_E is fixed at $\sigma_E = 0.01$. Note that for low T_{obs} , discriminability is maximized for a type-B network. However, as T_{obs} increases, discriminability is maximized by type-A networks. The graph shown plots ξ_{opt} for $T_{\text{obs}} \in [1, 40]$, as for this particular simple example the derivative of $\xi_{\text{opt}}(T_{\text{obs}})$ reaches a vertical asymptote at $T_{\text{obs}} \sim 40$. doi:10.1371/journal.pcbi.1002622.g005

system is measured for, the larger the effective number of samples T_{obs} is, assuming that measurements of the system that are sufficiently well-separated in time are effectively independent. We use the notation T_{obs} rather than N_{obs} to draw analogy with the observation-time discussed in the section entitled “A combination of vesicle-depletion and presynaptic-inhibition is required to optimize discriminability over short observation-times within a large-scale model”. Let us denote by $\langle ISI_A \rangle_{T_{\text{obs}}}$ this estimate for the mean of ISI_A . Second, we assume that our measurement

process includes some external noise modeled by a random variable X_E . For the sake of presentation, let’s assume that X_E is drawn from $\mathcal{N}(0, \sigma_E)$ (i.e., a Gaussian distribution with mean 0 and variance σ_E^2). Thus, our final measurement of the mean of $ISI_A(\mathcal{O}, \mu, \xi)$ is some estimate $\langle ISI_A(\mathcal{O}, \mu, \xi) \rangle_{T_{\text{obs}}} + X_E$.

Our goal is to determine from this measurement $\langle ISI_A(\mathcal{O}, \mu, \xi) \rangle_{T_{\text{obs}}} + X_E$ whether the input to the system is \mathcal{O}_1 or \mathcal{O}_2 . By analyzing the signal to noise ratio of this measurement process (see the section entitled “Analysis of signal-to-noise ratio in

a general discrimination task” in Methods), we can show that the discrimination error $E(T_{\text{obs}}, \mu, \xi)$ associated with the best linear-classifier for this problem is well-approximated by:

$$E(T_{\text{obs}}, \mu, \xi) \approx 1 - \text{erf} \left(\frac{1}{2\sqrt{2}} \cdot \frac{\Delta\lambda(\mu, \xi)}{\sqrt{\sigma^2(\mu, \xi)/T_{\text{obs}} + \sigma_E^2}} \right), \quad (2)$$

where $\Delta\lambda(\mu, \xi)$ is the difference in the means of $ISI_A(\mathcal{O}_1, \mu, \xi)$ and $ISI_B(\mathcal{O}_1, \mu, \xi)$, and $\sigma^2(\mu, \xi)$ is the average variance of $ISI_A(\mathcal{O}_1, \mu, \xi)$ and $ISI_B(\mathcal{O}_1, \mu, \xi)$. Because erf is monotonic increasing, $E(T_{\text{obs}}, \mu, \xi)$ is minimized when the ratio $\Delta\lambda/\sqrt{\sigma^2/T_{\text{obs}} + \sigma_E^2}$ is maximized. Recall the structure of the simple 2-neuron networks described above — both $\Delta\lambda(\mu, \xi)$ and $\sigma(\mu, \xi)$ are monotonically increasing functions of ξ , and $\mu(\xi)$ can be defined implicitly through ξ (by fixing \bar{m}) as a monotonically decreasing function of ξ . For these simple networks, when $T_{\text{obs}} \gg 1$, this discriminability error is minimized when $\Delta\lambda$ is as large as possible, and the maximum $\Delta\lambda$ is achieved when μ, ξ correspond to a type-A network. Conversely, when $T_{\text{obs}} \sim 1$, then the discriminability error is minimized for μ, ξ corresponding to a network in between the type-A and type-B extremes. In this case, one can show that if $\Delta\eta, \sigma_E$ are sufficiently small, then the optimal ξ_{opt} (for which the error $E(T_{\text{obs}}, \mu(\xi), \xi)$ is minimized) increases as T_{obs} increases. (see Fig. 5F, which displays $\xi_{\text{opt}}(T_{\text{obs}})$ for the case $\bar{\eta} = 1$, $\Delta\eta = 0.2$, $\sigma_E = 0.01$).

As we mentioned earlier, The argument given in this section is quite general, and similar reasoning can be applied whenever any measurement is made by sampling from a distribution and adding an observation error X_E . Given a 1-parameter family of networks indexed by α , and a measurement of any dynamical feature λ , with sensitivity described by $\Delta\lambda(\alpha)$ and reliability described by $\sigma^2(\alpha)$, the error associated with the best linear-classifier can be approximated by an equation similar to Eq. 2, where T_{obs} increases as the observation-time of the measurement increases.

A population-dynamics approach towards verifying Hypothesis 2 within more general networks. In the preceding sections we have described in detail a specific 2-neuron network which exhibits the phenomena associated with hypothesis-2. However, the reasoning used in these sections cannot readily be applied to more complicated heterogeneous networks composed of more realistic model neurons. Indeed, while there exist networks for which hypothesis-2 holds (e.g., the large-scale networks described earlier on), there also exist networks for which hypothesis-2 does not hold. A natural question is: given a specific network architecture, what dynamic phenomena will that network exhibit? In the remainder of this section we will apply a rather general method [24,25] which can be used to assess the equilibrium dynamics of pulse-coupled networks, and which can be used to determine which network architectures exhibit phenomena associated with hypothesis-2 (e.g., the tradeoff between reliability and sensitivity discussed above). With this analysis we will be able to see that hypothesis-2 holds for a rather large class of networks, and in particular holds for a class of sparse randomly connected networks, provided that the network size is sufficiently large.

For the purposes of illustration, let us consider a network of N discrete-state glomeruli (LNIs), each driven by a different ORN. We will model each ORN-LNI pair as a discrete-state discrete-time Markov process which is as simple as possible, while still retaining the following features: (i) each LNI generates spikes, (ii) each ORN input spike contributes to the vesicle-depletion of that

ORN→LNI synapse, and (iii), each LNI spike gives rise to presynaptic-inhibition of some subset of ORN→LNI synapses. This model does not take into account excitatory interactions; PNs and LNEs are not included. While these excitatory interactions certainly contribute to hypothesis-2, they do not substantially change the following analysis, and we delay discussion of their effects until the end of this section.

Within this simple network model we will model the j^{th} ORN-LNI pair using the state-variables $V_j(t)$, $\mu_j(t)$ and $\xi_j(t)$ which represent LNI membrane-potential, ORN vesicle-depletion and ORN presynaptic-inhibition, respectively. The architecture of the model is determined by the connectivity matrix Δ_{jk} which encodes the presynaptic-inhibitory coupling between ORN-LNI pairs. The other parameters of the model include the feedforward input rates η_j to each ORN-LNI pair, as well as the overall strength of vesicle-depletion κ_μ and the overall strength of presynaptic-inhibition κ_ξ . The details of the model are given in a section entitled “A discrete state model used to analyze hypothesis-2 within general networks with arbitrary architecture” in Methods (see Eq. 16).

We are interested in how the dynamics of such a network depends on the connectivity matrix Δ , and also on other parameters such as the inputs η_j , the vesicle-depletion strength κ_μ and the presynaptic-inhibition strength κ_ξ . For reference, let us specify precisely what we mean by ‘hypothesis-2’. Let us say that the j^{th} LNI satisfies ‘hypothesis 2.0’ if there exists a 1-parameter family of small variations in κ_ξ, κ_μ which maintain the firing rate of the j^{th} LNI (denoted by m_j), such that this 1-parameter family ranges from high κ_ξ and low κ_μ (i.e., type-A) to low κ_ξ and high κ_μ (i.e., type-B). Let us say that the j^{th} LNI satisfies ‘hypothesis-2.1’ if, given a small increase in κ_ξ along this 1-parameter family, the reliability of the j^{th} LNI decreases (i.e., $\text{var}(ISI_j)$ increases as the network parameters are shifted towards a type-A network). Finally, let us say that the j^{th} LNI satisfies ‘hypothesis-2.2’ if, given a small increase in κ_ξ along this 1-parameter family, the sensitivity of the j^{th} LNI to its own input increases (i.e., $\partial_{\eta_j} m_j$ increases as the network parameters are shifted towards a type-A network).

Let us assume that we have some large network in which $\eta_j = \bar{\eta} = 0.5$ (i.e., the input to each LNI is the same), and that Δ is given (but otherwise arbitrary). One can readily show that hypothesis-2.0 holds — namely, for sufficiently small κ_ξ , each ORN-LNI pair with at least one presynaptic-inhibitory input has the property that there exists a 1-parameter family of parameters (ranging from high κ_ξ and low κ_μ to low κ_ξ and high κ_μ) for which the firing rate m_j remains fixed. This is simply because, as either κ_ξ or κ_μ increases, the firing rate m_j decreases (i.e., $\partial_{\kappa_\xi} m_j$ and $\partial_{\kappa_\mu} m_j$ are both negative) as long as κ_ξ is sufficiently small.

In the rest of this section, we will analyze reliability (i.e., hypothesis-2.1). Let us concentrate on a single ORN-LNI pair (say, the j^{th} such pair) embedded within this larger network, and assume for the moment that the j^{th} ORN is presynaptically-inhibited by the k^{th} LNI (with $k \neq j$). If we were to increase the strength of the presynaptic-inhibitory connection between LNI k and ORN j (i.e., if we were to increase Δ_{jk}) without decreasing the strength of vesicle-depletion κ_μ , then the firing rate m_j would drop. If, instead, we were to increase Δ_{jk} while decreasing κ_μ simultaneously so as to maintain m_j (as required by hypothesis-2.1), then the ISI distribution of the j^{th} LNI would change (but the firing rate m_j would remain constant by construction). In this case the differential shift in the ISI distribution of the j^{th} LNI associated with increasing the strength of connection Δ_{jk} (while appropriately decreasing κ_μ) gives rise to an increase in $\text{var}(ISI_j)$. Thus, if coupling strengths are sufficiently weak, then the derivative of

$\text{var}(ISI_j)$ with respect to increasing Δ_{jk} (while appropriately decreasing κ_μ) is positive. Using population-dynamics techniques from [24,25], this reasoning can be systematically extended to consider every connection in the network, not merely the connection Δ_{jk} .

Formally speaking, this analysis is nothing more than a Taylor-expansion of $\text{var}(ISI_j)$ in terms of the coupling strengths of the network. Namely, $\text{var}(ISI_j)$ depends on many parameters (e.g., Δ , η_j , κ_ξ , κ_μ), and if we assume that κ_μ is implicitly dependent on κ_ξ, Δ in such a way that m_j is constant, then we can Taylor-expand $\text{var}(ISI_j)$ in terms of κ_ξ and the components of Δ . If we retain all terms up to second order, such an expansion has the form

$$\begin{aligned} \text{var}(ISI_j) = & \alpha + \kappa_\xi \beta_{jj} \Delta_{jj} + \kappa_\xi \sum_{k \neq j} \beta_{jk} \Delta_{jk} + \kappa_\xi^2 \gamma_{jjjj} \Delta_{jj}^2 + \kappa_\xi^2 \sum_{k \neq j} \gamma_{jkjk} \Delta_{jk}^2 \\ & + \kappa_\xi^2 \sum_{k \neq j} \gamma_{jjkk} \Delta_{jj} \Delta_{kk} + \kappa_\xi^2 \sum_{k \neq j} \gamma_{jkkj} \Delta_{jk} \Delta_{kj} + \kappa_\xi^2 \sum_{k,l,j \text{ distinct}} \gamma_{jkjl} \Delta_{jk} \Delta_{jl} \quad (3) \\ & + \kappa_\xi^2 \sum_{k,l,j \text{ distinct}} \gamma_{jkkk} \Delta_{jk} \Delta_{kk}, \end{aligned}$$

where α is a 0th-order contribution, each β term corresponds to a 1st-order contribution, and each γ term corresponds to a 2nd-order contribution. Each β and γ term is a correction to $\text{var}(ISI_j)$ associated with a particular subnetwork containing the j^{th} LNI (see Fig. 6 for an example). For example, the term $\kappa_\xi \beta_{jj}$ corresponds to the subnetwork in which the j^{th} LNI (see Fig. 6 for an example). For example, the term $\kappa_\xi \beta_{jj}$ corresponds to the subnetwork in which the j^{th} LNI presynaptically-inhibits its own ORN (the j^{th} ORN) — $\kappa_\xi \beta_{jj}$ is the differential correction to $\text{var}(ISI_j)$ associated with increasing the connection strength Δ_{jj} , while appropriately decreasing κ_μ . This correction $\kappa_\xi \beta_{jj}$ is actually negative (i.e., if one were to increase Δ_{jj} , then the j^{th} LNI would become more reliable). As another example, the term $\kappa_\xi^2 \gamma_{jkjl}$ (with k, l, j distinct) corresponds to the subnetwork in

which both the l^{th} LNI and the k^{th} LNI presynaptically-inhibit the j^{th} ORN. If both Δ_{jk} and Δ_{kl} were to increase, then the change in $\text{var}(ISI_j)$ would be well-approximated by the change in the terms $\kappa_\xi \beta_{jk} \Delta_{jk} + \kappa_\xi \beta_{jl} \Delta_{jl} + 2\kappa_\xi^2 \gamma_{jkjl} \Delta_{jk} \Delta_{jl}$ (in this manner, the 2nd-order correction $\kappa_\xi^2 \gamma_{jkjl}$ captures the change to $\text{var}(ISI_j)$ which is not accounted for by $\kappa_\xi \beta_{jk}$ and $\kappa_\xi \beta_{jl}$). Indeed, given any specific network containing the j^{th} ORN-LNI pair, one can determine the effect of increasing κ_ξ (or, equivalently, increasing all of the components of Δ simultaneously) by dissecting the specific network and determining the contributions made by the various comprised subnetworks to $\text{var}(ISI_j)$. See Fig. 7 for an illustration of this technique.

By analyzing the various terms in this expansion, one can determine that by increasing the strength of certain elements of Δ , it is possible to actually lower $\text{var}(ISI_j)$, and make the j^{th} LNI more reliable. For example, by exclusively strengthening Δ_{jj} without increasing the other Δ_{jk} , the reliability of the j^{th} LNI would increase, in seeming contradiction to hypothesis-2.1. However, in a typical random network (containing many ORN-LNI pairs, and many presynaptic-inhibitory connections), the subnetworks which increase $\text{var}(ISI_j)$ dominate those that lower $\text{var}(ISI_j)$, and thus a uniform increase in κ_ξ will increase $\text{var}(ISI_j)$. By analyzing the magnitudes of the various α, β, γ terms in Eq. 3, one can quantify this statement for any particular class of networks. For example, consider a random network of N neurons for which each $\eta_j = \bar{\eta}$, and each element of Δ is independently chosen to be either 1 or 0 with probability p and $1-p$ respectively (i.e., a Erdos-Renyi random graph with sparsity coefficient p). For any fixed LNI j , which does not presynaptically-inhibit its own ORN, there will be approximately Np subnetworks of the form Δ_{jk} , Np^2 subnetworks of the form $\Delta_{jk} \Delta_{kj}$, Np^2 subnetworks of the form $\Delta_{jk} \Delta_{kk}$, $N^2 p^2$ subnetworks of the form $\Delta_{jk} \Delta_{jl}$, and $N^2 p^2$ subnetworks of the form $\Delta_{jk} \Delta_{kl}$ (where we assume k, j, l are distinct). If κ_ξ and κ_μ are modified for such a network so as to maintain m_j , then as κ_ξ is increased $\text{var}(ISI_j)$ will increase as well (since the reduction in $\text{var}(ISI_j)$ caused by the subnetworks of the form

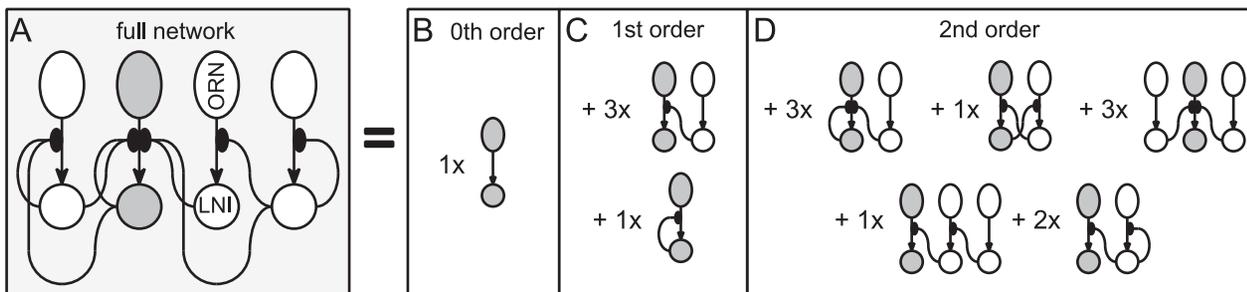


Figure 6. An example of subnetworks which come into play when considering the sensitivity or reliability of the j^{th} LNI. On the left in panel-A we show a particular network, with various ORN-LNI pairs (shown as ovals and circles respectively) connected via presynaptic-inhibitory connections. We will adopt the convention that the j^{th} ORN-LNI pair is fixed (highlighted in dark gray), whereas the indices k, l are not fixed, but are considered distinct from j and from each other. Several dynamic features associated with the j^{th} LNI can be determined by considering an expansion of the full network which are relevant for determining the sensitivity and reliability of the j^{th} LNI. The 0th-order subnetwork consists of the j^{th} ORN-LNI pair alone. The two 1st-order subnetworks shown are those incorporating a single presynaptic-inhibitory connection — namely Δ_{jk} (top) and Δ_{jj} (bottom). The full network has embedded within it three 1st-order subnetworks of the form Δ_{jk} , and one 1st-order subnetwork of the form Δ_{jj} . The five 2nd-order subnetworks shown are those incorporating two presynaptic-inhibitory connections. Listed in reading order, these subnetworks are denoted by $\Delta_{jj} \Delta_{jk}$, $\Delta_{jk} \Delta_{kj}$, $\Delta_{jk} \Delta_{jl}$, $\Delta_{jk} \Delta_{kl}$, and $\Delta_{jk} \Delta_{kk}$. The full network has embedded within it 3, 1, 3, 1 and 2 of these subnetworks, respectively. doi:10.1371/journal.pcbi.1002622.g006

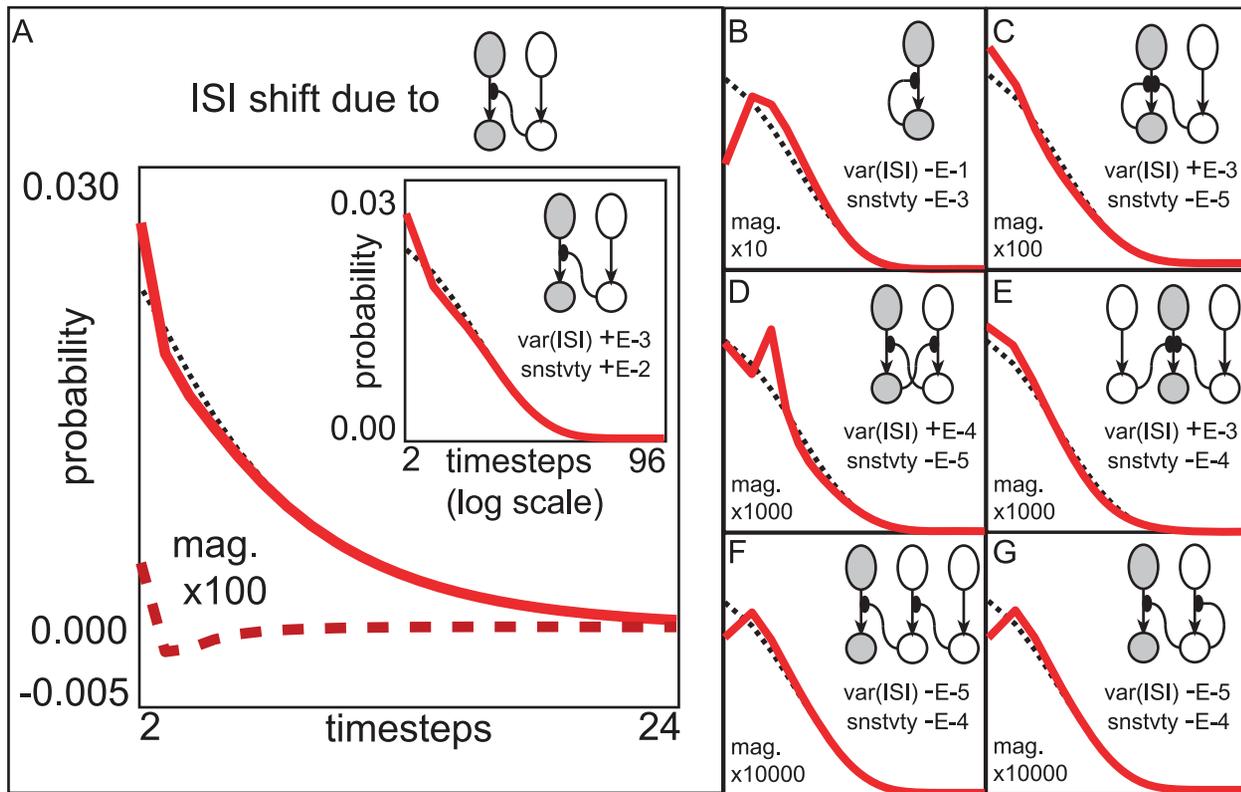


Figure 7. An example of subnetworks which influence reliability. In this example we assume a network of the form explained in the section entitled “A population-dynamics approach towards verifying Hypothesis 2 within more general networks”. The input to each LNI $\eta_j = \bar{\eta} = 0.5$ is constant, and the strength of vesicle-depletion $\kappa_\mu = 4$. Note however, that we do not assume that the connectivity Δ is fixed. We adopt the convention that j, k, l are distinct indices. [A] Here we illustrate the shift in the ISI-distribution of the j^{th} LNI (i.e., ISI_j) that would occur (up to 2nd-order) if the connectivity Δ_{jk} were increased while decreasing κ_μ so as to maintain the firing-rate of the j^{th} LNI (denoted by m_j). The ISI-distribution of the j^{th} LNI when uncoupled from the rest of the network is shown with a dotted-line for reference. The rate at which ISI_j changes with respect to an infinitesimal increase in the coupling strength Δ_{jk} is shown with a dashed-line. This rate is magnified by a factor of 100 for visibility. The sum of this rate and the uncoupled ISI_j is shown with a solid-line for a qualitative representation of the new ISI_j that would occur if the connectivity Δ_{jk} were increased by 1. The inset shows this same data (dotted and solid lines) with time plotted on a logarithmic scale for ease of view. For this particular term in the subnetwork-expansion, as Δ_{jk} increases (and the dotted ISI_j shifts to more closely resemble the solid ISI_j) the $\text{var}(ISI_j)$ increases. The rate at which $\text{var}(ISI_j)$ increases as Δ_{jk} is increased is approximately 10^{-3} for this system (as indicated by the legend ‘var(ISI)+E-3’). A separate calculation can be performed which shows that the rate at which the sensitivity of the j^{th} LNI (i.e., $\partial_{\eta_j} m_j$) changes as Δ_{jk} is increased is approximately 10^{-2} (as indicated by the legend ‘snstvtv+E-2’). Thus, by strengthening the presynaptic-inhibitory connections from several other LNIs onto the j^{th} ORN-LNI pair (while simultaneously reducing κ_μ so as to maintain m_j), we can readily show that, to 2nd-order, these shifts collectively increase both $\text{var}(ISI_j)$ and the sensitivity $\partial_{\eta_j} m_j$. [B–G] In these panels we show similar plots illustrating the influence of various other subnetworks on the reliability of the j^{th} LNI. These plots use axes identical to those shown on the inset in panel-A. Listed in reading order, these subnetworks are denoted by Δ_{jj} , $\Delta_{jj}\Delta_{jk}$, $\Delta_{jk}\Delta_{kj}$, $\Delta_{jk}\Delta_{jl}$, $\Delta_{jk}\Delta_{kl}$, and $\Delta_{jk}\Delta_{kk}$. Note that the contribution of the autapse Δ_{jj} actually decreases $\text{var}(ISI_j)$, and the contributions of the 2-edge subnetworks all decrease the sensitivity $\partial_{\eta_j} m_j$ (albeit with magnitudes that are dwarfed by the contribution of the 1-edge subnetwork Δ_{jk}). doi:10.1371/journal.pcbi.1002622.g007

$\Delta_{jk}\Delta_{kl}$ and $\Delta_{jk}\Delta_{kk}$ is more than cancelled out by the subnetworks of the form $\Delta_{jk}\Delta_{jl}$. If, on the other hand, the j^{th} LNI presynaptically-inhibits its own ORN, then in addition to the various subnetworks mentioned in the previous case, there will be a single subnetwork of the form Δ_{jj} , and approximately Np subnetworks of the form $\Delta_{jj}\Delta_{jk}$. If N, p are not sufficiently large, then the contribution to $\text{var}(ISI_j)$ will be dominated by the Δ_{jj} subnetwork, and increasing κ_ξ (while decreasing κ_μ to maintain m_j) will actually reduce $\text{var}(ISI_j)$. An example of the critical N below which hypothesis-2.1 fails is shown in Fig. 8, which illustrates that hypothesis-2.1 holds with high likelihood for all LNIs within Erdos-Renyi random networks obeying the dynamics specified in Eq. 16 (assuming that each $\eta_j = \bar{\eta}$), so long as N is sufficiently large.

This type of approach can also be used to analyze hypothesis-2.2. Using similar techniques as above, one can readily show that hypothesis-2.2 holds with high likelihood for all LNIs in an E-R random network of any size regardless of whether or not those LNIs presynaptically-inhibit their own ORNs, as long as their ORNs are presynaptically-inhibited by at least one other LNI.

Through the application of the above analysis, one can show that all sufficiently large E-R random networks obeying the dynamics of Eq. 16 have a high likelihood of satisfying the three properties associated with hypothesis-2 — namely, (0) there exists a 1-parameter family of network parameters ranging from type-A to type-B for which the firing rate of any particular LNI is maintained, (1) as the network parameters are shifted along this 1-parameter family in the type-A direction, that LNI’s reliability decreases, and

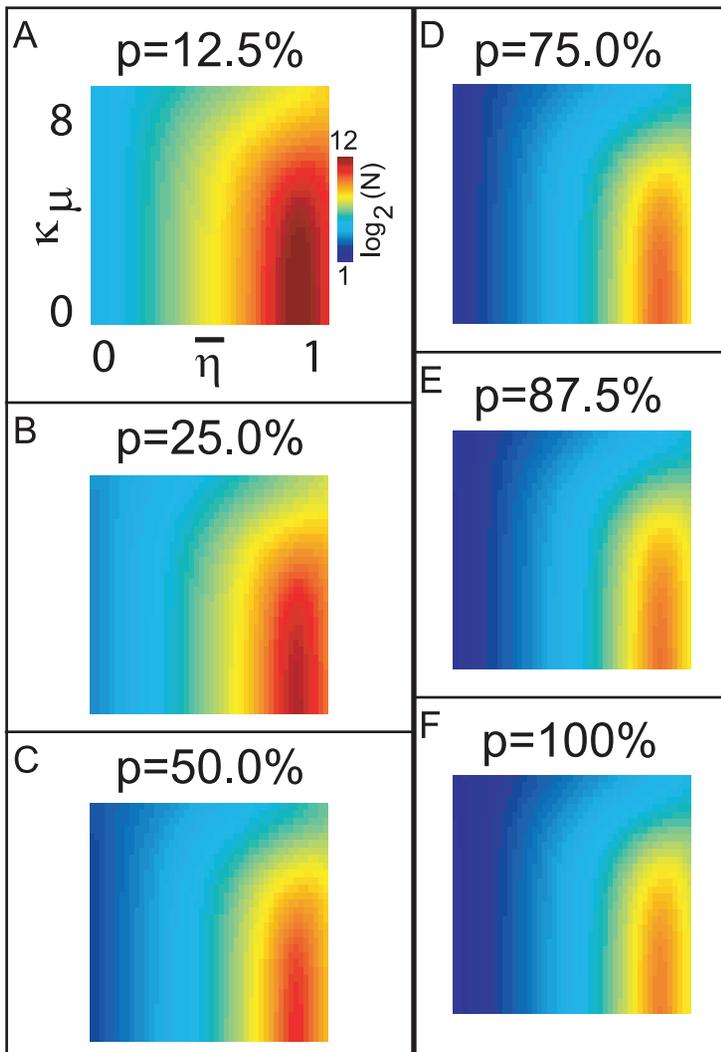


Figure 8. An analysis of sparsely-coupled Erdos-Renyi random networks, using a 2nd-order subnetwork-expansion. Given a network of the form described in the section entitled “A population-dynamics approach towards verifying Hypothesis 2 within more general networks”, with $\eta_j = \bar{\eta}$ and κ_μ fixed, one may ask if, for the j^{th} LNI, the reliability of this LNI would decrease if the presynaptic-inhibitory strength κ_ξ were to be increased (while simultaneously decreasing κ_μ so as to maintain the firing-rate m_j). Let us denote this condition by ‘hypothesis-2.1’. By analyzing the terms in the 2nd-order subnetwork-expansion, one can readily conclude that hypothesis-2.1 holds if the j^{th} LNI does not presynaptically-inhibit its own ORN, and there is at least one other LNI which does presynaptically-inhibit the j^{th} ORN. However, if the j^{th} LNI presynaptically-inhibits its own ORN, then hypothesis-2.1 holds only if the size of the network is sufficiently large. This critical network size N_{crit} (above which hypothesis-2.1 holds with high probability) is a function of the background firing-rate of the ORNs η_j , the strength of vesicle-depletion κ_μ , and the sparsity-coefficient p of the random network. In panel-A we plot $N_{\text{crit}}(\bar{\eta}, \kappa_\mu, p=0.125)$, where we have calculated N_{crit} such that, for values of $N \geq N_{\text{crit}}$, a randomly selected LNI within an E-R random network generated with sparsity-coefficient $p=0.125$ is highly likely (probability $> 75\%$) to obey hypothesis-2.1, given that the LNI in question presynaptically-inhibits its own ORN. Values of N_{crit} are displayed according to the colorscale shown on the right. In the remaining panels B–F we plot $N_{\text{crit}}(\bar{\eta}, \kappa_\mu, p)$ for different values of p . Note that, unless κ_μ is small and $\bar{\eta}$ is large, it is highly likely that hypothesis-2.1 holds (even for LNIs which presynaptically-inhibit their own ORNs) for all LNIs within an E-R random network of size ≥ 100 .
doi:10.1371/journal.pcbi.1002622.g008

(2) as the network parameters are shifted along this 1-parameter family in the type-A direction, that LNIs sensitivity to its dedicated ORN channel increases. For the class of networks considered in Fig. 8 the critical N above which hypothesis-2 holds is roughly $N_{\text{crit}} \sim 100$. Although this critical value obtained from our idealized model should not be compared quantitatively with the number of LNIs in the real fly AL, we expect that a similar qualitative result should hold for more realistic models — namely that hypothesis-2 should hold for any model as long as the presynaptic-inhibitory network includes sufficiently many inter-glomerular connections and the number of LNIs is sufficiently large.

This type of subnetwork analysis can also be used to probe the relationship between connectivity and dynamics that exists in more complicated heterogeneous networks, including scale-free and small-world networks, as well as networks in which the different neurons are governed by different equations. In each case, one can use the distribution of subnetworks within the larger network to determine dynamical features associated with neurons inside that network. For example, one can easily show that, for a scale-free network obeying the dynamics specified in Eq. 16, if the j^{th} ORN-LNI pair has low incoming-degree and is presynaptically-inhibited by ORN-LNI pairs with high incoming degree, then

the j^{th} LNI is likely to violate hypothesis-2.1. This is simply because there will be an overwhelming number of subnetworks of the form $\Delta_{jk}\Delta_{kl}$, which each contribute to the decrease of $\text{var}(ISI_j)$ when κ_ξ is scaled up (assuming as before that κ_μ is appropriately decreased to maintain the firing-rate m_j).

Finally, before we conclude, we address the effect of excitatory interconnectivity in the AL, which was not taken into account in the simple model analyzed above. Let's assume for a moment that, in the above model, we had added excitatory cells (say, LNEs) which were also affected by the ORN input channels. When analyzing the effect of modifying κ_ξ on the LNIs, the excitatory interactions associated with the LNEs only come into play at the 2nd-order. For example, if indices j,k correspond to LNIs and index l corresponds to an LNE, the simplest coupling terms by which LNE- l can affect LNI- j are of the form $\Delta_{jl}\Delta_{lk}$. Similar to the 2nd-order contributions shown in Fig. 7, The contributions of this term are dwarfed by the contributions of the LNI-LNI terms Δ_{jk} and Δ_{jj} . When analyzing the effects of modifying κ_ξ on the LNEs themselves, the situation is similar, with excitatory interactions playing a secondary role to the dominant contribution of first-order presynaptic-inhibitory connections. Thus, as long as the number of LNIs which affect each AL-neuron is not much less than the number of LNEs which affect that neuron (and as long as excitation and inhibition have effects of comparable magnitude), we expect the analytical results of this section to hold for networks which include both excitatory and inhibitory neurons.

Discussion

Using our large-scale model for the fly antennal lobe, we have been able to put forth two hypotheses which concern the functional role played by synaptic-depression at the ORN synapses within the fly AL. The first of these hypotheses is that ORN→AL synaptic-depression may contribute to a form of variance coding when ORN firing-rates are high. The second hypothesis is that two network mechanisms which both give rise to synaptic-depression (namely presynaptic-inhibition and vesicle-depletion) participate in establishing a balance between the network's reliability and sensitivity.

The first hypothesis hinges on the fact that, as ORN firing-rates increase beyond the value at which PN firing-rates saturate, the corresponding increase in ORN→PN synaptic-depression diminishes the postsynaptic impact of each individual ORN firing-event. In many situations (e.g., when many ORNs stimulate each PN, and different ORNs fire independently of one another), this scenario directly implies that, as ORN firing-rates increase, the ratio between the variance in input to the PNs and the mean input to the PNs decreases. If the ORN→PN synaptic-depression is sufficiently large, the variance in input to the PNs can actually decrease as ORN firing-rates increase, while the mean input to the PNs remains the same. When the variance to the PNs decreases sufficiently, the PNs behave as though they are driven by an input current. This current-like drive typically gives rise to very regular (periodically firing) PN dynamics. This first hypothesis is very general, and evidence of this hypothesis manifests in our simulations for a wide range of parameter values. As mentioned above, synaptic-depression (including vesicle-depletion and presynaptic-inhibition) is just one of several mechanisms which could contribute to this variance-coding effect.

One of the most straightforward predictions of our first hypothesis is that, within the real AL, the trial-to-trial variability of PN responses should decrease as ORN activity increases beyond the point where PN firing-rates have saturated. This should be measurable at the level of individual PN responses, as long as the

odor stimuli are chosen at the appropriate concentrations (see, e.g., [18]). Many of the experimental protocols which are currently employed are sufficient to test this prediction.

If such a variance code is indeed found in the fly AL, then the next question is whether or not the fly makes use of it. There are many ways in which an appropriately constructed downstream neuronal system could discriminate between two different stimuli that give rise to PN activity with similar firing-rates but differing degrees of regularity. For example, a downstream neuron may be wired to a subset of PNs in a given glomerulus in such a manner that (i) the steady nearly-periodic sequence of PN firing events associated with low-variance high-firing-rate PN activity (observed when m_{ORN} is very high) does not stimulate the downstream neuron sufficiently to induce firing, whereas (ii) the occasional clusters of PN firing-events which occur during high-variance high-firing-rate PN activity (observed when m_{ORN} is not quite so high) do stimulate the downstream neuron to fire. This is one of many potential 'readout' mechanisms that could serve to discriminate amongst stimuli which generate high ORN firing-rates. Such a readout mechanism may be at work within Kenyon cells, which have been observed to act like coincidence detectors [26–28]. A Kenyon cell that has a 'high-threshold' for firing may function as a coincidence detector, while also responding preferentially to spike-clusters within high-variance high-rate PN activity, but not to low-variance high-rate PN activity. An illustration of this principle is given in Fig. 2H. There are many other potential readout mechanisms that could serve to detect changes in the variance of PN activity, and effective readout mechanisms are likely nonlinear in nature [29,30].

Our first hypothesis may have significance for concentration-coding within the fly olfactory system. As an odor's concentration increases there are two typical changes to the ORN activity induced by that odor: (i) the number of activated ORN classes increases, and (ii) the ORN firing-rates within each activated class increase and eventually saturate [18,31,32]. It has been proposed that the firing-rates of the PN population across many glomeruli can collectively encode both stimulus identity and stimulus intensity [32–34], but it is still unclear how such a population firing-rate code could be deciphered by a downstream network within the fly. Our first hypothesis suggests an alternative to such a population firing-rate code, as the activity of a single glomerulus can also encode information regarding stimulus intensity, even when the glomerular firing-rates are saturated. Thus, a downstream decoder need not necessarily integrate information from multiple glomeruli in order to assess the intensity of a stimulus. In cases where a pair of high-concentration odors do in fact saturate the responses of most of the glomeruli which they target, typical readout mechanisms designed for rate-coding may not serve to discriminate these odors [35], and variance-coding mechanisms relying on transient bursts of activity (as mentioned above) may be more useful (see, for example, [36]).

The second hypothesis contains two statements. The first is that feedback-dominated synaptic-depression (i.e., presynaptic-inhibition) within type-A networks allows type-A networks to be more sensitive than type-B networks to shifts in input. Thus, type-A networks outperform type-B networks at odor-discrimination tasks when long observation-times are allowed. This statement seems to be true within our model network over a wide range of parameter values. The second statement is that feedforward-dominated synaptic-depression (i.e., vesicle-depletion) within type-B networks allows type-B networks to be more reliable than type-A networks (over multiple presentations of the same odor). This seems to be true within our large-scale model when (i) the typical number of ORN firing-events presynaptic to each PN within a glomerulus is

larger than the typical number of LNI firing-events presynaptic to the ORNs which target that glomerulus, (ii) the vesicle-depletion experienced by the ORN synapses is relatively well stereotyped across ORN firing-events, and (iii) the presynaptic-inhibitory network includes many ‘inter-glomerular’ connections. These three conditions seem to hold for the fly AL [12], but in other situations these conditions may be violated, and the converse of this second statement may hold. For example, we conjecture that in a hypothetical system containing (i’) very many high-firing-rate LNIs and PNs driven by only a few low-firing-rate ORNs and (ii’) a vesicle-depletion mechanism that is highly variable from one ORN firing-event to the next, it is likely that LNI→ORN presynaptic-inhibition can control the ORN input to the PNs much more reliably than vesicle-depletion of the ORN synapses could. In this hypothetical system type-B networks would be quite unreliable. Furthermore, even in a system with many ORNs and few LNIs, if (iii’) the presynaptic-inhibitory network is dominated by strong intra-glomerular connections (between LNIs and ORNs associated with the same odorant channel), then the subnetwork analysis carried out above indicates that type-A networks would actually be more reliable than type-B networks.

Obviously, the real AL for any given fly has a fixed architecture, and it does not seem likely that a fly could vary the effect of vesicle-depletion and presynaptic-inhibition on the ORN synapses to take advantage of a tradeoff between reliability and sensitivity. However, the hypothesis above does apply to the fly AL in the sense that fly physiology may be balanced to achieve some optimal compromise between gain (i.e., responsiveness), reliability, sensitivity and discriminability over short observation-times.

Suggestions for future study

The intuition gained in this study may be useful for understanding the coding properties of olfactory systems in other insects, or even in mammals, many of which also exhibit synaptic-depression.

In the olfactory system of many other insects, such as the locust and honeybee, the antennal lobe activity is characterized by oscillations which develop soon after the onset of stimulus. These oscillations are thought to be a key feature of the AL-response in these animals [26,33]. Since such oscillations do not manifest quickly in the fly [17], the dynamical regimes studied in this paper are not of this nature. However, we can retune our computational model to produce oscillations by increasing the density of lateral connectivity within the AL (thus, bringing our model closer in structure to that of [33]). The analytical techniques used in this paper may also be useful for studying some of the phenomena associated with such an oscillatory regime.

Within certain mammals, such as the mouse, primary olfactory input to the olfactory bulb can be presynaptically inhibited by interneurons [37–39]). Because the architecture of the mammalian olfactory system is different from that of the fly, the hypotheses investigated in this paper may not directly apply. For example, in the mouse it has been found that inhibitory neurons in the olfactory bulb strongly presynaptically-inhibit the olfactory sensory neurons stimulating their own glomerulus, but not those stimulating other glomeruli [40], a situation markedly different from the presynaptic-inhibitory network of the fly AL [12]. Thus, as hinted in (iii’) above, one might expect type-A networks to be both more sensitive *and* more reliable than type-B networks in the mouse. The extra coding power afforded by ‘feedback-induced’ synaptic-depression in this scenario may be necessary for an animal which is forced to sample its olfactory environment using short observation times (e.g., nose-pokes and sniffs). This sort of speculation begs for

a more detailed investigation of the structural and dynamic mechanisms at work in the mammalian olfactory system.

Materials and Methods

Overview of large-scale point-neuron model

In this section we describe the point-neuron model we used to investigate the dynamics of the fly AL. This model incorporates 5 glomerular channels, each with 6 PNs, 6 LNEs, 6 LNIs and 60 ORNs, and incorporates presynaptic-inhibition as well as vesicle-depletion (see Fig. 1). Each PN, LNE, LNI and ORN is modeled using single-compartment Hodgkin-Huxley type kinetics using standard sodium and potassium currents that give rise to fast sodium spikes [22] similar in shape to those observed experimentally [41].

Model synaptic currents. ORN, PN and LNE excitatory synapses and LNI GABAergic synapses are modeled by fast-activating synaptic currents [34,42]. Excitatory and GABAergic transmission are both modeled via stereotyped instantaneous neurotransmitter release in response to a presynaptic action potential, with ORN synapses experiencing vesicle-depletion. Although there is evidence for a slowly activating inhibitory current in the honeybee [43], we are not aware of any similar evidence for such slowly activating inhibitory current in the fly AL. There is, however, evidence for both long-timescale GABA-B type inhibitory currents, as well as short-timescale GABA-A type inhibitory currents in the fly AL [11]. To account for these two timescales, both a long- and short-timescale inhibitory current are incorporated into this model — the relevant synaptic currents include fast excitation (nACh-type, timescale 5–10 ms), fast inhibition (gabaA-type, timescale 10–15 ms) and slow inhibition (gabaB-type, timescale 100–400 ms).

Model connectivity. Experimental observations indicate that all the ORNs that express the same odorant receptor gene have similar odor responses and project to the same glomerulus in the brain [18,44]. Each PN receives direct ORN input from a single glomerulus [6], and thus all the ORNs and PNs corresponding to a given glomerulus constitute a discrete processing channel. We have designed the architecture of our model to reflect these experimental observations — our model ORNs and PNs only project to PNs and LNIs associated with their own glomerulus.

Experimental observations also indicate that glomeruli are interconnected by a network of local interneurons. There have been several experimental results implying that the lateral connectivity between glomeruli in the AL has a strong inhibitory component, and that this inhibitory component is partly due to presynaptic-inhibition (i.e., LNIs synapsing on the ORN axons presynaptically) [12,13,41]. It has also been revealed that ORN-PN synapses are quite strong, and likely experience substantial synaptic-depression through vesicle-depletion [11]. Our computational model reflects these observations, and model LNEs and LNIs project to neurons both within and outside their own glomerulus. In addition, the model LNIs only affect the ORNs presynaptically. That is, deposition of neurotransmitter from LNIs onto ORNs only suppresses the efficacy of the synapse at the ORN axon (without suppressing the membrane potential at the ORN soma). In order to model vesicle-depletion, the efficacy of the synapse associated with each ORN is depleted each time that ORN fires [22].

There is debate as to whether or not there is a functionally relevant large-scale structure to the lateral connections within the AL, such as center-surround excitation/inhibition, or chemotopy [7,45]. The lateral connectivity in our model network is structured so that the interconnections between ORNs, PNs, LNEs and LNIs

are sparse and randomly determined. Thus, the results of our model may be expected to generalize to a variety of networks with a degree-distribution similar to that of an Erdos-Renyi random-graph. The connectivity of the network is encoded by a matrix Δ , with $\Delta_{ij} \in \{0,1\}$ labelling the existence of a connection between cell j and cell i . Each entry in the connectivity matrix Δ is randomly chosen to be either 0 or 1 independently, with connection probabilities specific to the cell types and glomerular channel assignments of neurons i and j . Both the connection probabilities, as well as the coupling strengths are chosen in a manner consistent with the literature (details given in the sections regarding benchmarking below).

Model odor input. The model ORNs themselves are each stimulated by Poisson input. In background (i.e., in the absence of an odor stimulus) the Poisson input rate is 350 Hz, and the input strength is low enough that the ORN background firing-rate is 1–5 Hz, and the corresponding PN background firing-rate is 1–2 Hz (consistent with experimentally observed firing-rates — [6,41]). Odor presentation within this network is modeled by stimulating the ORNs corresponding to a subset of glomerular channels (typically around half of the glomerular channels) with additional high-rate Poisson input, in addition to the background Poisson input (consistent with experiments which indicate that an odorant typically activates multiple ORN types and triggers activity in $\sim 1/3$ to $1/2$ of the glomeruli — [18,20,46,47]). In this manner, a model odor is represented in a combinatorial fashion — an odor is defined by the degree to which that odor drives the various ORN input channels. To simulate odors of different chemical composition, we stimulate different subsets of ORN input channels, whereas to simulate odors of the same chemical composition (but of differing concentration) we stimulate the same input channels to differing degrees, as motivated by the observation that varying the concentration of a given odor tends to modulate the firing-rates of responding ORNs in vivo [45,48–50]. The odor-dependent noisy input signal is sufficient to drive individual ORNs strongly stimulated by the odor to firing-rates of 20–30 Hz. We choose a time-course of the odor-specific ORN current stimulus that is comparable to the time-course observed experimentally [18]. Specifically, at the time of ‘odor onset’ we increase the input-specific drive to the ORNs slowly (over ~ 400 ms), and at odor offset we decrease the input-specific ORN drive even more slowly (over ~ 1000 ms). While there is some evidence of more complicated odor-specific temporal structure to ORN odor response, we will model only the temporally simplistic ORN response detailed above, so as to focus on emergent dynamics within the AL which manifest solely as a result of AL interconnectivity.

Benchmarking the model

We have tuned the model so that, with a single set of parameters, the model exhibits a dynamic regime that is consistent with a variety of experimentally observed phenomena.

We attempted to ensure that the model network architecture is consistent with the literature. For example, motivated by [13], We chose the inhibitory postsynaptic coupling strengths from LNIs \rightarrow PNs so that the lateral inhibitory IPSC to a PN has both a fast and slow component (in our model IPSCs incorporate $\sim 50\%$ fast (gabaA type) and $\sim 50\%$ slow (gabaB type) inhibition). Similarly, we chose the inhibitory postsynaptic coupling strengths from LNIs \rightarrow LNEs and LNIs \rightarrow LNIs to be 100% fast-type. As another example, motivated by [41], we have chosen LNE \rightarrow PN intra-glomerular coupling to be sparse enough ($\sim 15\%$ - 25%) to align with the fact that direct LN \rightarrow PN connections are rarely observed. Nevertheless, LNE \rightarrow PN inter-glomerular coupling is

dense enough that lateral excitatory input is still observed between most pairs of glomeruli [8,9]. The lateral excitation between glomeruli is sufficiently strong that, even when ORNs belonging to a particular glomerulus are removed, some PNs and LNIs within that glomerulus can still fire after odor presentation.

Stimulus driven dynamic transients reflect PN reliability and saturation. We tuned the ORN \rightarrow PN connection probabilities and connection strengths in our model so as to produce stimulus-triggered dynamic transients which are qualitatively similar to experiment. This stimulated set of ORNs directly activates a corresponding subset of glomeruli. As the glomeruli respond to the ORN input, the glomerular activity pattern shifts and spreads to include other glomeruli not directly stimulated by the odor.

As shown in Fig. 9, the stimulus triggered firing-sequence of a typical model PN is more reliable than the corresponding sequence for a typical ORN associated with the same glomerulus, and the PN PSTH typically peaks after ~ 200 ms but before the ORN PSTH peaks (consistent with [20]). The sensitivity and reliability of PN activity within our model is critically dependent on three features: (i) the convergence ratio of ORNs to PNs must be sufficiently high that any given PN receives strong reliable input (summed over ORNs) immediately after odor onset, (ii) there must be some synaptic-depression at the ORN synapses, otherwise the PN activity does not peak before the ORN activity peaks (rather, the PN activity saturates and remains constant), and (iii) the combination of ORN activity and synaptic-depression at the ORN synapses must give rise to essentially ‘mean-driven’ (i.e., low-variance super-threshold) input to the PNs during odor presentation. These three network mechanisms seem to be necessary and sufficient to give rise to high-firing-rate PN activity which peaks prior to the peak in ORN activity, and which is more reliable than any single ORN’s activity. We note that the source of synaptic-depression is not constrained by this particular phenomena — either vesicle-depletion or presynaptic-inhibition (generated by the recruitment of LNIs in other glomeruli) can give rise to mean-driven ORN \rightarrow PN input, and hence to reliable PN firing sequences. As shown in Fig. 10, the PN firing-rate is a nonlinear function of ORN firing-rate when averaged across odors and trials and PN/ORN pairs within any given glomeruli, as is consistent with experiment [20].

PNs are more broadly responsive to odors than their ORN class would indicate. Experimentally, it has been observed that most PNs respond to multiple odors [45,51], and are more broadly responsive than their input ORNs [9,20] — there are many different PNs which respond to odors that do not stimulate their respective ORNs, and there are PNs which do not respond very strongly even when their respective ORNs are strongly stimulated. Thus, due to the lateral connectivity within the AL, the total set of activated glomeruli corresponding to a particular odor is generally not in one-to-one correspondence with the set of activated ORN olfactory receptors stimulated by that odor. Many believe that the inter-glomerular crosstalk is critically important for redistributing the glomerular activity within the AL ([41], but also see [7,45]). It has been postulated that this re-expression of the odor at the glomerular level is advantageous for the fly, as the combinatorial code linking different odors to different glomerular subsets serves to separate similar odors more efficiently than the corresponding combinatorial code at the olfactory receptor level [20].

We tuned the lateral connection probabilities and connection strengths within our model to produce a dynamic regime capable of rich glomerular activation patterns which are qualitatively similar to experiment. As shown in Fig. 11, the lateral connectivity

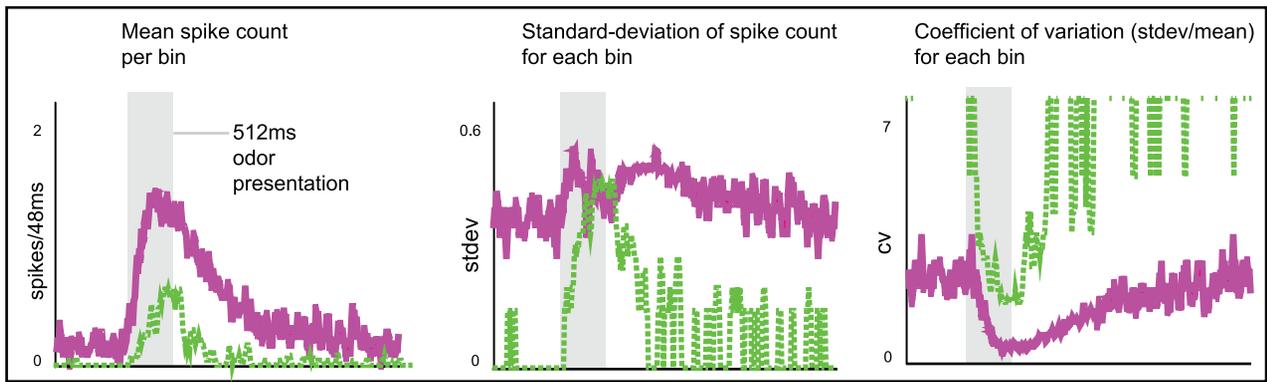


Figure 9. PNs are more reliable than their individual ORN inputs. Shown are averaged response curves for a typical model PN (magenta, solid) and model ORN (green, dashed) associated with the same glomerulus in our model. The grey overlay indicates the 512ms odor presentation period. Spikes were counted in 48ms bins. The mean spike-count per 48ms bin (averaged over 16 trials) is shown on the left. The standard-deviation in spike-count per 48ms bin is shown in the center, and the coefficient of variation (standard deviation \div mean) is shown on the right. Note that, qualitatively similar to experiment [20], the model PN activates more quickly, has higher firing-rates, and is more reliable than the ORN. doi:10.1371/journal.pcbi.1002622.g009

in our model is sparse enough that oscillations do not develop during the initial odor-response, and yet strong enough that the activity of PNs and LNs within each glomerulus is not in direct correspondence with the activity of their respective ORN inputs. This lack of correlation between PN response and ORN response across odors can be quantified by measuring the PN-PN and PN-ORN rank-correlation, which is in good qualitative agreement with experiment [20]. In our model, as in experiment, LN activity is similar to PN activity (data not shown). For example, LN firing-rate can get quite high, and typically peaks shortly after odor onset [13].

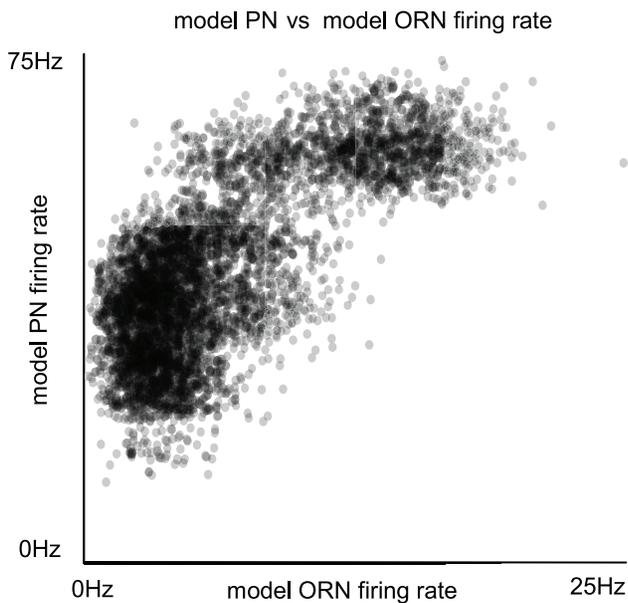


Figure 10. The relationship between ORN firing-rates and PN firing-rates is nonlinear. Shown is a scatterplot of model PN and model ORN firing-rates associated with a typical glomerulus in our model. Spike rates were measured during the 96ms epoch during which PN firing-rates peak following odor presentation. Note that, qualitatively similar to experiment [20], the model PN firing-rates saturate for relatively small values of ORN firing-rates. doi:10.1371/journal.pcbi.1002622.g010

ORN \rightarrow PN induced EPSCs attenuate as frequency of ORN synapse activation increases.

Our model of vesicle-depletion and presynaptic-inhibition is phenomenological, and is intended to allow us to qualitatively reproduce and investigate the functional role of synaptic-depression at the ORN synapses. We have chosen a parsimonious model for vesicle-depletion, involving only one timescale of $\sim 400ms$, and one parameter S^{vesdep} . Similarly, our model of presynaptic-inhibition involves only the timescales of synaptic inhibition and a coupling strength S^{presyn} (see the section regarding vesicle-depletion and presynaptic-inhibition below). In order to ensure that our phenomenological model of synaptic-depression at the ORN synapses (a combination of vesicle-depletion and presynaptic-inhibition) was qualitatively accurate, we followed the experimental paradigm of [11]. We constructed a numerical experiment to measure the attenuation timescale of ORN \rightarrow PN input (see Fig. 12). We first forced the ORN synapses to activate periodically at 8Hz (to simulate ORN background firing-rates), and then once the system equilibrated to this periodic input, we increased the frequency of ORN stimulation (to ωHz) and measured the input current to each PN as a function of time. For a larger given frequency ω , the attenuation of the PN EPSC will occur more quickly. We tuned the coupling strengths S^{presyn}, S^{vesdep} so that the attenuation time scale (as a function of ω) qualitatively matched experimental observations [11]. Our modeling work indicates that the attenuation time-scale match experiment as long as S^{presyn} and S^{vesdep} are sufficiently high (the exact ratio of S^{presyn} to S^{vesdep} is not strongly constrained by this particular experiment).

A further constraint on our model of synaptic-depression can be obtained by considering the correlation between total ORN activity, and reduction in subthreshold voltage observed at any given PN [12]. To obtain a roughly linear relationship between ORN activity and PN inhibition (as shown in Fig. 13), our model requires $S^{presyn} > 0$.

Details pertaining to Neuronal Model

The membrane potential of each ORN is governed by equations of the form

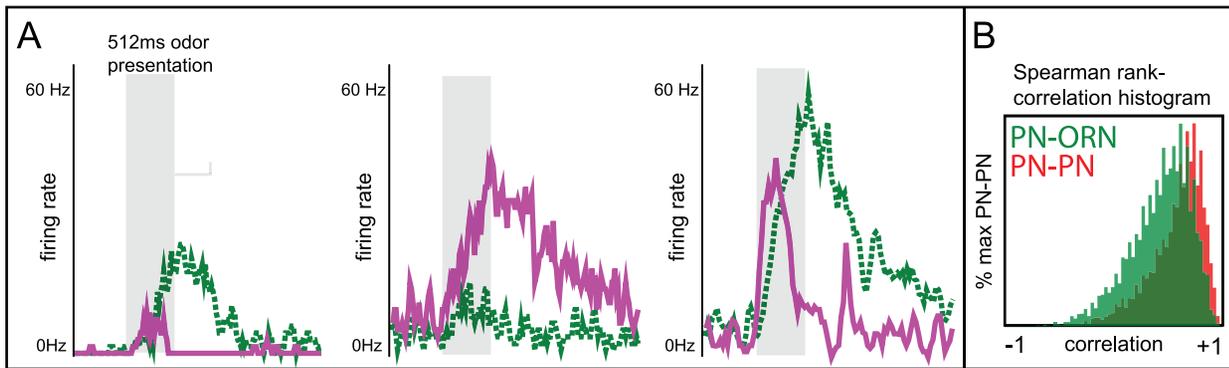


Figure 11. PNs exhibit broader odor responses than their associated ORNs. [A] Shown are trial-averaged firing-rate curves for various model PNs (magenta, solid) and associated model ORNs (green, dashed) in response to various model odors. Note that, qualitatively similar to experiment [12], the activity of the model PNs does not necessarily reflect the activity of the associated model ORNs. [B] Shown are the PN-ORN (green) and PN-PN (red) Spearman rank-correlation histograms for the model PNs and associated model ORNs (averaged over all PN and ORN pairs associated with each given glomerulus, and then further averaged over glomeruli — see [20] for the statistical methods used). Note that, qualitatively similar to experiment, the mean of the PN-ORN histogram is closer to 0 than the mean of the PN-PN histogram, indicating that, while PNs associated with a given glomerulus tend to respond to the same odors, they do not necessarily respond to the same set of odors which stimulate their associated ORNs.

doi:10.1371/journal.pcbi.1002622.g011

$$C_m \frac{dV}{dt} = -g_L(V - E_L) - I_{Na} - I_K - I^{stim},$$

with stimulus current described in a section entitled “Odor Stimulation” below. The membrane potential for each PN, LNE and LNI is governed by equations of the form:

$$C_m \frac{dV}{dt} = -g_L(V - E_L) - I_{Na} - I_K - I^{excitatory} - I^{GABA_A} - I^{GABA_B}.$$

The parameters for the passive leak current are $C_m = 1.0 \mu F$, $g_L = 0.3 \mu S$, $E_L = -64 mV$.

Intrinsic currents. The intrinsic currents for each neuron consist of fast sodium and potassium currents I_{Na} and I_K . These currents obey equations of the following form:

$$I_{Na} = g_{Na} m^3 h (V - E_{Na}), \quad I_K = g_K n (V - E_K)$$

The maximal conductances are $g_{Na} = 120 \mu S$, and $g_K = 3.6 \mu S$. The reversal potentials are $E_{Na} = 40 mV$, $E_K = -87 mV$.

The gating variables m, h, n take values between 0 and 1 and obey equations of the following form:

$$\frac{dm}{dt} = \alpha_m(V)(1 - m) - \beta_m(V)m, \quad \frac{dh}{dt} = \alpha_h(V)(1 - h) - \beta_h(V)h,$$

$$\frac{dn}{dt} = \alpha_n(V)(1 - n) - \beta_n(V)n$$

I_{Na} and I_K are described in [52]:

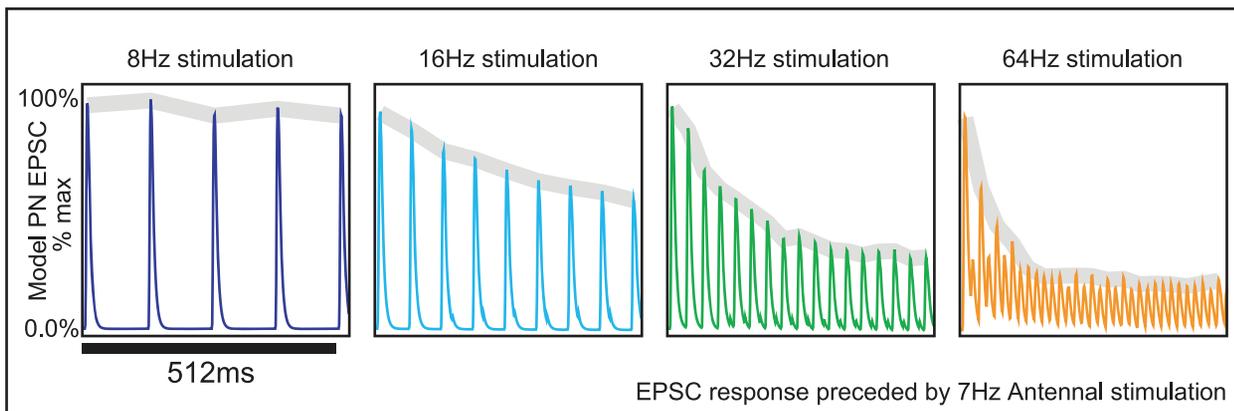


Figure 12. Synaptic-depression at the ORN synapses. Shown are current traces associated with a model PN in response to direct current stimulation of the model ORNs associated with that PN. Analogous to experiment [11], the model ORNs associated with the model PN have been stimulated by periodic 8Hz input current prior to the epoch shown in the figure. At the start of the epoch shown in this figure, the ORN stimulation is increased to 8Hz, 16Hz, 32Hz, or 64Hz. The trial-averaged model PN EPSCs in response these different stimulations are plotted (over a time interval of 550ms). Above each EPSC curve, we show the envelope of the response in gray. This envelope is calculated by fitting a piecewise linear function to the maxima of the EPSC response sampled at the rate of stimulation. Note that, similar to experiment, the envelope of the PN EPSC attenuates more quickly when stimulated at 64Hz than when stimulated at 8Hz.

doi:10.1371/journal.pcbi.1002622.g012

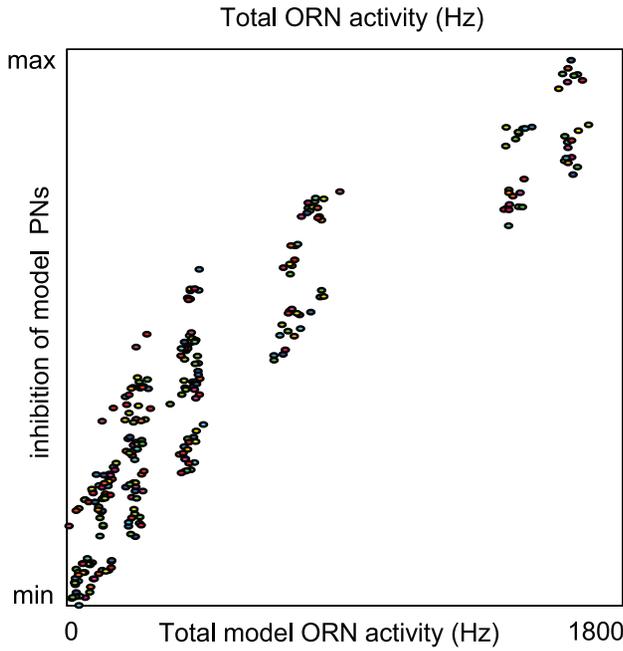


Figure 13. Presynaptic-inhibition is partly responsible for ORN→PN synaptic-depression. Shown is a scatterplot displaying the correlation between total ORN activity across all glomeruli in response to various odors, and the suppression of spontaneous EPSPs associated with a particular PN associated with a glomerulus which has been ‘shielded’ (i.e., the odor stimulus chosen does not affect the input drive to that glomerulus). In analogy with [12]. The PN suppression is measured as the difference in integrated PN membrane potential between (i) the scenario in which the PN receives spontaneous spikes from its associated ORNs in the absence of any odor, and (ii) the scenario in which the glomerulus associated with that PN is shielded and an odor is presented, in which case the activity generated within the other glomeruli reduce the effect of the spontaneous spikes impinging on the PN, and the spontaneous EPSPs are absent or greatly diminished. Note that, due to presynaptic-inhibition within the model, the correlation between PN EPSP magnitude and total ORN activity is qualitatively similar to experiment [12]. doi:10.1371/journal.pcbi.1002622.g013

$$\begin{aligned} \alpha_m &= 0.1[1 - \exp(-0.1(V + 35))]/[1 - \exp(-0.1(V + 35))] \\ \beta_m &= 4.0 \exp(-(V + 60)/18) \\ \alpha_h &= 0.07 \exp(-(V + 60)/20) \\ \beta_h &= 1.0/[1.0 + \exp(-0.1(V + 30))] \\ \alpha_n &= 0.01[1 - \exp(-0.1(V + 50))]/[1 - \exp(-0.1(V + 50))] \\ \beta_n &= 0.125 \exp(-(V + 60)/80). \end{aligned}$$

Synaptic currents. Given two connected neurons in the network, the synaptic conductances of the postsynaptic neuron are increased whenever the presynaptic neuron’s membrane potential rises above a threshold of $V_T = 0mV$ (i.e., when the intrinsic currents of the presynaptic neuron generate an action potential). The excitatory-synaptic current associated with the i^{th} neuron in the network is governed by an equation of the following form:

$$I_i^{excitatory} = \sum_j [S_{ij}^{excitatory} \Delta_{ij} O_j^{excitatory}] (V_i - E_{excitatory}), \quad (4)$$

$$I_i^{GABA_A} = \sum_j [S_{ij}^{GABA_A} \Delta_{ij} O_j^{GABA_A}] (V_i - E_{GABA_A}), \quad (5)$$

with the excitatory reversal potential $E_{excitatory} = 0mV$. The efficacy of the excitatory-type synapse $O_j^{excitatory}$ associated with the j^{th} neuron obeys the equation:

$$\begin{aligned} \frac{dO_j^{excitatory}(t)}{dt} &= -\frac{O_j^{excitatory}(t)}{\tau_{excitatory}} + \\ \sum_k A_{syn} \delta(t - [T_k^{spike,j} + \tau_{delay}]) &\left[1 - \lim_{t' \rightarrow t^-} O_j^{excitatory}(t')\right] \cdot d_j, \end{aligned} \quad (6)$$

where $\tau_{excitatory} = 5ms$, the times $T_k^{spike,j}$ are the spiketimes of the j^{th} neuron, $\tau_{delay} = 0.125ms$, and $A_{syn} = 1.5$ (adapted from [33,34,53,54]). The term d_j represents synaptic-depression at the ORN synapses, and is used to model both vesicle-depletion and presynaptic-inhibition at these synapses. This term d_j is identically 1 if j corresponds to a PN, LNE or LNI. The dynamics of d_j when j corresponds to an ORN will be discussed later. The differential equation Eq. 6 is structured so that $0 \leq O_j^{excitatory}(t) < 1$ for all time, and in the absence of firing $O_j^{excitatory}(t) \rightarrow 0^+$. The synaptic GABA-A current for the PNs, LNEs and LNIs obeys equations analogous to Eqs. 4,6, with $E_{GABA_A} = -70mV$. The efficacy of the GABA-A-type synapse $O_j^{GABA_A}$ associated with the j^{th} neuron obey the equation

$$\begin{aligned} \frac{dO_j^{GABA_A}(t)}{dt} &= -\frac{O_j^{GABA_A}(t)}{\tau_{GABA_A}} + \\ \sum_k A_{syn} \delta(t - [T_k^{spike,j} + \tau_{delay}]) &\left[1 - \lim_{t' \rightarrow t^-} O_j^{GABA_A}(t')\right], \end{aligned} \quad (7)$$

with $\tau_{GABA_A} = 15ms$. The synaptic GABA-B current for the PNs, LNEs and LNIs obeys an equation of the following form:

$$I_i^{GABA_B} = \sum_j [S_{ij}^{GABA_B} \Delta_{ij} \tilde{O}_j^{GABA_B}] (V_i - E_{GABA_B}),$$

with inhibitory reversal potential $E_{GABA_B} = -95mV$. The quantity $\tilde{O}_j^{GABA_B}$ for the j^{th} neuron obeys the equation:

$$\frac{d\tilde{O}_j^{GABA_B}(t)}{dt} = -\frac{\tilde{O}_j^{GABA_B}(t)}{\tilde{\tau}_{GABA_B}} + O_j^{GABA_B}(t),$$

with $\tilde{\tau}_{GABA_B} = 400ms$. The efficacy of the synapse $O_j^{GABA_B}$ for the j^{th} neuron obeys an equation analogous to Eq. 7, with $\tau_{GABA_B} = 30ms$. Thus, the slow GABA-B type synaptic current has a rise and decay time-scale, whereas the fast excitatory and GABA-A type synaptic currents only have decay time-scales (again adapted from [33,34,53,54]).

The synaptic coupling strengths $S_{ij}^{excitatory}$, $S_{ij}^{GABA_A}$, and $S_{ij}^{GABA_B}$ depend only on the cell types of neurons i and j , and are chosen so that only presynaptic ORNs, PNs and LNEs give rise to excitatory-type currents, and only LNIs give rise to GABA-A and GABA-B-type currents. The strengths are given by the following arrays:

| | excitatory fromORN | excitatory fromPN | excitatory fromLNE | GABA-A fromLNI | GABA-B fromLNI |
|-------|-----------------------|----------------------|-----------------------|----------------------------------|-----------------------------------|
| toORN | 0 | 0 | 0 | $S^{\text{presyn}} \cdot 0.0016$ | $S^{\text{presyn}} \cdot 0.00035$ |
| toPN | 0.052 | 0.001 | 0.025 | 0.0001 | 0.00005 |
| toLNE | 0.048 | 0.015 | 0.025 | 0.15 | 0 |
| toLNI | 0.130 | 0.001 | 0.030 | 0.0001 | 0.0001 |

The parameter S^{presyn} varies between 0 (no presynaptic-inhibition) and 10 (strong presynaptic-inhibition). Note that the PN→PN and LNI→PN coupling strengths are all negligible, to account for the experimental observations that PNs may not be targeted by other PNs, or by local inhibitory interneurons. However, we do allow for PNs to connect to LNEs (and LNIs), as observed in [41]. If these strengths are set to 0 we can retune the remaining connectivity strengths so that our major conclusions still hold (data not shown).

Network connectivity. The intra-glomerular connection probabilities are given by the array:

| | fromORN | fromPN | fromLNE | fromLNI |
|-------|---------|--------|---------|---------|
| toORN | 0.00 | 0.00 | 0.00 | 0.85 |
| toPN | 1.00 | 0.50 | 0.15 | 0.50 |
| toLNE | 0.75 | 0.50 | 0.15 | 0.50 |
| toLNI | 0.25 | 0.50 | 0.50 | 0.50 |

and the inter-glomerular connection probabilities Δ are given by the array:

| | fromORN | fromPN | fromLNE | fromLNI |
|-------|---------|--------|---------|---------|
| toORN | 0.00 | 0.00 | 0.00 | 0.85 |
| toPN | 0.00 | 0.00 | 0.75 | 0.00 |
| toLNE | 0.00 | 0.00 | 0.15 | 0.25 |
| toLNI | 0.00 | 0.00 | 0.50 | 0.50 |

Modeling vesicle-depletion and presynaptic-inhibition at the ORN→PN and ORN→LN synapses. The ORNs in this model do not directly experience synaptic conductances from either the PNs or the LNEs. The excitatory-type from PNs and LNEs onto ORNs is identically 0 (see the array shown in (8) above). The GABA-A and GABA-B-type conductances associated with an ORN (say, with index j) alter the term d_j in Eq. 6, thus affecting the efficacy of synapses from the j^{th} ORN onto the AL. The evolution of d_j in our model is governed by two parameters: S^{presyn} (described above), and S^{vesdep} (described below). These two parameters will allow us to consider a 2-parameter family of model networks in which the strength of vesicle-depletion and presynaptic-inhibition can be altered independently (see Fig. 4). The equations governing the evolution of d_j are

$$d_j = [1 - \mu_j] \exp \left[- \sum_k [S_{j,k}^{\text{GABA}_A} \Delta_{j,k} O_k^{\text{GABA}_A}] - \sum_k [S_{j,k}^{\text{GABA}_B} \Delta_{j,k} O_k^{\text{GABA}_B}] \right],$$

with the vesicle-depletion parameter μ_j obeying the differential equation:

$$\frac{d\mu_j(t)}{dt} = -\frac{\mu_j(t)}{\tau_\mu} + \sum_k S^{\text{vesdep}} \delta(t - T_k^{\text{spike},j}) \left[\lim_{t' \rightarrow t^-} \mu_j(t') - 1 \right],$$

with $\tau_\mu = 400ms$. The parameter S^{vesdep} varies from 0 (no vesicle-depletion, $\mu_j \equiv 0$) to 1.0 (complete vesicle-depletion with each firing-event). As the vesicle release rate per synaptic event is likely quite high within the real fly AL [11], values of $S^{\text{vesdep}} \in [0.5, 0.9]$ are most reasonable from a physiological standpoint. Note that the vesicle-depletion parameter μ_j is bounded between 0 and 1.

Odor simulation. The stimulus current to the j^{th} ORN is governed by the equation

$$I^{\text{stim}} = \sum_k S^{\text{bkg}} \alpha(t - T_{j,k}^{\text{bkg}}) (V_j - E_{\text{excitatory}}) + \sum_k S^{\text{odor}} \alpha(t - T_{j,k}^{\text{odor}}) (V_j - E_{\text{excitatory}}),$$

where $\alpha(t)$ is a response-function such that $\alpha(t) = 0$ if $t < 0$, and $\alpha(t) = \exp(-t/\tau_{\text{ORN}})$ if $t > 0$, with $\tau_{\text{ORN}} = 5ms$. The spiketimes $T_{j,k}^{\text{bkg}}$ are drawn from a Poisson-process with rate $\nu^{\text{bkg}} = 350Hz$. The strength of this background input is $S^{\text{bkg}} = 0.067$. The spiketimes $T_{j,k}^{\text{odor}}$ are drawn from a Poisson-process with rate $\nu_j^{\text{odor}}(t)$ which depends on the time since odor onset, the odor being presented, as well as the ORN under consideration. Typical values for ν_j^{orn} range from 0 (when the odor does not directly stimulate the j^{th} ORN) to 150Hz (when the j^{th} ORN is being strongly stimulated by the odor). The strength of this odor-specific input is $S^{\text{odor}} = 0.027$.

The time-dependence of the odor-specific input-rate $\nu_j^{\text{ORN}}(t)$ is governed by the factor

$$\nu_j^{\text{ORN}}(t) \propto [1 - \exp(-[t - t_{\text{onset}}]/\tau_{\text{rise}})] \cdot \exp(-[t - t_{\text{offset}}]/\tau_{\text{decay}}),$$

with t_{onset} and t_{offset} representing the onset and offset times of the odor stimulus (respectively), and with $\tau_{\text{rise}} = 400ms$, $\tau_{\text{decay}} = 1000ms$. A typical odor (stimulation of the ORNs within $\sim 1/2$ of the glomerular channels) activates $\sim 2/3$ of the PNs, which fire at about 50–150Hz. Not all of the PNs exhibit increased activity upon odor stimulation — a given odor will typically cause a few PNs which are not directly stimulated to actually decrease in activity (as a result of inter-glomerular inhibition). On presentation of a typical odor, the typical PN PSTH rises very quickly, and peaks after $\sim 200ms$, before the ORN PSTH peaks (at $\sim 500ms$ by construction). The typical PN PSTH decays more quickly than the ORN PSTH ($\sim 200 - 300ms$ vs $\sim 1000ms$). Due to lateral connectivity, PNs respond to a broader selection of odors than their ORN inputs and most PNs have a ‘rank order’ of odors that is different from the rank order of their ORN inputs. Note that PNs exhibit a rise in firing-rate after odor offset only very rarely. This is a consequence of the fact that the current model only incorporates odors which increase the firing-rate of the stimulated ORNs. It has been observed that some odors actually decrease the firing-rate of certain ORN types, with perhaps a resurgence of ORN firing-rates after odor offset (as shown in DL1 response to cis-3-hexen-1-ol, cyclohexane and ethyl acetate — [20]). These types of inhibitory Odor→ORN responses could potentially give rise to richer PN dynamics (potentially

triggered by PNs/LNs which fire greatly after odor offset), and these phenomena will be studied in more detail in future work.

Odor discrimination. In order to estimate the model network's ability to discriminate different odors, we measure the odor-dependent probability distribution of each PN's firing-rate, and use standard methods from classification theory [55]. For completeness we describe our procedure applied to firing-rate vectors.

Assume that we are estimating the network's ability to discriminate between 2 different odors. For each $i \in \{1, 2\}$, we perform multiple trials of odor i , and estimate the probability distribution

$$P_k^i(z) = P(\text{the } k^{\text{th}} \text{ PN fires } z \text{ times during a single trial of odor } i).$$

Once the P_k^i are sufficiently well estimated, we perform and classify individual odor trials. A single trial of odor i (randomly chosen to be either 1 or 2 with 50% probability) will give rise to a vector \vec{p} such that p_k is the number of firing-events produced by the k^{th} PN during that trial. By looking at a fixed k and comparing $P_k^1(p_k)$ and $P_k^2(p_k)$, we can use the k^{th} PN to identify a possible candidate stimulus (either odor 1 if $P_k^1(p_k) > P_k^2(p_k)$, or odor 2 if $P_k^1(p_k) < P_k^2(p_k)$). This process can be performed for each k , and in this way each PN 'votes' for a candidate stimulus. We tally these votes, weighting each one by the log of the information ratio associated with each PN. We use the weighting

$$\log(A_k/B_k),$$

where A_k is the 'hit-rate' associated with the k^{th} PN:

$$A_k = \frac{1}{2} \sum_{z=0}^{\infty} \max(P_k^1(z), P_k^2(z)),$$

and $B_k = 1 - A_k$ is the 'error-rate'. This particular weighting is chosen so that N votes for stimulus 1 with error-rate B have the same combined weight as a single vote for stimulus 2 with a far smaller error-rate of $B^N / [B^N + (1 - B)^N]$. The sum of the weighted votes is compared to determine the candidate stimulus underlying this particular trial. If the candidate stimulus matches the true stimulus, the trial is classified correctly. If not, the trial has been classified incorrectly. By going through this process with multiple trials, we can generate a probability that any given trial will be classified correctly. To perform N -way discriminability tasks, we go through an analogous procedure, performing all $\binom{N}{2}$ pairwise discriminability tasks for each sample observation, and ultimately selecting the candidate stimulus corresponding to the majority (with ties automatically counted as incorrect).

We have chosen this particular procedure because it allows us to take advantage of components of the firing-rate vector which carry substantial information (as measured by the information ratio), without requiring an estimate of the joint distribution of firing-rates (across PNs) for any particular odor. Thus, this measure of discriminability is more sensitive than typical linear discriminators which use the Euclidian distance between firing-rate vectors (see [20]), but does not succumb to the curse of dimensionality associated with the large number of distinct PNs.

An idealized model used to illustrate variance coding

Here we describe in detail the idealized model used in the section entitled "A simple cartoon of variance coding" in the main

text. This model includes a single conductance-based integrate-and-fire PN, driven by a set of 4 ORNs, each endowed with a simple model of synaptic-depression. Each of the 4 ORNs is modeled as a Poisson process with fixed rate η ($0\text{Hz} \leq \eta \leq 180\text{Hz}$). The coupling strength $S^{PN \leftarrow ORN}$ between the ORNs and the PN is modulated by a term $\mu(t)$ ($0 < \mu < 1$), which is intended to model vesicle-depletion at the ORN synapses. As each ORN fires, this μ term will give rise to synaptic-depression between the ORNs and the PN.

The membrane potential V and conductance G_E for this single PN obey the following differential equations:

$$dV/dt = -G_{LEAK}(V - V_R) - G_E(V - V_{EX})$$

$$dG_E/dt = -G_E/\tau_{G_E} + \sum_{j,k} S^{PN \leftarrow ORN} \left[1 - \lim_{t' \rightarrow t^-} \mu(t') \right] \delta(t - T_k^{ORNj}),$$

where $G_{LEAK} = 1/20\text{ms}$, $V_R = 0$ is the reset potential, $V_{EX} = 14/3$ is the excitatory reversal potential, and $\tau_{G_E} = 2\text{ms}$ is the conductance time-constant. The voltage V evolves continuously until V reaches a threshold $V = V_T = 1$, at which point the PN fires, and V is reset to V_R . The conductance G_E evolves continuously except when an ORN (say, the j^{th} ORN) spikes, at which point G_E jumps. The time T_k^{ORNj} is the k^{th} spiketime of the j^{th} ORN, and $S^{PN \leftarrow ORN} [1 - \lim_{t' \rightarrow t^-} \mu(t')]$ is the coupling strength associated with the j^{th} ORN at time t . If $\mu = 1$, the synapses between the j^{th} ORN and the PN are 100% exhausted. If $\mu = 0$, the synapses between the j^{th} ORN and the PN are completely refreshed. For this simple model $S^{PN \leftarrow ORN} = 0.75$, and the equation for $\mu(t)$ is given by:

$$\frac{d\mu(t)}{dt} = -\mu(t)/\tau_\mu + \sum_{j,k} \kappa_\mu \left[1 - \lim_{t' \rightarrow t^-} \mu(t') \right] \delta(t - T_k^{ORNj}),$$

where $\tau_\mu = 400\text{ms}$ is the time-constant associated with vesicle-depletion. The term μ decays to 0 continuously, except when the j^{th} ORN fires at some time T , at which point μ jumps by an amount proportional to $\kappa_\mu [1 - \lim_{t' \rightarrow T^-} \mu(t')]$ (the limit $\lim_{t' \rightarrow T^-} \mu(t')$ is used since $\mu(T)$ is not technically defined). The parameter $\kappa_\mu = 0.1325$ governs the relative increase in μ associated with each spike T_k^{ORNj} , and hence μ is bounded between 0 and 1. The parameters G_{LEAK} , τ_{G_E} and τ_μ are chosen to be consistent with typical point-neuronal models, and the parameter κ_μ is chosen essentially arbitrarily (different choices for κ_μ do not qualitatively change the results).

A simple analyzable cartoon of variance coding

As a simple cartoon which illustrates Hypothesis 1, consider a single PN modeled by a conductance-based integrate-and-fire neuron [22], driven by a single ORN modeled as a Poisson process (with firing rate η). The state variables of the PN are the membrane-potential V , the excitatory conductance G_E , and the vesicle-depletion parameter μ . The equations governing the state of the PN are

$$\begin{aligned}
\frac{dV(t)}{dt} &= -G_{LEAK}(V(t) - V_R) - G_E(t)(V(t) - V_{EX}) \\
\frac{dG_E(t)}{dt} &= -\frac{G_E(t)}{\tau_{G_E}} + \sum_k \frac{1}{\lim_{t' \rightarrow t^-} \mu(t')} \delta(t - T_k^{ORN}) \\
\frac{d\mu(t)}{dt} &= -\frac{\mu(t)}{\tau_\mu} + \sum_k \kappa_\mu \delta(t - T_k^{ORN}),
\end{aligned} \quad (11)$$

where G_{LEAK} is the leakage conductance, $V_R=0$ is the reset potential, $V_{EX}>1$ is the excitatory reversal potential, and τ_{G_E} is the conductance time-constant. The voltage V evolves continuously until V reaches a threshold $V=V_T=1$, at which point the PN fires, and V is reset to V_R . The conductance G_E decays to 0 continuously except when the ORN fires. The time T_k^{ORN} is the k^{th} spiketime of the ORN, and is produced by a Poisson process with rate η . The conductance G_E jumps by $\left[\lim_{t' \rightarrow [T_k^{ORN}]^-} \mu(t') \right]^{-1}$ at time T_k^{ORN} (the limit is used since $\mu(T_k^{ORN})$ is not technically defined). The vesicle-depletion parameter μ decays to 0 continuously with time-constant τ_μ , except when the ORN fires, at which point μ jumps by κ_μ . In this simple model the vesicle-depletion parameter μ and conductance G_E are both bounded to lie within $(0, \infty)$. As the vesicle-depletion parameter μ increases, the effect of ORN spikes on the PN conductance decreases. It should be noted that the functional form for the ORN \rightarrow PN synapse — in this case modeled by $1/\mu$ — is chosen to make the subsequent analysis easier, and is not particularly realistic (as very small values for μ imply a very strong ORN \rightarrow PN synapse). Nevertheless, the general picture implied by this cartoon holds for more realistic models of vesicle-depletion (see the section entitled “A simple cartoon of variance coding” in the main text).

The simple model Eq. 11 can be analyzed by considering the long-time evolution of the PN. For sufficiently small τ_{G_E} , τ_μ , κ_μ and sufficiently large η (with $\kappa_\mu \eta$ fixed), it can be shown [56] that the equilibrium-distribution of μ (collected by sampling over a sufficiently long time interval) is well-approximated by a Gaussian, with mean $\bar{\mu}$ and standard-deviation σ_μ given by:

$$\rho^{eq}(\mu) \propto \exp\left[-\frac{\mu - \bar{\mu}}{2\sigma_\mu^2}\right], \quad \bar{\mu} = \tau_\mu \kappa_\mu \eta, \quad \sigma_\mu^2 = \frac{\tau_\mu \kappa_\mu^2 \eta}{2}.$$

It can also be shown that, under these conditions, the equilibrium-distribution of G_E is also well-approximated by a Gaussian, with mean \bar{G}_E and standard-deviation σ_{G_E} given by:

$$\rho^{eq}(G_E) \propto \exp\left[-\frac{G_E - \bar{G}_E}{2\sigma_{G_E}^2}\right], \quad \bar{G}_E = \tau_{G_E} \frac{1}{\bar{\mu}}, \quad \sigma_{G_E}^2 = \frac{\tau_{G_E} \eta}{2\bar{\mu}^2}.$$

Using the expression for $\bar{\mu}$, the expressions for \bar{G}_E and σ_{G_E} can be simplified to

$$\bar{G}_E = \frac{\tau_{G_E}}{\tau_\mu \kappa_\mu}, \quad \sigma_{G_E}^2 = \frac{\tau_{G_E}}{2\tau_\mu^2 \kappa_\mu^2 \eta}.$$

Thus, for a sufficiently small κ_μ , as $\eta \rightarrow \infty$ the variance $\sigma_{G_E}^2$ of equilibrium-distribution $\rho^{eq}(G_E)$ shrinks to 0, and the mean \bar{G}_E

remains constant. Thus, as $\eta \rightarrow \infty$, the long-time conductance-distribution becomes sharply peaked around $\bar{G}_E = \tau_{G_E} / \tau_\mu \kappa_\mu$; so much so that, for sufficiently large η , the PN effectively has a fixed excitatory-conductance \bar{G}_E and will fire perfectly regularly with a period of

$$\Delta T = \frac{1}{G_L + \bar{G}_E} \log \left[\frac{(G_L V_R + \bar{G}_E V_{EX}) / (G_L + \bar{G}_E)}{(G_L V_R + \bar{G}_E V_{EX}) / (G_L + \bar{G}_E) - V_T} \right].$$

The excitatory conductance G_E is, in this case, independent of the activity of the PN because the ORN input is only affected by vesicle-depletion, and not by presynaptic-inhibition. Nevertheless, the conclusions we draw from this simple model are quite general, and will hold for more realistic models of synaptic-depression.

Note also that synaptic-depression is critical to hypothesis-1 within this model. If μ were fixed to be 1 (i.e., no synaptic-depression of the ORN synapses), and η were sufficiently large, then the equilibrium-distribution of G_E would be Gaussian with a mean and variance that grow unbounded as $\eta \rightarrow \infty$.

A simple model illustrating the tradeoff between reliability and sensitivity

In this section we analyze the model used in the section entitled “A simple analyzable cartoon of the tradeoff between reliability and sensitivity”.

To analyze the solutions of Eq. 1, let’s assume for the moment that $\mu=1$, and $\eta_A > \eta_B$ (i.e., neuron A fires more frequently than neuron B). If $\xi=0$, then A and B do not affect one another. The steady-state firing-rate m_A of neuron A is η_A , and the steady-state firing-rate m_B of neuron B is η_B . Let us define $\gamma_A = 1/\eta_A$ and $\gamma_B = 1/\eta_B$. Since both A and B are perfect phase-oscillators, A and B fire perfectly regularly every γ_A and γ_B time-units (respectively). If there is a difference in ORN inputs to these two neurons (say, $\eta_A - \eta_B = \Delta\eta > 0$), then the difference in firing-rates is $\Delta m = \Delta\eta$. Thus, if $\xi=0$, this system is perfectly reliable (in the sense that the ISI distribution of A and the ISI distribution of B both have 0 variance), and somewhat sensitive to shifts in the input (in the sense that any difference $\Delta\eta$ in input is reflected in the difference Δm of the output firing-rates).

If $\xi > 0$, then A and B affect one another with several consequences: (i) the steady-state firing-rates m_A and m_B will be lower than η_A and η_B (respectively), (ii) the ISI distributions of A and B will have nonzero variance, and (iii) the difference in steady-state firing-rates Δm will be greater than $\Delta\eta$. Indeed, as ξ increases away from 0 the system becomes less reliable while becoming more sensitive to shifts in the input. More specifically, for a given fixed $\xi > 0$, the system will settle down to a steady-state dynamics in which neuron A fires either $\lfloor (\gamma_B - \xi) / (\gamma_A - \xi) \rfloor$ or $\lfloor (\gamma_B - \xi) / (\gamma_A - \xi) \rfloor - 1$ times in between each pair of B -firing-events. The steady-state sequence of spike-times is independent of the initial state of the network and, while not generally periodic, can be solved for explicitly.

To show why this is true, we consider the return-maps f_A and f_B . We define the return-map $f_A(V_A)$ as follows: given a spike of neuron B (say, T_1^B), let T^A be the first spike of neuron A which occurs after T_1^B , and let T_2^B be the first spike of neuron B which occurs after T^A — we define $f_A(V_A(T_1^B)) = V_A(T_2^B)$. Similarly, given a spike T_1^A of neuron A , let T^B be the first spike of neuron B after T_1^A , and let T_2^A be the first spike of neuron A after T^B — we define $f_B(V_B(T_1^A)) = V_B(T_2^A)$. For the return map f_A , we can also define the numbers N_A and N_B as follows: N_B is the number of times B fires in between T_1^B and T^A , and N_A is the number of

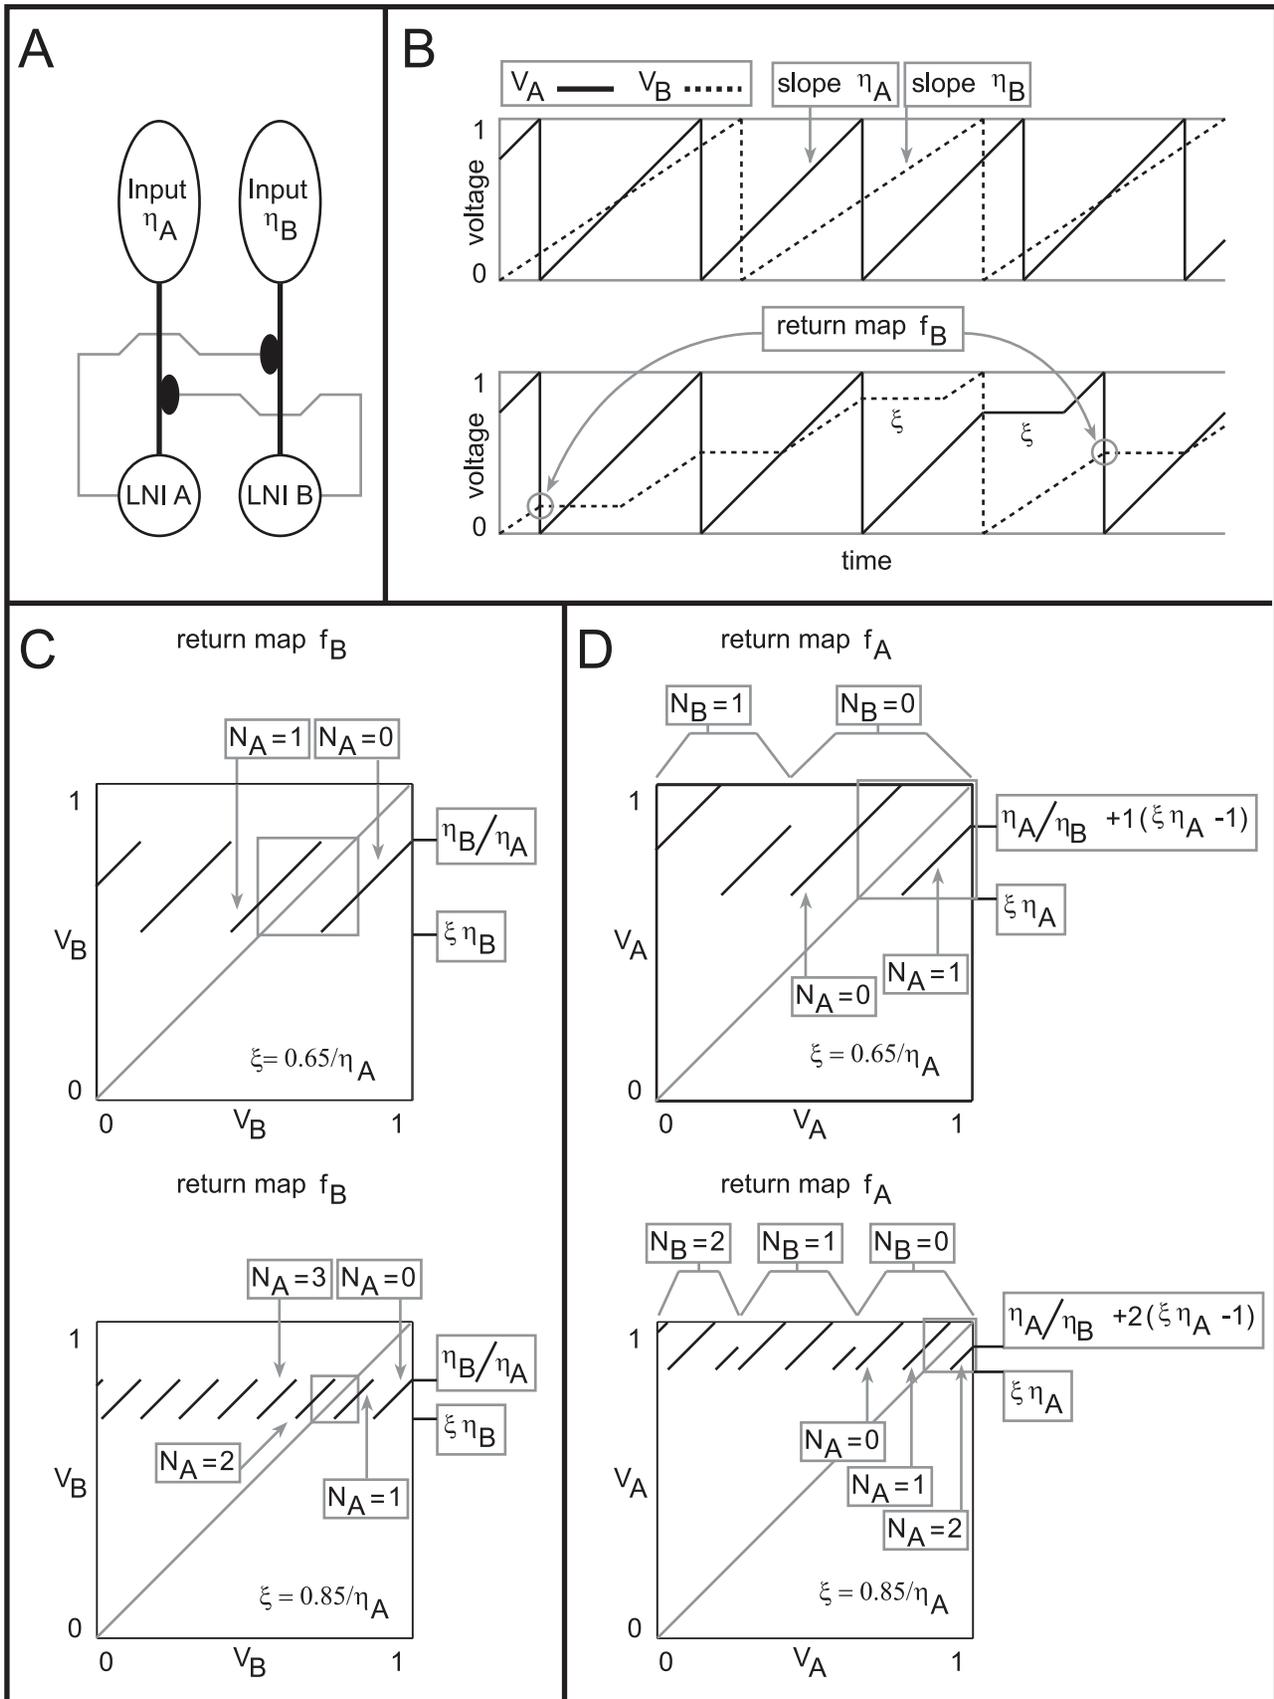


Figure 14. A simple analyzable cartoon of the tradeoff between reliability and sensitivity. [A] Shown is a schematic of the simple network, consisting of 2 ORN→LNI pairs, each of which presynaptically inhibits the other. [B] Shown on top are sample voltage-traces for the two LNIs (represented by V_A and V_B) for the case $\xi=0$. Shown on the bottom are sample voltage-traces for the two LNIs in the case that ξ is nonzero. Note that after LNI A fires, V_B is constant for ξ -time. Similarly, after LNI B fires V_A is constant for ξ -time. A pair of voltages for LNI B are circled. This pair of voltages $[V_B(T_1^A), V_B(T_2^A)]$ corresponds to a point on the graph of the return map f_B , namely $f_B(V_B(T_1^A)) = V_B(T_2^A)$. For this point on the graph of f_B , $N_A=2$, and $N_B=0$. [C] Shown on the top and bottom are return maps $f_B(V)$ for the values $\xi=0.65/\eta_A$, and $\xi=0.85/\eta_A$, respectively. [D] Shown on the top and bottom are return maps $f_A(V)$ for the values $\xi=0.65/\eta_A$, and $\xi=0.85/\eta_A$, respectively.
doi:10.1371/journal.pcbi.1002622.g014

times A fires in between T^A and T_2^B . Similarly, for the return map f_B , we can define N_A as the number of times A fires in between T_1^A and T^B , and N_B as the number of times B fires in between T^B and T_2^A . See Fig. 14 for an example of these return maps.

Recall that, without loss of generality, we have assumed $\eta_A > \eta_B$. By considering the return maps f_A and f_B , one can easily show that the maximum and minimum of f_A are 1 and $\xi\eta_A$, respectively, and that the maximum and minimum of f_B are η_B/η_A and $\xi\eta_B$, respectively. Moreover, f_A maps the interval $I_A = [\xi\eta_A, 1]$ into I_A . Similarly, f_B maps the interval $I_B = [\xi\eta_B, \eta_B/\eta_A]$ into I_B . It is straightforward to show that the image of I_A under f_A is composed of 2 sub-intervals: (i) an interval of length l_1^A for which $N_A = \lfloor (\gamma_B - \xi)/(\gamma_A - \xi) \rfloor$, and (ii) an interval of length $l_2^A = (1 - \xi\eta_A) - l_1^A$ for which $N_A = \lfloor (\gamma_B - \xi)/(\gamma_A - \xi) \rfloor - 1$. Similarly, the image of I_B under f_B is composed of 2 sub-intervals: (i) an interval of length l_1^B for which $N_A = \lfloor (\gamma_B - \xi)/(\gamma_A - \xi) \rfloor$, and (ii) an interval of length $l_2^B = (\eta_B/\eta_A - \xi\eta_B) - l_1^B$ for which $N_A = \lfloor (\gamma_B - \xi)/(\gamma_A - \xi) \rfloor - 1$. Letting

$$\bar{N} = \lfloor (\gamma_B - \xi)/(\gamma_A - \xi) \rfloor, \quad (12)$$

one can show that the lengths l_1^A and l_1^B are given by

$$\begin{aligned} l_1^A &= \eta_A/\eta_B - \xi\eta_A - \bar{N}(1 - \xi\eta_A) \\ l_1^B &= 1 - \eta_B(\gamma_A + (\bar{N} - 1)(\gamma_A - \xi)). \end{aligned} \quad (13)$$

For both the sub-maps $f_A : I_A \rightarrow I_A$ and $f_B : I_B \rightarrow I_B$ the number of extra spikes $N_B=0$. These observations allow us to conclude that the steady-state ISI distribution for neuron B (i.e., $ISI_B(\Delta T)$) has a peak of magnitude $l_1^A/(l_1^A + l_2^A)$ at $\Delta T = \gamma_B + (\bar{N} + 1)\xi$, and a peak of magnitude $l_2^A/(l_1^A + l_2^A)$ at $\Delta T = \gamma_B + \bar{N}\xi$. Similarly, the steady-state ISI distribution for neuron A (i.e., $ISI_A(\Delta T)$) has a peak of magnitude $(\bar{N}l_1^B + (\bar{N} - 1)l_2^B)/(\bar{N}l_1^B + (\bar{N} - 1)l_2^B + 1)$ at $\Delta T = \gamma_A$, and a peak of $1/(\bar{N}l_1^B + (\bar{N} - 1)l_2^B + 1)$ at $\Delta T = \gamma_A + \xi$. The steady-state firing-rates associated with these ISI-distributions can be expressed in closed form and directly computed:

$$\begin{aligned} m_A &= \left[\frac{\bar{N}l_1^B + (\bar{N} - 1)l_2^B}{\bar{N}l_1^B + (\bar{N} - 1)l_2^B + 1} [\gamma_A] + \frac{1}{\bar{N}l_1^B + (\bar{N} - 1)l_2^B + 1} [\gamma_A + \xi] \right]^{-1} \\ m_B &= \left[\frac{l_1^A}{l_1^A + l_2^A} [\gamma_B + (\bar{N} + 1)\xi] + \frac{l_2^A}{l_1^A + l_2^A} [\gamma_B + \bar{N}\xi] \right]^{-1}. \end{aligned} \quad (14)$$

Similarly, the means and variances associated with these steady-state ISI-distributions can be expressed in closed form (see Fig. 5).

By considering these expressions for small ξ , one can see that the steady-state return-map $f_A : I_A \rightarrow I_A$ consists of segments of length $(\eta_A/\eta_B - 1)$ and $(2 - \xi\eta_A - \eta_A/\eta_B)$, corresponding to inter-spike-intervals for neuron B of length $\gamma_B + 2\xi$ and $\gamma_B + \xi$, respectively. Thus, ISI_B has 2 distinct peaks (at $\gamma_B + \xi$ and $\gamma_B + 2\xi$, respectively), and as ξ increases the distance between

these two peaks increases. As a consequence, as ξ increases, the variance in ISI_B increases. In effect, a larger ξ implies that extra spikes from A have a larger effect on the ISI of B . A similar argument applies to $f_B : I_B \rightarrow I_B$ and ISI_A , and one can also show that, as ξ increase the difference in firing rates $\Delta m = m_A - m_B$ also increases (see Fig. 5). Thus, within this simple network, presynaptic-inhibition between the neurons disrupts their natural regularly-firing behavior, and increases the variance of their ISI distribution (thus decreasing their reliability).

In the discussion above, we assumed that $\mu=1$. If we assume $\mu \geq 1$, we can express the system firing-rate $\bar{m} = (m_A + m_B)/2$ in closed form as a function of μ, ξ (simply by replacing γ_A with $\mu\gamma_A$ and γ_B with $\mu\gamma_B$ in Eqs. 12,13,14). By requiring the system firing-rate to be constant, we can define μ implicitly (as a function of ξ, η_A, η_B and the system firing-rate \bar{m}). Thus, we can directly compute the 1-parameter family of networks which, for fixed η_A, η_B , attain a fixed system firing-rate \bar{m} . As shown in Fig. 5, this 1-parameter family of networks does indeed range from type-A networks (with high ξ and low μ) to type-B networks (with low ξ and high μ). Moreover, as one moves along this 1-parameter family of networks by increasing ξ (and decreasing μ appropriately), the variance in ISI_A and ISI_B increases, and the sensitivity $\Delta m = m_A - m_B$ also increases. In conclusion, this simple network illustrates that presynaptic-inhibition is capable of increasing the variance of the ISI distributions of the neurons within that network (hence reducing their reliability), while at the same time increasing the sensitivity of the neurons' firing-rates to subtle shifts in input.

Analysis of signal-to-noise ratio in a general discrimination task

In this section we provide details regarding the analysis in the section entitled "A simple cartoon of optimizing discriminability over short observation-times". Our goal is to determine from a measurement $\langle ISI_A(\mathcal{O}, \mu, \xi) \rangle_{T_{\text{obs}}} + X_E$ whether the input to the system is \mathcal{O}_1 or \mathcal{O}_2 . Let us denote by $\lambda_{\mathcal{O}}(\mu, \xi)$ and $\sigma_{\mathcal{O}}^2(\mu, \xi)$ the mean and variance of $ISI_A(\mathcal{O}, \mu, \xi)$. As long as T_{obs} is sufficiently large, the estimate $\langle ISI_A(\mathcal{O}, \mu, \xi) \rangle_{T_{\text{obs}}}$ can be considered to be drawn from $\mathcal{N}(\lambda_{\mathcal{O}}, \sigma_{\mathcal{O}}/\sqrt{T_{\text{obs}}})$. Thus, as long as T_{obs} is sufficiently large, the measurement $\langle ISI_A \rangle_{T_{\text{obs}}} + X_E$ can be considered to be drawn from $\mathcal{N}(\lambda_{\mathcal{O}}, \sqrt{\sigma_{\mathcal{O}}^2/T_{\text{obs}} + \sigma_E^2})$ (since $\mathcal{N}(0, \alpha) \star \mathcal{N}(0, \beta) = \mathcal{N}(0, \sqrt{\alpha^2 + \beta^2})$). If we attempt to discriminate between the two possible inputs by using a linear-classifier, then the error $E(\mu, \xi)$ associated with the best linear-classifier is simply given by the overlap of the distributions $\mathcal{N}(\lambda_{\mathcal{O}_1}, \sqrt{\sigma_{\mathcal{O}_1}^2/T_{\text{obs}} + \sigma_E^2})$ and $\mathcal{N}(\lambda_{\mathcal{O}_2}, \sqrt{\sigma_{\mathcal{O}_2}^2/T_{\text{obs}} + \sigma_E^2})$. Because $ISI_A(\mathcal{O}_1, \mu, \xi) = ISI_B(\mathcal{O}_2, \mu, \xi)$ and $ISI_A(\mathcal{O}_2, \mu, \xi) = ISI_B(\mathcal{O}_1, \mu, \xi)$ for this simple scenario, and the variance of these ISI distributions is very similar for $\xi < \gamma_A - \gamma_B/2$ (see Fig. 5), the error $E(T_{\text{obs}}, \mu, \xi)$ is well-approximated by

$$E(T_{\text{obs},\mu,\xi}) \approx 1 - \text{erf} \left(\frac{1}{2\sqrt{2}} \cdot \frac{\Delta\lambda(\mu,\xi)}{\sqrt{\sigma^2(\mu,\xi)/T_{\text{obs}} + \sigma_E^2}} \right),$$

where $\Delta\lambda(\mu,\xi)$ is the difference in the means of $ISI_A(\mathcal{O}_1,\mu,\xi)$ and $ISI_B(\mathcal{O}_1,\mu,\xi)$, and $\sigma^2(\mu,\xi)$ is the average variance of $ISI_A(\mathcal{O}_1,\mu,\xi)$ and $ISI_B(\mathcal{O}_1,\mu,\xi)$.

A discrete state model used to analyze hypothesis-2 within general networks with arbitrary architecture

In this section we describe the point-neuron model used in the section entitled ‘‘A population-dynamics approach towards verifying Hypothesis 2 within more general networks’’. This model is a stripped down version of the fly AL, consisting of N discrete-state LNIs, each driven by a different ORN. We will model each ORN-LNI pair as a discrete-state discrete-time Markov process which is as simple as possible, while still retaining the following features: (i) each LNI generates spikes, (ii) each ORN input spike contributes to the vesicle-depletion of that ORN→LNI synapse, and (iii), each LNI spike gives rise to presynaptic-inhibition of ORN→LNI synapses. We will model the j^{th} ORN-LNI pair using the state-variables $V_j(t)$, $\mu_j(t)$ and $\xi_j(t)$ which represent LNI membrane-potential, ORN vesicle-depletion and ORN presynaptic-inhibition, respectively.

At each discrete time, each state variable is either 0 or 1, thus, at each time, the j^{th} ORN-LNI pair is in one of $s=8$ states. The input from the j^{th} ORN to the j^{th} LNI is modeled as a bernoulli-random-variable $I_j(t)$, which is 1 with probability η_j (and 0 otherwise). The state-variables undergo transitions of the following form:

$$P(\xi_j(t+1)=x|\xi_j(t)=y) = [L_j^\xi]_{x,y},$$

$$P(\mu_j(t+1)=x|\mu_j(t)=y) = [L_j^\mu]_{x,y},$$

$$P(V_j(t+1)=x|V_j(t)=y) = [L_j^V]_{x,y}, \text{ where}$$

$$L_j^\mu(t) = (1 - I_j(t)) \begin{bmatrix} 1 & 1/(1 + \tau_\mu) \\ 0 & \tau_\mu/(1 + \tau_\mu) \end{bmatrix} + I_j(t) \begin{bmatrix} 1 - I(\kappa_\mu) & (1 - I(\kappa_\mu))/(1 + \tau_\mu) \\ I(\kappa_\mu) & (\tau_\mu + I(\kappa_\mu))/(1 + \tau_\mu) \end{bmatrix},$$

$$L_j^\xi(t) = \begin{bmatrix} 1 - I(\kappa_\xi \sum_k \Delta_{jk} V_k(t)) & (1 - I(\kappa_\xi \sum_k \Delta_{jk} V_k(t)))/(1 + \tau_\xi) \\ I(\kappa_\xi \sum_k \Delta_{jk} V_k(t)) & (\tau_\xi + I(\kappa_\xi \sum_k \Delta_{jk} V_k(t)))/(1 + \tau_\xi) \end{bmatrix}, \quad (15)$$

$$L_j^V(t) = (1 - I_j(t)) \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} +$$

$$I_j(t) \begin{bmatrix} (1 + \mu_j(t) + \xi_j(t))/(2 + \mu_j(t) + \xi_j(t)) & 1 \\ 1/(2 + \mu_j(t) + \xi_j(t)) & 0 \end{bmatrix},$$

and the function $I(x)$ is the logistic function

$$I(x) = 1/(1 + \exp(-x)).$$

For this system τ_μ and τ_ξ are the typical persistence times of the $\mu_j=1$ and $\xi_j=1$ states (respectively), and we will assume that $\tau_\mu = \tau_\xi = 1$. The $V_j=1$ state is considered a ‘firing’ state for LNI j . If $V_j(t)=1$, then $V_j(t+1)$ always equals 0. If $V_j(t)=0$ and the j^{th}

LNI does not receive input (i.e., $I_j(t)=0$), then $V_j(t+1)=0$. However, if $V_j(t)=0$ and $I_j(t)=1$, then V_j may transition to the firing state. Given that $I_j(t)=1$, the probability of V_j transitioning from 0 to the firing-state is typically 1/2, but is lowered if either $\mu_j=1$ or $\xi_j=1$. The vesicle-depletion parameter μ_j is likely to transition to the $\mu=1$ state whenever the j^{th} LNI receives input (i.e., $I_j=1$). The presynaptic-inhibition parameter ξ_j is likely to transition to the $\xi=1$ state whenever many other LNIs in the network fire. Note that the connectivity matrix Δ_{jk} encodes the connectivity of the network, and can be chosen to encode many different network architectures (e.g., a densely connected homogeneous network, or a sparsely connected heterogeneous network). If Δ_{jk} is nonzero, then the k^{th} LNI presynaptically-inhibits the j^{th} ORN, making it more likely that $\xi_j=1$, and thus less likely that ORN input from the j^{th} ORN to the j^{th} LNI will cause the j^{th} LNI to fire. For this model $\kappa_\mu \cdot \kappa_\xi$ are the overall strengths of vesicle-depletion and presynaptic-inhibition. As κ_μ increases, the likelihood of μ_j transitioning to the 1 state increases. Similarly, as κ_ξ increases, the likelihood of ξ_j transitioning to the 1 state increases as long as $\Delta_{jk} > 0$ for some k . Ultimately, we will assume that the probability that the j^{th} neuron will transition from the state $\omega_y = \{V_j(t)=y_V, \mu_j(t)=y_\mu, \xi_j(t)=y_\xi\}$ at time t to the state $\omega_x = \{V_j(t+1)=x_V, \mu_j(t+1)=x_\mu, \xi_j(t+1)=x_\xi\}$ at time $t+1$ is given by

$$[L_j]_{\omega_y, \omega_x} = [L_j^V]_{x_V, y_V} [L_j^\mu]_{x_\mu, y_\mu} [L_j^\xi]_{x_\xi, y_\xi}. \quad (16)$$

Note that L_j is an $s \times s$ state-transition matrix which depends on the state of the k^{th} neuron in the system as long as $\Delta_{jk} > 0$.

Analysis of reliability and sensitivity using a subnetwork expansion

This section reviews a diagrammatic approach to analyzing network dynamics, and presents the salient calculations relevant to analyzing ISI-distribution and firing-rate (which can then be used to analyze reliability and sensitivity, respectively). One way to understand the equilibrium dynamics of a network such as Eq. 15 is to first picture the network as a point in phase-space Ω , with the network’s current state determined by the collection of parameters

$$\omega = \{\omega_1, \omega_2, \dots, \omega_N\} = \{V_1, \mu_1, \xi_1, \dots, V_N, \mu_N, \xi_N\}$$

at the current time, where we denote by ω_j the state of the j^{th} ORN-LNI pair. As time passes this network will trace out a trajectory $\omega(t)$ in phase-space, and this trajectory will depend on the network’s architecture (i.e., η_j , Δ_{jk} , κ_μ , κ_ξ). If one could determine the ‘typical’ phase-trajectories exhibited by this network (over very long times) then, in particular, one could determine this network’s spike-time reliability. If one could determine how this network’s typical trajectories shift as the input η_j changes, then one could determine this network’s sensitivity. The typical trajectories of a network can be determined by considering both the evolution-operator of the network L , and the frequency ρ with which the network visits each part of phase-space (i.e., the network’s equilibrium-distribution). The full evolution-operator $L_{\hat{\omega}, \omega}$ is the probability that the network moves from state ω to state $\hat{\omega}$ over one timestep. In this case L is an $N^s \times N^s$ matrix such that each entry has the form

$$[L]_{\omega,\omega} = \prod_{j=1,\dots,N} [L_j]_{\omega_j,\omega_j}.$$

The probability that the network will be in state ω_θ at time $t_0 + \theta$, given that the network was in state ω_0 at time t_0 is

$$P(\omega_\theta, t_0 + \theta | \omega_0, t_0) = \sum_{\omega_1, \dots, \omega_{\theta-1}} \prod_{j=0}^{\theta-1} L_{\omega_j, \omega_{j+1}}. \quad (17)$$

Eq. 17 can be thought of as an integral over all possible paths in state-space connecting ω_0 to ω_θ (i.e., each path traverses the system-states $\omega_0, \omega_1, \dots, \omega_\theta$ in sequence). The equilibrium-distribution ρ is an eigenfunction of L with eigenvalue 1 — namely

$$\rho = L\rho,$$

and in this case ρ is an $N^s \times 1$ matrix (i.e., an N^s -dimensional vector). In this discussion we will assume that ρ is unique (i.e., L only has a single equilibrium-distribution). Note that both L and ρ are functions of the network's architecture.

Given both L and ρ , one can determine many properties of the network's equilibrium dynamics. For example, the probability that the j^{th} neuron fires at any given time (i.e., the steady-state firing-rate of the j^{th} neuron) is given by

$$m_j = f_j \cdot \rho,$$

where f_j is a $1 \times N^s$ operator such that $[f_j]_\omega = 0$, except for states in which $V_j = 1$, in which case $[f_j]_\omega = 1$ (i.e., $[f_j]_\omega = \delta_{V_j,1}$). Similarly, the probability that the j^{th} neuron fires at times t and $t + \theta$, without firing at any intermediate times (denoted by $[j \rightarrow \theta j]$) is given by

$$[j \rightarrow \theta j] = f_j \cdot L \cdot \left((I - F_j^T) \cdot L \right)^{\theta-1} \cdot F_j^T \cdot \rho,$$

where F_j^T is a $N^s \times N^s$ operator such that $[F_j^T]_{\omega,\omega} = \delta_{V_j,1} \delta_{V_j,1}$. The sensitivity of the network's firing-rates can be calculated via the N^2 derivatives $\partial_{\eta_j} m_k$. The reliability $\text{var}(ISI_j)$ of the j^{th} neuron in the network can be characterized by calculating the variance of $[j \rightarrow \theta j]$ (considered as a distribution with respect to θ).

Ideally, one might wish to determine how dynamic sensitivity (i.e., $\partial_{\eta_j} m_k$) and reliability (i.e., $\text{var}(ISI_j)$) vary as functions of a network's architecture. Unfortunately, the explicit functional dependence of m_j and $[j \rightarrow \theta j]$ on architectural parameters (such as $\eta, \Delta, \kappa_\mu, \kappa_\xi$) cannot be directly determined for most typical networks. However, it is possible to approximate these quantities by considering a weak-coupling expansion of L and ρ in terms of κ_ξ .

If $\kappa_\xi = 0$, then each ORN-LNI pair is independent from the rest of the network, and the full state-evolution operator $L^{[0]}$ can be constructed by taking an operator-direct-product of the various L_j , where L_j is the $s \times s$ -dimensional state-evolution operator associated with the j^{th} ORN-LNI pair shown in Eq. 16 (note that if $\kappa_\xi = 0$, then each L_j is independent of all other neurons). Similarly, if $\kappa_\xi = 0$ the full equilibrium-distribution $\rho^{[0]}$ can be constructed by taking the product of the various ρ_j , where ρ_j is the s -dimensional

equilibrium-distribution of the j^{th} ORN-LNI pair (note that ρ_j is the eigenvector of L_j with eigenvalue 1). Both L_j and ρ_j only depend on η_j and κ_μ . The sensitivity and reliability for this uncoupled $\kappa_\xi = 0$ network can be determined simply by computing the sensitivity and reliability for individual (uncoupled) ORN-LNI pairs.

If κ_ξ is small, then the network's full state-evolution operator L is no longer a direct product of the L_j (and ρ is no longer a product of the ρ_j). Nevertheless, by taking a Taylor-expansion of L in terms of κ_ξ (around $\kappa_\xi = 0$) one can approximate L and ρ via a series

$$L \approx L^{[0]} + \kappa_\xi L^{[1]} + \kappa_\xi^2 L^{[2]} + \dots$$

$$\rho \approx \rho^{[0]} + \kappa_\xi \rho^{[1]} + \kappa_\xi^2 \rho^{[2]} + \dots$$

It can be shown that the M^{th} -order terms in these series (corresponding to $L^{[M]}$ and $\rho^{[M]}$) incorporate subnetworks of the original network spanning up to $M + 1$ ORN-LNI pairs [24,25]. Specifically, the 0^{th} -order terms capture the equilibrium dynamics of each single ORN-LNI pair in the absence of the rest of the network. The $M = 1$ terms capture the first-order corrections associated with a single presynaptic-inhibitory connection of the form $\kappa_\xi \Delta_{jk}$. The $M = 2$ terms capture both the second-order corrections associated with a single presynaptic-inhibitory connection (of the form $\kappa_\xi^2 \Delta_{jk}^2$), as well as the second-order corrections associated with 2 presynaptic-inhibitory connections (of the form $\kappa_\xi^2 \Delta_{jk} \Delta_{j'k'}$). In the main text (Fig. 7), we have grouped the 2^{nd} -order terms corresponding to only 1 presynaptic-inhibitory connection with the 1^{st} -order terms associated with that connection. For example, when presenting the term associated with the subnetwork Δ_{jj} , we implicitly include both the 1^{st} -order term proportional to Δ_{jj} , as well as the 2^{nd} -order term proportional to Δ_{jj}^2 . When presenting the term associated with the subnetwork Δ_{jk} , we implicitly include both the 1^{st} -order term proportional to Δ_{jk} , as well as the 2^{nd} -order term proportional to Δ_{jk}^2 .

Using the series-expansion for L and ρ , one can compute a series-expansion for many quantities of interest (e.g., firing rate m_k , autocorrelation $[j \rightarrow \theta j]$, sensitivity $\partial_{\eta_j} m_k$ or reliability $\text{var}(ISI_j)$) in terms of subnetworks of the original network. One attractive feature of this approach is that the formal series-expansion can be constructed without specifying the connectivity matrix Δ . The terms of the series expansion can then be analyzed to determine which connectivity matrices will give rise to various dynamic phenomena. As an example, the terms $L^{[1]}$ and $L^{[2]}$ in the series expansion for L can be written as:

$$L^{[1]} = \sum_{a,b \text{ distinct}} \Delta_{ab} \left[L'_a \otimes F_b^R L_b F_b^T \otimes_{z \neq a,b} L_z \right] + \sum_a \Delta_{aa} \left[F_a^R L'_a F_a^T \otimes_{z \neq a} L_z \right],$$

$$\begin{aligned}
L^{[2]} = & \sum_a \Delta_{aa}^2 \left[F_a^R \frac{L_a''}{2} F_a^T \otimes_{z \neq a} L_a \right] \\
& + \sum_{a,b \text{ distinct}} 2\Delta_{ab}\Delta_{aa} \left[F_a^R \frac{L_a''}{2} F_a^T \otimes F_b^R L_b F_b^T \otimes_{z \neq a,b} L_z \right] \\
& + \sum_{a,b \text{ distinct}} \Delta_{ba}\Delta_{aa} \left[F_a^R L_a' F_a^T \otimes L_b' \otimes_{z \neq a,b} L_z \right] \\
& + \sum_{a,b \text{ distinct}} \Delta_{bb}\Delta_{aa} \left[F_a^R L_a' F_a^T \otimes F_b^R L_b' F_b^T \otimes_{z \neq a,b} L_z \right] \\
& + \sum_{a,b \text{ distinct}} \Delta_{ab}\Delta_{ab} \left[\frac{L_a''}{2} \otimes F_b^R L_b F_b^T \otimes_{z \neq a,b} L_z \right] \\
& + \sum_{a,b \text{ distinct}} \Delta_{ba}\Delta_{ab} \left[F_a^R L_a' F_a^T \otimes F_b^R L_b' F_b^T \otimes_{z \neq a,b} L_z \right],
\end{aligned}$$

References

- Axel R (1995) The molecular logic of smell. *Sci Am* 273: 154–159.
- Hildebrand J, Shepherd G (1997) Mechanisms of olfactory discrimination: converging evidence for common principles across phyla. *Annu Rev Neurosci* 20: 595–631.
- Vosshall L, Wong A, Axel R (2000) An olfactory sensory map in the fly brain. *Cell* 102: 147–159.
- Gao Q, Yuan Y, Chess A (2000) Convergent projections of *Drosophila* olfactory neurons to specific glomeruli in the antennal lobe. *Nat Neurosci* 8: 780–785.
- Treloar H, Feinstein P, Mombaerts P, Greer C (2002) Specificity of glomerular targeting by olfactory sensory axons. *J Neurosci* 22: 2469–2477.
- Couto A, Alenius M, Dickson B (2005) Molecular, anatomical, and functional organization of the *Drosophila* olfactory system. *Curr Biol* 15: 1535–1576.
- Root C, Semmelhack J, Wong A, Flores J, Wang J (2007) Propagation of olfactory information in *drosophila*. *Proc Natl Acad Sci U S A* 104: 11826–11831.
- Olsen S, Bhandawat V, Wilson R (2007) Excitatory interactions between olfactory processing channels in the *Drosophila* antennal lobe. *Neuron* 54: 89–103.
- Shang Y, Claridge-Chang A, Sjulson L, Pypaert M, Miesenbock G (2007) Excitatory local circuits and their implications for olfactory processing in the fly antennal lobe. *Cell* 128: 431–432.
- Stocker R, Lienhard M, Borst A, Fischbach K (1990) Neuronal architecture of the antennal lobe in *drosophila melanogaster*. *Cell Tissue Res* 262: 9–34.
- Kazama H, Wilson R (2008) Homeostatic matching and nonlinear amplification at identified central synapses. *Neuron* 58: 401–413.
- Olsen S, Wilson R (2008) Lateral presynaptic-inhibition mediates gain control in an olfactory circuit. *Nature* 452: 956–960.
- Wilson R, Laurent G (2005) Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the *Drosophila* antennal lobe. *J Neurosci* 25: 9069–9079.
- Vickers N, Baker T (1996) Latencies of behavioral response to interception of filaments of sex pheromone and clean air influence flight track shape in *heliophilis virescens* males. *J Comp Physiol A* 178: 831–847.
- Bhandawat V, Maimon G, Dickinson M, Wilson R (2010) Olfactory modulation of flight in *drosophila* is sensitive, selective and rapid. *J Exp Biol* 213: 3625–3635.
- Lei H, Riffell J, Gage S, Hildebrand J (2009) Contrast enhancement of stimulus intermittency in a primary olfactory network and its behavioral significance. *J Bio* 8:21.
- Tanaka N, Ito K, Stopfer M (2009) Odor-evoked neural oscillations in *drosophila* are mediated by widely branching interneurons. *J Neurosci* 29: 8595–8603.
- Halle M, Ho M, Carlson J (2004) The molecular basis of odor coding in the *drosophila* antenna. *Cell* 117: 965–979.
- Nagel K, Wilson R (2011) Biophysical mechanisms underlying olfactory receptor neuron dynamics. *Nature Neurosci* 14: 208–218.
- Bhandawat V, Olsen S, Schlieff M, Gouwens N, Wilson R (2007) Sensory processing in the *drosophila* antennal lobe increases the reliability and separability of ensemble odor representations. *Nat Neurosci* 10: 1474–1482.
- Olsen S, Bhandawat V, Wilson R (2010) Divisive normalization in olfactory population codes. *Neuron* 66: 287–299.
- Dayan P, Abbott L (2001) *Theoretical Neuroscience*. Cambridge, MA: MIT press.
- Masse N, Turner G, Jefferis G (2009) Olfactory information processing in *drosophila*. *Curr Biol* 19: R700–R713.
- Rangan A (2009) Diagrammatic expansion of pulse-coupled network dynamics in terms of subnetworks. *Phys Rev E* 80: 036101.
- Rangan A (2009) Diagrammatic expansion of pulse-coupled network dynamics. *Phys Rev Lett* 102: 158101.
- Perez-Orive J, Bazhenov M, Laurent G (2004) Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *J Neurosci* 24: 6037–6047.
- Wang Y, Guo H, Pologruto T, Hannan F, Hakker I, et al. (2004) Stereotyped odor-evoked activity in the mushroom body of *Drosophila* revealed by green fluorescent protein-based Ca^{2+} imaging. *J Neurosci* 24: 6507–6514.
- Turner G, Bazhenov M, Laurent G (2008) Olfactory representations by *drosophila* mushroom body neurons. *J Neurophys* 99: 734–746.
- Shamir M, Sompolinsky H (2004) Nonlinear population codes. *Neural Comput* 16: 1105–36.
- Barak O, Tsodyks M (2006) Recognition by variance: learning rules for spatiotemporal patterns. *Neural Comput* 18: 2343–58.
- Sachse S, Galizia G (2003) The coding of odour-intensity in the honeybee antennal lobe: local computation optimizes odour representation. *Eur J Neurosci* 18: 2119–2132.
- Ito I, Bazhenov M, Ong R, Raman B, Stopfer M (2009) Frequency transitions in odor-evoked neural oscillations. *Neuron* 64: 629–706.
- Bazhenov M, Stopfer M, Rabinovich M, Abarbanel H, Sejnowski T, et al. (2001) Model of cellular and network mechanisms for odor-evoked temporal patterning in the locust antennal lobe. *Neuron* 30: 569–581.
- Bazhenov M, Stopfer M, Rabinovich M, Huerta R, Abarbanel H, et al. (2001) Model of transient oscillatory synchronization in the locust antennal lobe. *Neuron* 30: 553–567.
- Luo S, Axel R, Abbot L (2010) Generating sparse and selective third-order responses in the olfactory system of the fly. *Proc Natl Acad Sci U S A* 107: 10713–8.
- Lei H, Christensen T, Hildebrand J (2002) Local inhibition modulates odor-evoked synchronization of glomerulus-specific output neurons. *Nat Neurosci* 5: 557–565.
- Wachowiak M, Cohen L (1999) Presynaptic inhibition of primary olfactory afferents mediated by different mechanisms in lobster and turtle. *J Neurosci* 19: 8808–8817.
- Wachowiak M, McGann J, Heyward P, Shao Z, Puche A, et al. (2005) Inhibition of olfactory receptor neuron input to olfactory bulb glomeruli mediated by suppression of presynaptic calcium influx. *J Neurophysiol* 94: 2700–2712.
- Pirez N, Wachowiak M (2008) In vivo modulation of sensory input to the olfactory bulb by tonic and activity-dependent presynaptic inhibition of receptor neurons. *J Neurosci* 28: 6360–6371.
- McGann J, Pirez N, Gainey M, Muratore C, Elias A, et al. (2005) Odorant representations are modulated by intra- but not interglomerular presynaptic inhibition of olfactory sensory neurons. *Neuron* 48: 1039–1053.
- Wilson R, Turner G, Laurent G (2004) Transformation of olfactory representations in the *Drosophila* antennal lobe. *Science* 303: 366–370.
- Patel M, Rangan A, Cai D (2009) A large scale model of the locust antennal lobe. *J Comp Neurosci* in press.
- Barbara G, Zube C, Rybak J, Gauthier M, Grunewald B (2005) Acetylcholine, GABA and glutamate induce ionic currents in cultured antennal lobe neurons of the honeybee, *Apis mellifera*. *J Comp Physiol A* 191: 823–836.

where we use the notation that L_j is equal to the $s \times s$ single-neuron operator shown in Eq. 16, with $\kappa_\xi \equiv 0$, and L_j' is the derivative of L_j with respect to $\kappa_\xi \sum_k \Delta_{jk} V_k(t)$ (the coupling parameter which appears in L_j^ξ). We also use the operator

$$\left[F_j^R \right]_{\omega, \omega} = \delta_{V_j, 0} \delta_{V_j, 0, a} N^s \times N^s \text{ operator.}$$

In the above representation of $L^{[1]}$ and $L^{[2]}$, we use \otimes to denote an operator-direct-product, and \otimes to denote an accumulation of operator-direct-products (analogous to the use of ‘+’ and ‘ Σ ’ respectively).

Author Contributions

Conceived and designed the experiments: AVR. Performed the experiments: AVR. Analyzed the data: AVR. Contributed reagents/materials/analysis tools: AVR. Wrote the paper: AVR.

44. Clyne P, Grant A, O'Connell R, Carlson J (1997) Odorant response of individual sensilla on the drosophila antenna. *Invert Neurosci* 3: 127–135.
45. Wang J, Wong A, Flores J, Vosshall L, Axel R (2003) Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* 112: 271–282.
46. Mombaerts P (2004) Genes and ligands for odorant, vomeronasal and taste receptors. *Nat Rev Neurosci* 5: 263–278.
47. Bargmann C (2006) Comparative chemosensation from receptors to ecology. *Nature* 444: 295–301.
48. Friedrich R, Korsching S (1997) Combinatorial and chemotopic odorant coding in the zebrafish olfactory bulb visualized by optical imaging. *Neuron* 18: 737–752.
49. Meister M, Bonhoeffer T (2001) Tuning and topography in an odor map on the rat olfactory bulb. *J Neurosci* 21: 1351–1360.
50. deBruyne K, Foster M, Carlson J (2001) Odor coding in the drosophila antenna. *Neuron* 30: 537–552.
51. Ng M, Roorda R, Lima S, Zemelman B, Morcillo P, et al. (2002) Transmission of olfactory information between three populations of neurons in the antennal lobe of the fly. *Neuron* 36: 463–474.
52. Hodgkin A, Huxley A (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 117: 500–544.
53. Destexhe A, Bal T, McCormick D, Sejnowski T (1996) Ionic mechanisms underlying synchronized oscillations and propagating waves in a model of ferret thalamic slices. *J Neurophysiol* 76: 2049–2070.
54. Bazhenov M, Timofeev I, Steriade M, Sejnowski T (1998) Cellular and network models for intrathalamic augmenting responses during 10hz stimulation. *J Neurophysiol* 79: 2730–2748.
55. Rangan A, Cai D, McLaughlin D (2008) Quantifying neuronal network dynamics through coarsegrained event trees. *Proc Nat Acad Sci U S A* 105: 10990.
56. Rangan A, Cai D (2006) Maximum-entropy closures for kinetic theories of neuronal network dynamics. *Phys Rev Lett* 96: 178101.